

Lemma 3.11. *If $\tilde{d}_X(\cdot, \cdot)$ is either Hausdorff distance or max-distance, then $d_X(x_1, x_2) \leq \tilde{d}_X(X_1, X_2)$ for all $X_1, X_2 \subseteq X$, $x_1 \in X_1$, $x_2 \in X_2$, such that $x_1 \overset{X_2}{\sim} x_2$.*

Proof. Let $\tilde{d}_X(\cdot, \cdot)$ be either Hausdorff distance or max-distance.

- For max-distance, $\tilde{d}_X(X_1, X_2) = \max_{x_1 \in X_1, x_2 \in X_2} d_X(x_1, x_2)$, so $d_X(x_1, x_2) \leq \tilde{d}_X(X_1, X_2)$ holds for all $x_1 \in X_1$, $x_2 \in X_2$.
- For Hausdorff distance, $\tilde{d}_X(X_1, X_2) = \max_{i \in \{1,2\}, j=3-i} \{ \max_{x_i \in X_i, x_j \in X_j, x_i \overset{X_j}{\sim} x_j} d_X(x_i, x_j) \}$, and hence $d_X(x_1, x_2) \leq \tilde{d}_X(X_1, X_2)$ holds for all x_1 and x_2 such that $x_1 \overset{X_2}{\sim} x_2$. \square

Let us henceforth assume that $\tilde{d}_X(\cdot, \cdot)$ is either Hausdorff distance or max-distance. The distance $\tilde{d}_Y(\cdot, \cdot)$ can in general be arbitrary, and we will mention when some particular instantiation is needed.

Theorem 3.12. *If $\tilde{d}_Y(\cdot, \cdot)$ is Hausdorff distance, then A-sensitivity is equivalent to C-sensitivity, and B-sensitivity is equivalent to D-sensitivity.*

Proof. Let ρ be a component of X , and $X_1, X_2 \in X/\rho$. Let $f : X \rightarrow Y$ be c_I I-sensitive, where $I \in \{A, B, C, D\}$.

\Leftarrow Let $x_1 \in X_1$, $x_2 \in X_2$. At least for max-distance and Hausdorff distance, we have $d_X(x_1, x_2) = \tilde{d}_X(\{x_1\}, \{x_2\})$ and $d_Y(f(x_1), f(x_2)) = \tilde{d}_Y(f(\{x_1\}), f(\{x_2\}))$. Hence, we may always estimate $(\{x_1\}, \{x_2\})$ instead of (x_1, x_2) . Thus $c_A \leq c_C$ and $c_B \leq c_D$.

\Rightarrow Let $X'_1 \overset{X_2}{\sim} X'_2$. The A- and B-sensitivities give us bounds only for the elements that are friends in X_2 , but $\tilde{d}_Y(f(X'_1), f(X'_2))$ may also depend on non-friends of X'_1 and X'_2 , since they may become friends after f is applied to them. We will show that the bounds c_A and c_B of A- and B-sensitivities are still valid. Let $P := \{(x_1, x_2) \mid x_1 \in X'_1, x_2 \in X'_2, x_1 \overset{X_2}{\sim} x_2\}$. We are able to apply A- and B-sensitivity to all $(x_1, x_2) \in P$.

Since $X'_1 \overset{X_2}{\sim} X'_2$, the first coordinates of pairs of P cover the entire set X'_1 , and the second coordinates of P cover the entire set X'_2 , possibly with repetitions. By Lemma 3.9, for all $(x_1, x_2) \in P$, we have $x_1 \overset{X'_2}{\sim} x_2$, and hence, by Lemma 3.11, we have $d_X(x_1, x_2) \leq \tilde{d}_X(X'_1, X'_2)$. Applying Lemma 3.11 directly to $x_1 \overset{X_2}{\sim} x_2$, we also have $d_X(x_1, x_2) \leq \tilde{d}_X(X_1, X_2)$. There are now two cases possible:

- If at least one $(x_1, x_2) \in P$ is such, that $\tilde{d}_Y(f(X'_1), f(X'_2)) \leq d_Y(f(x_1), f(x_2))$, then we get a valid bound $\tilde{d}_Y(f(X'_1), f(X'_2)) \leq d_Y(f(x_1), f(x_2)) \leq c_A \cdot d_X(x_1, x_2) \leq c_A \cdot \tilde{d}_X(X'_1, X'_2)$, implying $c_C \leq c_A$. Similarly, we have $\tilde{d}_Y(f(X'_1), f(X'_2)) \leq d_Y(f(x_1), f(x_2)) \leq c_B \cdot \tilde{d}_X(X_1, X_2)$ with the implication $c_D \leq c_B$.
- Suppose that for all $(x_1, x_2) \in P$ we have $\tilde{d}_Y(f(X'_1), f(X'_2)) > d_Y(f(x_1), f(x_2))$. Assuming that all maximums exist, and Hausdorff distance is used, let us take x_1 and x_2 such that $f(x_1) \overset{f(X'_2)}{\sim} f(x_2)$ and $\tilde{d}_Y(f(X'_1), f(X'_2)) = d_Y(f(x_1), f(x_2))$. Since the first coordinates of pairs of P cover the entire set X'_1 , there is $x'_2 \in X'_2$ such that $(x_1, x'_2) \in P$. By assumption, $\tilde{d}_Y(f(X'_1), f(X'_2)) > d_Y(f(x_1), f(x'_2))$, which implies $d_Y(f(x_1), f(x_2)) > d_Y(f(x_1), f(x'_2))$. However, since $f(x_1) \overset{f(X'_2)}{\sim} f(x_2)$, the element $f(x'_2) \in f(X'_2)$ cannot be strictly closer to $f(x_1)$ than $f(x_2)$ is, which is a contradiction. \square

We note that the \Rightarrow implication does not hold in general if $\tilde{d}_Y(\cdot, \cdot)$ max-distance, since x_1 and x_2 that achieve $\tilde{d}_Y(f(X'_1), f(X'_2)) = d_Y(f(x_1), f(x_2))$ are not necessarily friends, and the equality $\tilde{d}_Y(f(x_1), f(x_2)) > d_Y(f(x_1), f(x'_2))$ does not give any contradictions.

Example 3.3. A tiny counterexample would be $X = Y = \mathbb{R}$, $d_X(x, y) = d_Y(x, y) = |x - y|$, $X'_1 = \{0, 11\}$, $X'_2 = \{2, 9\}$, and $f(x) := x \bmod 2$. Since 0 is closer to 2, and 11 is closer to 9, using A- and B-sensitivity, we could give bounds for $d_Y(f(0), f(2)) = d_Y(0, 0)$ and $d_Y(f(11), f(9)) = d_Y(1, 1) = 0$, for

which an arbitrarily small c is suitable. However, for max-distance we have $\tilde{d}_Y(f(\{0, 11\}), f(\{2, 9\})) = \tilde{d}_Y(\{0, 1\}, \{0, 11\}) = 1$, and $\tilde{d}_X(\{0, 11\}, \{2, 9\}) \leq 9$ (whenever Hausdorff distance or max-distance is used for $d_X(\cdot, \cdot)$), which allows only $c \geq \frac{1}{9}$.

Theorem 3.13. *Let (ρ, σ) be a adjacent component pair. Let the mapping $f : X \rightarrow Y$ be c_ρ and c_σ A-sensitive w.r.t. ρ and σ respectively. Then, f is $(c_\rho + c_\sigma)$ A-sensitive w.r.t. $\rho \sqcap \sigma$.*

Proof. Let $x, x' \in X$ be such that $x \stackrel{[\rho \sqcap \sigma]}{\sim} x'$. We need to show that $d_Y(f(x), f(x')) \leq (c_\rho + c_\sigma) \cdot d_X(x, x')$. Since (ρ, σ) is a adjacent pair, there exists $x'' \in X$ such that $x \stackrel{[\rho]}{\sim} x''$ and $x' \stackrel{[\sigma]}{\sim} x''$. Since $x'' \in [x']_\rho$ and $x'' \in [x]_\sigma$, we have $[x'']_\rho = [x']_\rho$, $[x'']_\sigma = [x]_\sigma$, and $x \stackrel{[\rho]}{\sim} x''$, $x' \stackrel{[\sigma]}{\sim} x''$, so we can apply A-sensitivity:

- $d_Y(f(x), f(x'')) \leq c_\rho \cdot d_X(x, x'')$;
- $d_Y(f(x''), f(x')) \leq c_\sigma \cdot d_X(x'', x')$.

Since $x \stackrel{[\rho]}{\sim} x''$, the element x' cannot be closer to x than x'' is, so $d_X(x, x'') \leq d_X(x, x')$.

Since $x' \stackrel{[\sigma]}{\sim} x''$, the element x cannot be closer to x' than x'' is, so $d_X(x'', x') \leq d_X(x, x')$.

We now estimate $d_Y(f(x), f(x'))$ from above, using the triangle inequality.

$$\begin{aligned} d_Y(f(x), f(x')) &\leq d_Y(f(x), f(x'')) + d_Y(f(x''), f(x')) \\ &\leq c_\rho \cdot d_X(x, x'') + c_\sigma \cdot d_X(x'', x') \\ &\leq c_\rho \cdot d_X(x, x') + c_\sigma \cdot d_X(x, x') \\ &= (c_\rho + c_\sigma) \cdot d_X(x, x') . \end{aligned}$$

□

Corollary 3.14. *Let (ρ, σ) be an adjacent pair of equidistant components. Let the mapping $f : X \rightarrow Y$ be c_ρ and c_σ B-sensitive w.r.t. ρ and σ respectively. Then, f is $(c_\rho + c_\sigma)$ B-sensitive w.r.t. $\rho \sqcap \sigma$.*

Proof. Let $x, x' \in X$, $x \stackrel{[\rho \sqcap \sigma]}{\sim} x'$. Since ρ and σ are equidistant, B-sensitivity (Definition 3.20) implies A-sensitivity (Definition 3.19), so f is c_ρ and c_σ A-sensitive w.r.t. ρ and σ respectively. Applying Theorem 3.13, we get $d_Y(f(x), f(x')) \leq (c_\rho + c_\sigma) \cdot d_X(x, x')$. By Lemma 3.11, $d_X(x, x') \leq \tilde{d}_X([x]_\rho, [x']_\rho)$, and this gives us $d_Y(f(x), f(x')) \leq (c_\rho + c_\sigma) \cdot \tilde{d}_X([x]_\rho, [x']_\rho)$, which is the definition of B-sensitivity. □

Theorem 3.15. *Let (ρ, σ) be an adjacent expanding component pair. Let the mapping $f : X \rightarrow Y$ be c_ρ and c_σ B-sensitive w.r.t. ρ and σ respectively. Then, f is $(c_\rho + c_\sigma)$ B-sensitive w.r.t. $\rho \sqcap \sigma$.*

Proof. Constructing x'' from x and x' as in the proof of Theorem 3.13, we we can apply B-sensitivity of f to the pairs (x, x'') and (x'', x') separately, getting the following inequalities:

- $d_Y(f(x), f(x'')) \leq c_\rho \cdot \tilde{d}_X([x]_\rho, [x'']_\rho) = c_\rho \cdot \tilde{d}_X([x]_\rho, [x']_\rho)$;
- $d_Y(f(x''), f(x')) \leq c_\sigma \cdot \tilde{d}_X([x'']_\sigma, [x']_\sigma) = c_\sigma \cdot \tilde{d}_X([x]_\sigma, [x']_\sigma)$.

Since ρ and σ are expanding, we have $\tilde{d}_X([x]_\rho, [x']_\rho) \leq \tilde{d}_X([x]_{\rho \sqcap \sigma}, [x']_{\rho \sqcap \sigma})$ and $\tilde{d}_X([x]_\sigma, [x']_\sigma) \leq \tilde{d}_X([x]_{\rho \sqcap \sigma}, [x']_{\rho \sqcap \sigma})$.

We now estimate $d_Y(f(x), f(x'))$ from above, using the triangle inequality.

$$\begin{aligned} d_Y(f(x), f(x')) &\leq d_Y(f(x), f(x'')) + d_Y(f(x''), f(x')) \\ &\leq c_\rho \cdot \tilde{d}_X([x]_\rho, [x']_\rho) + c_\sigma \cdot \tilde{d}_X([x]_\sigma, [x']_\sigma) \\ &\leq c_\rho \cdot \tilde{d}_X([x]_{\rho \sqcap \sigma}, [x']_{\rho \sqcap \sigma}) + c_\sigma \cdot \tilde{d}_X([x]_{\rho \sqcap \sigma}, [x']_{\rho \sqcap \sigma}) \\ &= (c_\rho + c_\sigma) \cdot \tilde{d}_X([x]_{\rho \sqcap \sigma}, [x']_{\rho \sqcap \sigma}) . \end{aligned}$$

□

Theorem 3.16. *Let (ρ, σ) be an adjacent component pair. Let the mapping $f : X \rightarrow Y$ be c_ρ and c_σ C-sensitive w.r.t. ρ and σ respectively. Then, f is $(c_\rho + c_\sigma)$ C-sensitive w.r.t. $\rho \sqcap \sigma$.*

Proof. Let $X_1, X_2 \in X/(\rho \sqcap \sigma)$. Let $X'_1 \subseteq X_1$ and $X'_2 \subseteq X_2$ be such that $X'_1 \overset{X_2}{\approx} X'_2$. We need to show that $\tilde{d}_Y(f(X'_1), f(X'_2)) \leq (c_\rho + c_\sigma) \cdot \tilde{d}_X(X'_1, X'_2)$.

Since every class of $X/(\rho \sqcap \sigma)$ is an intersection of two classes of X/ρ and X/σ , for $i \in \{1, 2\}$, we can write $X_i = X_i^\rho \cap X_i^\sigma$, where $X_i^\rho \in X/\rho$ and $X_i^\sigma \in X/\sigma$. Hence, we can represent $X'_i \subseteq X_i$ as $X'_i = X_i^{\prime\rho} \cap X_i^{\prime\sigma}$, where $X_i^{\prime\rho} \subseteq X_i^\rho$ and $X_i^{\prime\sigma} \subseteq X_i^\sigma$.

Let $x \in X'_1$. Since $X'_1 \overset{X_2}{\approx} X'_2$, there exists $x' \in X'_2$ such that $x \overset{X_2}{\approx} x'$. Since $X_2 \in X/(\rho \sqcap \sigma)$, and (ρ, σ) is an adjacent pair, there exists x'' such that $x \overset{X_2^\rho}{\approx} x''$ and $x' \overset{X_2^\sigma}{\approx} x''$. Let $X'' := \{x'' \mid x \in X'_1\}$ be the set of all such x'' that exist for all elements of X'_1 (if there are several such elements for the same x , take all of them, so that a suitable $x'' \in X''$ would exist for all $x' \in X'_2$).

Since X'' is constructed in such a way that $\forall x \in X'_1 \exists x'' \in X'' : x \overset{X_2^\rho}{\approx} x''$, and there are no other elements in X'' , we have $X'_1 \overset{X_2^\rho}{\approx} X''$, so we can apply C-sensitivity of f to them. Similarly, since X'_2 by assumption contains only friends of X'_1 and no other elements, we have $\forall x' \in X'_2 \exists x'' \in X'' : x' \overset{X_2^\sigma}{\approx} x''$, and since there are no elements in X'' for which such x' does not exist, we have $X'_2 \overset{X_2^\sigma}{\approx} X''$, so we can apply C-sensitivity of f :

- $\tilde{d}_Y(f(X'_1), f(X'')) \leq c_\rho \cdot \tilde{d}_X(X'_1, X'')$;
- $\tilde{d}_Y(f(X''), f(X'_2)) \leq c_\sigma \cdot \tilde{d}_X(X'', X'_2)$.

Consider any pair $x \in X'_1, x'' \in X''$, such that $x \overset{X_2^\rho}{\approx} x''$. Since $x \overset{X_2^\rho}{\approx} x''$, no element $x' \in X'_2 \subseteq X_2$ can be closer to x than x'' is, so $d_X(x, x'') \leq d_X(x, x')$ for all $x' \in X'_2$. This inequality holds for all pairs of friends $x \in X'_1, x'' \in X''$, and hence also for the pair that gives the maximal distance, so $\tilde{d}_X(X'_1, X'') \leq d_X(x, x')$.

On the other hand, it holds for all $x' \in X'_2$, and hence also for x' such that $x \overset{X_2^\sigma}{\approx} x'$, and by Lemma 3.11 we have $d_X(x, x') \leq \tilde{d}_X(X'_1, X'_2)$. Putting all together, we get $\tilde{d}_X(X'_1, X'') \leq \tilde{d}_X(X'_1, X'_2)$. Similarly, we get $\tilde{d}_X(X'', X'_2) \leq \tilde{d}_X(X'_1, X'_2)$.

We now estimate $\tilde{d}_Y(f(X'_1), f(X'_2))$ from above, using the triangle inequality.

$$\begin{aligned} \tilde{d}_Y(f(X'_1), f(X'_2)) &\leq \tilde{d}_Y(f(X'_1), f(X'')) + \tilde{d}_Y(f(X''), f(X'_2)) \\ &\leq c_\rho \cdot \tilde{d}_X(X'_1, X'') + c_\sigma \cdot \tilde{d}_X(X'', X'_2) \\ &\leq c_\rho \cdot \tilde{d}_X(X'_1, X'_2) + c_\sigma \cdot \tilde{d}_X(X'_1, X'_2) \\ &= (c_\rho + c_\sigma) \cdot \tilde{d}_X(X'_1, X'_2) . \end{aligned}$$

□

Corollary 3.17. *Let (ρ, σ) be an adjacent pair of equidistant components. Let the mapping $f : X \rightarrow Y$ be c_ρ and c_σ D-sensitive w.r.t. ρ and σ respectively. Then, f is $(c_\rho + c_\sigma)$ D-sensitive w.r.t. $\rho \sqcap \sigma$.*

Proof. Similarly to Corollary 3.14, since the components are equidistant, we have $d_X(x, x') = \tilde{d}_X(X_1, X_2)$ for all $x \in X_1, x' \in X_2$ such that $x \overset{X_2}{\approx} x'$, and $X_1, X_2 \in X/\rho$ or $X_1, X_2 \in X/\sigma$. Hence, also $\tilde{d}_X(X'_1, X'_2) = \tilde{d}_X(X_1, X_2)$ for $X'_1 \subseteq X_1$ and $X'_2 \subseteq X_2$ such that $X'_1 \overset{X_2}{\approx} X'_2$, so C-sensitivity of ρ and σ implies their D-sensitivity. Applying Theorem 3.16, we get $\tilde{d}_Y(f(X'_1), f(X'_2)) \leq (c_\rho + c_\sigma) \cdot \tilde{d}_X(X'_1, X'_2)$. Since we have $X'_1 \overset{X_2}{\approx} X'_2$, every element $x \in X'_1$ has a friend in X'_2 that is also a friend in X_2 and hence, by Lemma 3.11, is at most $\tilde{d}_X(X_1, X_2)$ far away from it. This gives us $\tilde{d}_X(X'_1, X'_2) \leq \tilde{d}_X(X_1, X_2)$. □

Theorem 3.18. *Let (ρ, σ) be an adjacent expanding component pair. Let the mapping $f : X \rightarrow Y$ be c_ρ and c_σ D-sensitive w.r.t. ρ and σ respectively. Then, f is $(c_\rho + c_\sigma)$ D-sensitive w.r.t. $\rho \sqcap \sigma$.*

Proof. Constructing X'' from X'_1 and X'_2 as in the proof of Theorem 3.16, we we can apply C-sensitivity of f to the pairs (X'_1, X'') and (X'_2, X'') separately, getting the following inequalities:

- $\tilde{d}_Y(f(X'_1), f(X'')) \leq c_\rho \cdot \tilde{d}_X(X_1^\rho, X_2^\rho)$;

- $\tilde{d}_Y(f(X''), f(X'_2)) \leq c_\sigma \cdot \tilde{d}_X(X_1^\sigma, X_2^\sigma)$.

Since ρ and σ are expanding, we have $\tilde{d}_X(X_1^\rho, X_2^\rho) \leq \tilde{d}_X(X_1^\rho \cap X_1^\sigma, X_2^\rho \cap X_2^\sigma)$, and $\tilde{d}_X(X_1^\sigma, X_2^\sigma) \leq \tilde{d}_X(X_1^\rho \cap X_1^\sigma, X_2^\rho \cap X_2^\sigma)$. By construction, $\tilde{d}_X(X_1^\rho \cap X_1^\sigma, X_2^\rho \cap X_2^\sigma) = \tilde{d}_X(X_1, X_2)$.

We now estimate $\tilde{d}_Y(f(X'_1), f(X'_2))$ from above, using the triangle inequality.

$$\begin{aligned} \tilde{d}_Y(f(X'_1), f(X'_2)) &\leq \tilde{d}_Y(f(X'_1), f(X'')) + \tilde{d}_Y(f(X''), f(X'_2)) \\ &\leq c_\rho \cdot \tilde{d}_X(X_1^\rho, X_2^\rho) + c_\sigma \cdot \tilde{d}_X(X_1^\sigma, X_2^\sigma) \\ &\leq c_\rho \cdot \tilde{d}_X(X_1, X_2) + c_\sigma \cdot \tilde{d}_X(X_1, X_2) \\ &= (c_\rho + c_\sigma) \cdot \tilde{d}_X(X_1, X_2) . \end{aligned}$$

□

3.3.5.4 Example of Vectors with Independent Coordinates. We now consider metric spaces $(X, d_X(\cdot, \cdot))$ and components ρ of certain form.

- $X := X_1 \times \cdots \times X_n$;
- $d_X(\vec{x}, \vec{x}') := \|\vec{x} - \vec{x}'\|_p$ or $d_X(\vec{x}, \vec{x}') := \|\vec{x} - \vec{x}'\|_\infty$;
- $\tilde{d}_X(X_1, X_2)$ is the Hausdorff distance;
- ρ_i is such that $X_a \in X/\rho_i$ is defined as $X_a := \{(x_1, \dots, x_{i-1}, a, x_{i+1}, \dots, x_n) \mid x_j \in X_j\}$.

In this section, we will consider components of the form $\rho := \prod_{i \in \mathcal{I}} \rho_i$. Since the norms ℓ_p and ℓ_∞ do not depend on the ordering of the coordinates X_i , without loss of generality, let $\mathcal{I} = [m]$ for some $m \leq n$. Let the classes of $X/\prod_{i \in [m]} \rho_i$ be denoted by $X_{(a_i)_{i \in [m]}} = \{(a_1, \dots, a_m, x_{m+1}, \dots, x_n) \mid x_j \in X_j, m < j \leq n\}$.

Lemma 3.19. *Let $\rho = \prod_{i \in [m]} \rho_i$; $X_{(a_i)_{i \in [m]}}, X_{(a'_i)_{i \in [m]}} \in X/\rho$. For $\vec{x} = (a_1, \dots, a_m, x_{m+1}, \dots, x_n) \in X_{(a_i)_{i \in [m]}}$, if $\vec{x}' = (a'_1, \dots, a'_m, x_{m+1}, \dots, x_n)$, then $\vec{x} \stackrel{X_{(a'_i)_{i \in [m]}}}{\sim} \vec{x}'$.*

Proof. First, let us assume that ℓ_p norm is used. We get $d_X(\vec{x}, \vec{x}') = (\sum_{i=1}^m |a'_i - a_i|^p + \sum_{i=m+1}^n |x_i - x_i|^p)^{1/p} = (\sum_{i=1}^m |a'_i - a_i|^p)^{1/p}$. We cannot achieve a better bound since a'_i are fixed, and taking $x'_i \neq x_i$ in \vec{x}' may only increase the absolute value. Hence, \vec{x}' minimizes $d_X(\vec{x}, \vec{x}')$ in $X_{(a'_i)_{i \in [m]}}$, so $\vec{x} \stackrel{X_{(a'_i)_{i \in [m]}}}{\sim} \vec{x}'$.

Similarly, if ℓ_∞ norm is used, then $d_X(\vec{x}, \vec{x}') = \max(\max_{i=1}^m |a'_i - a_i|, \max_{i=m+1}^n |x_i - x_i|) = \max_{i=1}^m |a'_i - a_i|$, and taking $x'_i \neq x_i$ in \vec{x}' may only increase the maximum. □

Proposition 3.20. *Any component of the form $\rho = \prod_{i \in [m]} \rho_i$ is equidistant. In particular, $\tilde{d}_X(X_{(a_i)_{i \in [m]}}, X_{(a'_i)_{i \in [m]}}) = (\sum_{i=1}^m |a'_i - a_i|^p)^{1/p}$ for ℓ_p norm, and $\tilde{d}_X(X_{(a_i)_{i \in [m]}}, X_{(a'_i)_{i \in [m]}}) = \max_{i=1}^m |a'_i - a_i|$ for ℓ_∞ norm.*

Proof. Let $X_{(a_i)_{i \in [m]}}, X_{(a'_i)_{i \in [m]}} \in X/\rho$. For any element $\vec{x} = (a_1, \dots, a_m, x_{m+1}, \dots, x_n) \in X_{(a_i)_{i \in [m]}}$, we can take $\vec{x}' = (a'_1, \dots, a'_m, x_{m+1}, \dots, x_n) \in X_{(a'_i)_{i \in [m]}}$. By Lemma 3.19, $\vec{x} \stackrel{X_{(a'_i)_{i \in [m]}}}{\sim} \vec{x}'$.

We have $d_X(\vec{x}, \vec{x}') = (\sum_{i=1}^m |a'_i - a_i|^p + \sum_{i=m+1}^n |x_i - x_i|^p)^{1/p} = (\sum_{i=1}^m |a'_i - a_i|^p)^{1/p}$. Since such \vec{x}' exists for all $\vec{x} \in X_{(a_i)_{i \in [m]}}$, we also have $\max_{\vec{x} \in X_{(a_i)_{i \in [m]}}} \min_{\vec{x}' \in X_{(a'_i)_{i \in [m]}}} d_X(\vec{x}, \vec{x}') = (\sum_{i=1}^m |a'_i - a_i|^p)^{1/p}$. The same holds if we swap $X_{(a_i)_{i \in [m]}}$ and $X_{(a'_i)_{i \in [m]}}$. By definition of Hausdorff distance, $\tilde{d}_X(X_{(a_i)_{i \in [m]}}, X_{(a'_i)_{i \in [m]}}) = (\sum_{i=1}^m |a'_i - a_i|^p)^{1/p}$.

Similarly, if ℓ_∞ norm is used, then $\max(\max_{i=1}^m |a'_i - a_i|, \max_{i=m+1}^n |x_i - x_i|) = \max_{i=1}^m |a'_i - a_i|$ for all $\vec{x} \in X_{(a_i)_{i \in [m]}}$, so $\tilde{d}_X(X_{(a_i)_{i \in [m]}}, X_{(a'_i)_{i \in [m]}}) = \max_{i=1}^m |a'_i - a_i|$. □

Proposition 3.21. *The components $\rho := \prod_{i \in [m]} \rho_i$ and $\sigma := \rho_{m+1}$ are adjacent.*

Proof. Let $X_{(a_i)_{i \in [m]}}, X_{(a'_i)_{i \in [m]}} \in X/\rho$, and let $X_a, X_{a'} \in X/\sigma$. For any $\vec{x} = (a_1, \dots, a_m, a, x_{m+2}, \dots, x_n) \in X_{(a_i)_{i \in [m]}} \cap X_a$ and $\vec{x}' = (a'_1, \dots, a'_m, a', x_{m+2}, \dots, x_n) \in X_{(a'_i)_{i \in [m]}} \cap X_{a'}$, we can define $\vec{x}'' \in X_{(a'_i)_{i \in [m]}} \cap X_a$ as $\vec{x}'' = (a'_1, \dots, a'_m, a, x_{m+2}, \dots, x_n)$. By Lemma 3.19, we have $\vec{x} \stackrel{X_{(a'_i)_{i \in [m]}}}{\sim} \vec{x}''$ and $\vec{x} \stackrel{X_a}{\sim} \vec{x}''$. □

Corollary 3.22. For all $i \in [n]$, let c_i be I -sensitivity of a function f w.r.t. the component ρ_i , for $I \in \{A, B, C, D\}$. Then $\sum_{i \in \mathcal{I}} c_i$ is the I -sensitivity of f w.r.t. the component $\prod_{i \in \mathcal{I}} \rho_i$ for all $\mathcal{I} \subseteq [n]$.

Proof. Prop. 3.20 and Prop. 3.21 together give us the necessary preconditions for applying Theorem 3.13, Corollary 3.14, Theorem 3.16, Corollary 3.17 for all variants of sensitivity definition. \square

3.3.5.5 Example of Vectors with Linearly Dependent Coordinates. We consider metric spaces $(\hat{X}, d_X(\cdot, \cdot))$ of the same form as in Sec. 3.3.5.4, with the following additional conditions.

- $\forall i, j : X_i = X_j =: Z$;
- There is a matrix $\mathbf{A} \in Z^{p \times n}$ for $p \leq n$, and a vector $\vec{b} \in Z^p$ such that, $\forall \vec{x} \in \hat{X} : \mathbf{A}\vec{x} = \vec{b}$.

Similarly to Sec. 3.3.5.4, we consider components of the form $\hat{\rho} := \prod_{i \in \mathcal{I}} \hat{\rho}_i$, where $\hat{\rho}_i$ is the analog of ρ_i , defined over \hat{X} . For equivalence classes, we will use the notation

$$\hat{X}_{(a_i)_{i \in [m]}} = \{(a_1, \dots, a_m, x_{m+1}, \dots, x_n) \mid x_j \in X_j, m < j \leq n, \mathbf{A}\vec{x} = \vec{b}\}$$

to discern them from the analogous classes of independent components.

Let $\vec{x} = \vec{x}_1 | \vec{x}_2$ denote that \vec{x} is a concatenation of vectors \vec{x}_1 and \vec{x}_2 . For matrices, $\mathbf{A} = \mathbf{A}_1 | \mathbf{A}_2$ denotes concatenation of rows.

Lemma 3.23. Without loss of generality, let $\mathcal{I} = [m]$. Let $\hat{\rho} = \prod_{i \in [m]} \hat{\rho}_i$, and $\hat{X}_{(a_i)_{i \in [m]}}$, $\hat{X}_{(a'_i)_{i \in [m]}} \in \hat{X}/\hat{\rho}$. Let $\vec{x} \in \hat{X}$ be such that $x_i = a_i$ for $i \in [m]$, and $\mathbf{A}\vec{x} = \vec{b}$. Let $\vec{x}' \in \hat{X}$ be such that $\vec{x} \stackrel{X_{(a'_i)_{i \in [m]}}}{\sim} \vec{x}'$.

Let $\mathbf{A} = \mathbf{A}_1 | \mathbf{A}_2$, where \mathbf{A}_1 are the first m columns and \mathbf{A}_2 are the rest. We can write $\vec{x}' = \vec{x} + \vec{z}$, where $\vec{z} = \vec{z}_1 | \vec{z}_2$, $\vec{z}_1 = (a' - a)$, and \vec{z}_2 is an optimal solution to the task

$$\text{minimize } \|\vec{z}_2\|, \text{ subject to } \mathbf{A}_2 \vec{z}_2 = \mathbf{A}_1 (a' - a) .$$

Proof. In order to find the vector \vec{x}' that is the closest to \vec{x} , we need to solve the following task:

$$\text{minimize } \|\vec{x}' - \vec{x}\|, \text{ subject to } \mathbf{A}\vec{x}' = \vec{b} .$$

We would like to get rid of x_j that make this quantity dependent on the particular instance \vec{x} . We may introduce variables $\vec{z} := \vec{x}' - \vec{x}$. This gives us linear constraints $\mathbf{A}(\vec{x} + \vec{z}) = \vec{b}$. Since by assumption $\mathbf{A}\vec{x} = \vec{b}$, this is equivalent to $\mathbf{A}\vec{z} = \vec{0}$. We get the task

$$\text{minimize } \|\vec{z}\|, \text{ subject to } \mathbf{A}\vec{z} = \vec{0} .$$

We can rewrite $\mathbf{A}\vec{z}$ as $(\mathbf{A}_1 | \mathbf{A}_2)(\vec{z}_1 | \vec{z}_2) = \mathbf{A}_1 \vec{z}_1 + \mathbf{A}_2 \vec{z}_2$. For $j \in [m]$, we have $z_j = (a'_j - a_j)$, so $\mathbf{A}_1 \vec{z}_1 = \mathbf{A}_1 (a' - a)$. Substituting this value into the constraints, we get the statement of this lemma. \square

Like in the case of independent components, we could first define sensitivity for each single component, and then define their composition. However, it is no longer reasonable to define it for friends as in Sec. 3.3.5.4: due to the constraints, the closest vectors may have very different entries, and we might be unable to define sensitivity even for single components that we could use as building blocks. Moreover, using the same definitions as in Sec. 3.3.5.4, we will not achieve the necessary preconditions for the composition theorems, since dependent components are in general not adjacent.

Instead of defining new theory for dependent components, let us compute upper bounds directly from the theory of independent components. We assume that we only know how to compute initial sensitivity for *independent* components ρ_j , i.e. we can give upper bounds only for $d_X(f(x), f(x'))$ where $x_i = x'_i$ for all $i \neq j$. We use them to construct sensitivities for *dependent* components $\hat{\rho}_j$. We will derive the bounds not for friends, but for *matches* – the vectors that have maximal amount of similar coordinates. Let $d_{Ham}(\vec{x}, \vec{x}')$ denote the *Hamming distance* of \vec{x} and \vec{x}' , i.e. the number of their different coordinates.

Definition 3.26 (match in the set). Let $\hat{X}' \subseteq \hat{X}$, $\vec{x} \in \hat{X}$. An element $\vec{x}' \in \hat{X}'$ is called a *match* of \vec{x} in the set \hat{X}' , denoted $\vec{x} \stackrel{\hat{X}'}{\approx} \vec{x}'$, if

$$\vec{x}' \in \arg \min_{\vec{x}' \in \hat{X}'} d_{Ham}(\vec{x}, \vec{x}') .$$

Definition 3.27 (matching sets). Subsets $X_1, X_2 \subseteq X$ are called *matching* in X' , denoted $X_1 \stackrel{X'}{\approx} X_2$, if the following conditions hold:

- $\forall \vec{x}_1 \in X_1 \exists \vec{x}_2 \in X_2 : \vec{x}_1 \stackrel{X'}{\approx} \vec{x}_2$;
- $\forall \vec{x}_2 \in X_2 \exists \vec{x}_1 \in X_1 : \vec{x}_2 \stackrel{X'}{\approx} \vec{x}_1$;

The notions of a friend and a match (as well as friendly sets and matching sets) are equivalent for vectors with independent components, but they are different in the case of dependent components.

Lemma 3.24. *If $\vec{x} \stackrel{\hat{X}_{(a_i)_{i \in [m]}}}{\approx} \vec{x}'$, then there are at least $\max(0, n - p - m)$ equal entries in \vec{x} and \vec{x}' .*

Proof. Since $\vec{x}' \in \hat{X}_{(a_i)_{i \in [m]}}$, we already have $x'_i = a_i$ for $i \in [m]$, and it may happen that $a_i \neq x_i$ for all $i \in [m]$, so in general we do not get equal entries from here. The linear equation system has degree of freedom $n - p$, so even after m entries are fixed, if $p + m < n$, then there are at least $n - p - m$ entries whose value can be set to arbitrary, and $x'_j = x_j$ can be taken without violating the constraints. The remaining $\min(n - m, p)$ entries are uniquely determined by the constraints and the other entries. \square

Proposition 3.25. *Any component of the form $\hat{\rho} = \prod_{i \in [m]} \hat{\rho}_i$ is equidistant.*

Proof. Let $\hat{X}_{(a_i)_{i \in [m]}}$, $\hat{X}_{(a'_i)_{i \in [m]}} \in \hat{X}/\hat{\rho}$. For any element $\vec{x} = (a_1, \dots, a_m, x_{m+1}, \dots, x_n) \in \hat{X}_{(a_i)_{i \in [m]}}$, we can take $\vec{x}' \in \hat{X}$ such that $\vec{x} \stackrel{\hat{X}_{(a'_i)_{i \in [m]}}}{\approx} \vec{x}'$, using the construction of Lemma 3.23. As in the proof of Prop. 3.20, it is only important that $d_X(\vec{x}, \vec{x}')$ depends only on the classes $\hat{X}_{(a_i)_{i \in [m]}}$, $\hat{X}_{(a'_i)_{i \in [m]}}$, and on the matrix \mathbf{A} that is fixed, so it can be expressed as some quantity $g(\mathbf{A}, (a_i, a'_i)_{i \in [m]})$ that does not depend on any other entries of \vec{x} and \vec{x}' . Since such \vec{x}' exists for all $\vec{x} \in \hat{X}_{(a_i)_{i \in [m]}}$, we also have $\max_{\vec{x} \in \hat{X}_{(a_i)_{i \in [m]}}} \min_{\vec{x}' \in \hat{X}_{(a'_i)_{i \in [m]}}} d_X(\vec{x}, \vec{x}') = g(\mathbf{A}, (a_i, a'_i)_{i \in [m]})$, and also $\max_{\vec{x}' \in \hat{X}_{(a'_i)_{i \in [m]}}} \min_{\vec{x} \in \hat{X}_{(a_i)_{i \in [m]}}} d_X(\vec{x}, \vec{x}') = g(\mathbf{A}, (a_i, a'_i)_{i \in [m]})$. By definition of the Hausdorff distance, we get $\tilde{d}_X(\hat{X}_{(a_i)_{i \in [m]}}, \hat{X}_{(a'_i)_{i \in [m]}}) = g(\mathbf{A}, (a_i, a'_i)_{i \in [m]})$. \square

Proposition 3.26 (A-sensitivity w.r.t. component). *Let mapping $f : X \rightarrow Y$ be c_i A-sensitive w.r.t. the component ρ_i for all $i \in \{1, \dots, n\}$. Let $\hat{\rho} := \prod_{i \in \mathcal{I}} \hat{\rho}_i$. For all $\hat{X}_1, \hat{X}_2 \in \hat{X}/\hat{\rho}$ we have*

$$\forall \vec{x}_1 \in \hat{X}_1, \vec{x}_2 \in \hat{X}_2 : \vec{x}_1 \stackrel{\hat{X}_1}{\approx} \vec{x}_2 \implies d_Y(f(\vec{x}_1), f(\vec{x}_2)) \leq \left(p \cdot \max_{i \in [n] \setminus \mathcal{I}} (c_i) + \sum_{i \in \mathcal{I}} c_i \right) \cdot d_X(\vec{x}_1, \vec{x}_2) .$$

Proof. Since $\vec{x}_1 \stackrel{\hat{X}_1}{\approx} \vec{x}_2$, by Lemma 3.24, there are at least $n - p - |\mathcal{I}|$ equal entries in \vec{x}_1 and \vec{x}_2 . Therefore, \vec{x}_2 can be considered a friend of \vec{x}_1 in the set $X_{(a'_i)_{i \in \mathcal{I}} \cup (x'_j)_{j \in \mathcal{J}}}$ (note that it is a class of vectors with *independent* components), where \mathcal{J} , such that $\mathcal{J} \cap \mathcal{I} = \emptyset$, is the set of the remaining p entries that are not necessarily equal and are not part of $\hat{\rho}$. We can now apply Corollary 3.22, getting sensitivity $\sum_{i \in \mathcal{I} \cup \mathcal{J}} c_i = \sum_{i \in \mathcal{I}} c_i + \sum_{j \in \mathcal{J}} c_j$. In general, we do not know which \mathcal{J} exactly will be taken for these particular \vec{x} and \vec{x}' , but in any case $|\mathcal{J}| = \min(n - |\mathcal{I}|, p) \leq p$, and $p \cdot \max_{i \in [n] \setminus \mathcal{I}} (c_i)$ is a valid upper bound for $\sum_{j \in \mathcal{J}} c_j$. \square

Corollary 3.27 (B-sensitivity w.r.t. component). *Let mapping $f : X \rightarrow Y$ be c_i B-sensitive w.r.t. the component ρ_i for all $i \in \{1, \dots, n\}$. Let $\hat{\rho} := \prod_{i \in \mathcal{I}} \hat{\rho}_i$. For all $\hat{X}_1, \hat{X}_2 \in \hat{X}/\hat{\rho}$ we have*

$$\forall \vec{x}_1 \in \hat{X}_1, \vec{x}_2 \in \hat{X}_2 : \vec{x}_1 \stackrel{\hat{X}_1}{\approx} \vec{x}_2 \wedge \vec{x}_1 \stackrel{\hat{X}_2}{\approx} \vec{x}_2 \implies d_Y(f(\vec{x}_1), f(\vec{x}_2)) \leq \left(p \cdot \max_{i \in [n] \setminus \mathcal{I}} (c_i) + \sum_{i \in \mathcal{I}} c_i \right) \cdot \tilde{d}_X(\hat{X}_1, \hat{X}_2) .$$

Proof. As in the proof of Prop. 3.26, we have $\vec{x}_1 \stackrel{X_{(a_i)_{i \in \mathcal{I}} \cup (x'_j)_{j \in \mathcal{J}}}}{\sim} \vec{x}_2$. By Prop. 3.20, any composite component $\prod_{i \in \mathcal{I}} \rho_i$ is equidistant, so $d_X(\vec{x}, \vec{x}') = \tilde{d}_X(X_{(a_i)_{i \in \mathcal{I}} \cup (x_j)_{j \in \mathcal{J}}}, X_{(a'_i)_{i \in \mathcal{I}} \cup (x'_j)_{j \in \mathcal{J}}})$, and we may apply Prop. 3.26, getting the upper bound $(p \cdot \max_{i \in [n] \setminus \mathcal{I}}(c_i) + \sum_{i \in \mathcal{I}} c_i) \cdot d_X(\vec{x}_1, \vec{x}_2)$. Now, by Prop. 3.25, each composite component $\hat{\rho} := \prod_{i \in \mathcal{I}} \hat{\rho}_i$ is also equidistant, and since $\vec{x}_1 \stackrel{\hat{X}_2}{\sim} \vec{x}_2$, we have $d_X(\vec{x}_1, \vec{x}_2) = \tilde{d}_X(\hat{X}_1, \hat{X}_2)$. \square

We note that the conditions $\vec{x}_1 \stackrel{\hat{X}_2}{\sim} \vec{x}_2$ and $\vec{x}_1 \stackrel{\hat{X}_1}{\sim} \vec{x}_2$ are in general not simultaneously satisfiable. It may even happen that such \vec{x}_1 and \vec{x}_2 do not exist for some classes, since having similar coordinates may be harmful for reducing the distance due to the constraints.

For C- and D-sensitivities, we could use the assumption $\vec{X}_1 \stackrel{\hat{X}_2}{\approx} \vec{X}_2$ instead of $\vec{x}_1 \stackrel{\hat{X}_2}{\sim} \vec{x}_2$. However, when we reduce $\vec{X}_1 \stackrel{\hat{X}_2}{\approx} \vec{X}_2$ directly to $\vec{X}_1 \stackrel{X'}{\sim} \vec{X}_2$ for some set of vectors independent coordinates $X' \subseteq X$, it may happen that different pairs $x_1 \stackrel{\hat{X}_2}{\sim} x_2$ and $x'_1 \stackrel{\hat{X}_2}{\sim} x'_2$ have different similar components, and in the worst case we will only be able to take $X' = X$, which does not give any reasonable bounds. Adding restrictions on X'_1 and X'_2 , that similar components of different matching pairs should be located on the same positions, would help. In this case, the condition $\vec{X}_1 \stackrel{\hat{X}_2}{\approx} \vec{X}_2 \wedge \vec{X}_1 \stackrel{\hat{X}_1}{\approx} \vec{X}_2$ (needed for D-sensitivity) would be even more difficult to achieve.

3.3.5.6 Adding Restrictions on Constraints. In this section, we see which results we can achieve if we put some restrictions on $\mathbf{A}\vec{x} = \vec{b}$. Using Gaussian elimination, we may bring any underdetermined full-rank system to the form $\mathbf{A} = \mathbf{A}'\mathbf{I}$, where \mathbf{I} is an identity matrix, and \mathbf{A}' is some matrix of $n - p$ columns. Since \mathbf{I} does not necessarily have to be on the right, and can be scattered among any p columns, we may get different variants of \mathbf{A}' . Let us consider the case when it is possible to obtain $\|\mathbf{A}'\| < 1$, where $\|\mathbf{A}'\|$ is the *induced matrix norm*, i.e. $\|\mathbf{A}'\| = \sup_{\vec{x} \in \hat{X}} \frac{\|\mathbf{A}'\vec{x}\|}{\|\vec{x}\|}$.

In this section, we consider ℓ_1 norm and subsets of components indexed by $\mathcal{I} \subseteq [n - p]$. By properties of the induced ℓ_1 norm, $\|\mathbf{A}\| := \max_{j \in n} \sum_{i=1}^p |a_{ij}|$ for a $p \times n$ matrix \mathbf{A} . Note that the norm of a matrix does not change if we permute its columns in any way. Without loss of generality, we may assume that $\mathcal{I} = [m]$ for $m \leq n - p$, since the columns of \mathbf{A}' can be permuted arbitrarily without violating the condition $\|\mathbf{A}'\| < 1$.

Similarly to Sec. 3.3.5.5, we consider components of the form $\hat{\rho} := \prod_{i \in \mathcal{I}} \hat{\rho}_i$, where $\hat{\rho}_i$ is the analog of ρ_i , defined over \hat{X} . We use the same notation $\hat{X}_{(a_i)_{i \in [m]}}$ for equivalence classes.

Lemma 3.28. *Let $\mathcal{I} = [m]$, $m \leq n - p$. Let $\hat{\rho} = \prod_{i \in [m]} \hat{\rho}_i$, and $\hat{X}_{(a_i)_{i \in [m]}}$, $\hat{X}_{(a'_i)_{i \in [m]}} \in \hat{X}/\hat{\rho}$. Let $\vec{x} \in \hat{X}$ be such that $x_i = a_i$ for $i \in [m]$, and $\mathbf{A}\vec{x} = \vec{b}$. Let $\vec{x}' \in \hat{X}$ be such that $\vec{x} \stackrel{\hat{X}_{(a'_i)_{i \in [m]}}}{\sim} \vec{x}'$, where distance is defined using ℓ_1 norm. Let $\mathbf{A} = \mathbf{A}_1\mathbf{A}_2$, where \mathbf{A}_1 are the first m columns of \mathbf{A} .*

We can write $\vec{x}' = \vec{x} + \vec{z}$, where $\vec{z} = \vec{z}_1|\vec{z}_2$, $\vec{z}_1 = \vec{a}' - \vec{a}$, and $\vec{z}_2 = \vec{0}_{n-m-p}|\mathbf{A}_1(\vec{a} - \vec{a}')$.

Proof. From Lemma 3.23, we get that $\vec{x}' = \vec{x} + \vec{z}$, where $\vec{z} = \vec{z}_1|\vec{z}_2$, $\vec{z}_1 = (\vec{a}' - \vec{a})$, and \vec{z}_2 is an optimal solution to the task

$$\text{minimize } \|\vec{z}_2\|, \text{ subject to } \mathbf{A}_2\vec{z}_2 = \mathbf{A}_1(\vec{a} - \vec{a}'). \quad (12)$$

By assumption, we have $\mathbf{A} = \mathbf{A}'\mathbf{I}$, where \mathbf{A}' are the first $n - p$ columns of \mathbf{A} , such that $\|\mathbf{A}'\| < 1$. Hence, we have $\mathbf{A}_2 = \mathbf{A}''\mathbf{I}$, where \mathbf{A}'' are the remaining $n - m - p$ columns of \mathbf{A}' (if any). By definition of the induced ℓ_1 norm, $\|\mathbf{A}\| := \max_{j \in n} \sum_{i=1}^p |a_{ij}|$ for a $p \times n$ matrix \mathbf{A} , so $\mathbf{A}'' \leq \mathbf{A}'$ since \mathbf{A}' has more columns and, hence, more choices for the maximum. We propose that $\vec{z}_2 = \vec{0}_{n-m-p}|\mathbf{A}_1(\vec{a} - \vec{a}')$, where $\vec{0}_{n-m-p}$ is a vector of $n - m - p$ zeroes, will be the optimal solution.

1. Substituting \vec{z}_2 with this value, we get $\mathbf{A}_2\vec{z}_2 = \mathbf{A}'' \cdot \vec{0}_{n-m-p} + \mathbf{I} \cdot \mathbf{A}_1(\vec{a} - \vec{a}') = \mathbf{A}_1(\vec{a} - \vec{a}')$, so the constraints are satisfied. In this case, $\|\vec{z}_2\| = \|\mathbf{A}_1(\vec{a} - \vec{a}')\|$.
2. Let now \vec{z} be an arbitrary vector that satisfies the constraints. We have $\|\mathbf{A}_1(\vec{a} - \vec{a}')\| = \|\mathbf{A}_2\vec{z}_2\| = \|\mathbf{A}''\vec{z}_2 + \mathbf{I}\vec{z}_2\|$. Using the triangle inequality and norm submultiplicativity, we get

$\|\mathbf{A}''\vec{z}_{21} + \mathbf{I}\vec{z}_{22}\| \leq \|\mathbf{A}''\vec{z}_{21}\| + \|\mathbf{I}\vec{z}_{22}\| \leq \|\mathbf{A}''\| \cdot \|\vec{z}_{21}\| + \|\vec{z}_{22}\| \leq \|\vec{z}_{21}\| + \|\vec{z}_{22}\|$. For ℓ_1 norm, $\|\vec{z}_{21}\| + \|\vec{z}_{22}\| = \|\vec{z}_{21}\vec{z}_{22}\| = \|\vec{z}\|$.

Since we have a strict inequality $\|\mathbf{A}''\| < 1$, the last inequality in this chain is also strict whenever $\|\vec{z}_{21}\| \neq \vec{0}$, which holds for all $\vec{z}_{21} \neq \vec{0}$. In this way, any other solution is strictly worse than $\vec{z}_2 = \vec{0}_{n-m-p} | \mathbf{A}_1(\vec{a} - \vec{a}')$. \square

If we took $\|\mathbf{A}'\| \leq 1$ instead of $\|\mathbf{A}'\| < 1$ in Lemma 3.28, we would still get that $\vec{z}_2 = \vec{0}_{n-m-p} | \mathbf{A}_{\{m+1\}}(\vec{a} - \vec{a}')$ is an optimal solution, but there could possibly be some other optimal solutions that are of different form. If solutions of different form were allowed, we would not always be able to construct a suitable mutual friend whose existence is required for the adjacent components.

Proposition 3.29. *Let $m \leq n - p - 1$. The components $\hat{\rho} := \prod_{i \in [m]} \hat{\rho}_i$ and $\hat{\sigma} := \hat{\rho}_{m+1}$ are adjacent.*

Proof. Let $\hat{X}_{(a_i)_{i \in [m]}}$, $\hat{X}_{(a'_i)_{i \in [m]}} \in \hat{X}/\hat{\rho}$, and let $\hat{X}_a, \hat{X}_{a'} \in \hat{X}/\hat{\sigma}$. Let $\vec{x} = (a_1, \dots, a_m, a, x_{m+2}, \dots, x_n) \in \hat{X}_{(a_i)_{i \in [m]}} \cap \hat{X}_a$ and $\vec{x}' = (a'_1, \dots, a'_m, a', x_{m+2}, \dots, x_n) \in \hat{X}_{(a'_i)_{i \in [m]}} \cap \hat{X}_{a'}$. By definition of adjacent components, their mutual friend should be some $\vec{x}'' \in \hat{X}_{(a'_i)_{i \in [m]}} \cap \hat{X}_{a_{m+1}}$. Let $\vec{a} := (a_1, \dots, a_m)$, $\vec{a}' := (a'_1, \dots, a'_m)$, $\vec{z} := \vec{x}'' - \vec{x}$. Denote $\vec{a}_1 := \mathbf{A}_{\{1, \dots, m\}}(\vec{a} - \vec{a}')$, $\vec{a}_2 := \mathbf{A}_{\{m+1\}}(a - a')$, and $\mathbf{A}' := \mathbf{A}_{\{m+2, \dots, n\}}$.

The first $m+1$ coordinates of \vec{x}'' are in any case equal to (a'_1, \dots, a'_m, a) , and we cannot choose them. However, as far as $m+1 \leq n-p$, it is still possible to find $\vec{x}'' \in \hat{X}_{(a'_i)_{i \in [m]}} \cap \hat{X}_{a_{m+1}}$ that does not violate the constraints (it is uniquely determined by constraints if $m+1 = n-p$).

Let $m+1 \leq n-p$. By Lemma 3.28, the vector \vec{z} is such, that $\vec{z}_2 = \vec{0}_{n-(m+1)-p} | \mathbf{A}_{\{1, \dots, m+1\}}(\vec{a} | (a - a')) = \vec{0}_{n-(m+1)-p} | (\vec{a}_1 + \vec{a}_2)$. Let us now take $\vec{x}'' = \vec{x} + (\vec{y}_1 | \vec{y}_2)$, where $\vec{y}_2 = \vec{0}_{n-(m+1)-p} | \vec{a}_1$. By Lemma 3.28, $\vec{x} \stackrel{\hat{X}_{(a'_i)_{i \in [m]}}}{\sim} \vec{x}''$. On the other hand, $\|\vec{z}_2 - \vec{y}_2\| = \vec{0}_{n-(m+1)-p} | (\vec{a}_1 + \vec{a}_2 - \vec{a}_1) = \vec{0}_{n-(m+1)-p} | \vec{a}_2$. Hence, by Lemma 3.28, also $\vec{x}'' \stackrel{\hat{X}_{a_{m+1}}}{\sim} \vec{x}''$. \square

Proposition 3.30. *The components $\hat{\rho} := \prod_{i \in [m]} \hat{\rho}_i$ and $\hat{\sigma} := \hat{\rho}_{m+1}$ are in general NOT expanding.*

Proof. Let $\hat{X}_{(a_i)_{i \in [m]}}$, $\hat{X}_{(a'_i)_{i \in [m]}} \in \hat{X}/\hat{\rho}$, and let $\hat{X}_a, \hat{X}_{a'} \in \hat{X}/\hat{\sigma}$. By Lemma 3.28, the distance between all friends of the sets $\hat{X}_{(a_i)_{i \in [m]}}$ and $\hat{X}_{(a'_i)_{i \in [m]}}$ is $\|\mathbf{A}_{\{1, \dots, m\}}(\vec{a} - \vec{a}')\|$. For the finer component $\hat{\rho} \cap \hat{\sigma}$, this distance is $\|\mathbf{A}_{\{m+1\}}(a - a')\|$. In general, we do not know whether it is larger or smaller. It depends not only on the constraints, but also on the particular classes that we consider, so we also cannot give sufficient conditions for \mathbf{A} that would make all components expanding. \square

Similarly to vectors with independent components, the adjacent property allows us to apply A- and C-sensitivity using Theorem 3.13 and Theorem 3.16. Although the components are not expanding, since they are equidistant, we can also apply B- or D-sensitivity using Corollary 3.14 and Corollary 3.17. However, the components are adjacent only as far as we compose up to $n-p$ of them. After that, we may no longer use the common friend since the corresponding component may be empty due to the violation of constraints, and then we may only find the bounds as in Sec. 3.3.5.5.

3.3.5.7 Sensitivity of Function Composition. We want to know how sensitivity w.r.t. components depends on function composition. Given functions $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, and a component ρ of X and a component σ of Y , it may be too difficult to estimate the sensitivity of $g \circ f$ w.r.t. ρ if we know only sensitivity of g w.r.t. σ , since f may map elements of X to very different classes of σ . If x and x' belong to different ρ -classes, it does not mean that $f(x)$ and $f(x')$ belong to different σ classes, unless we make σ dependent on ρ and f . Even in this case, we would require that f is injective to avoid overlapping of σ classes. If we do not put any additional constraints on f , we will need to assume that we know the general sensitivity of g (not w.r.t. some component).

Theorem 3.31. *Let $f : X \rightarrow Y$ be c_f sensitive w.r.t. component ρ . Let $g : Y \rightarrow Z$ be c_g sensitive. Then, the function composition $g \circ f$ is $c_f \cdot c_g$ sensitive w.r.t. ρ , using any definition of sensitivity (A-, B-, C-, D-).*

Proof. We give the proof for different types of sensitivity of f . Let $X_1, X_2 \in X/\rho$.

1. Let $x \in X_1, x' \in X_2, x \stackrel{X_2}{\sim} x'$. If f is A-sensitive w.r.t. ρ , we have $d_Y(f(x), f(x')) \leq c_f \cdot d_X(x, x')$. Since g is c_g sensitive, we have $d_Z(g(y), g(y')) \leq c_g \cdot d_Y(y, y')$ for all $y, y' \in Y$, and hence also for $y = f(x), y' = f(x')$. Putting the two inequalities together, we get $d_Z(gf(x), gf(x')) \leq c_g \cdot d_Y(f(x), f(x')) \leq c_g \cdot c_f \cdot d_X(x, x')$.
2. The proof is similar for B-sensitivity, using $\tilde{d}_X(X_1, X_2)$ instead of $d_X(x, x')$.
3. If $\tilde{d}_Y(\cdot, \cdot)$ is Hausdorff distance, then by Theorem 3.12, C-sensitivity is equivalent to A-sensitivity, and the composition is reduced to composition of A-sensitivity. Let $\tilde{d}_Y(\cdot, \cdot)$ be max-distance. Let $X'_1 \subseteq X_1, X'_2 \subseteq X_2$ be two friendly sets. Denote $Y_1 := f(X'_1), Y_2 := f(X'_2)$. Let $y_1 \in Y_1$ and $y_2 \in Y_2$ be such that $g(y_1) \stackrel{g(Y_2)}{\sim} g(y_2)$ and $\tilde{d}_Z(g(Y_1), g(Y_2)) = d_Z(g(y_1), g(y_2))$. Since g is c_g sensitive, we have $\tilde{d}_Z(g(Y_1), g(Y_2)) = d_Z(g(y_1), g(y_2)) \leq c_g \cdot d_Y(y_1, y_2)$. Since $\tilde{d}_Y(\cdot, \cdot)$ is max-distance, $d_Y(y_1, y_2) \leq \tilde{d}_Y(Y_1, Y_2)$. Putting it together with $\tilde{d}_Y(f(X'_1), f(X'_2)) \leq c_f \cdot \tilde{d}_X(X'_1, X'_2)$, we get $\tilde{d}_Z(gf(X'_1), gf(X'_2)) \leq c_g \cdot c_f \cdot \tilde{d}_X(X'_1, X'_2)$.
4. The proof is similar for D-sensitivity, using $\tilde{d}_X(X_1, X_2)$ instead of $\tilde{d}_X(X'_1, X'_2)$. □

3.3.6 Categorical View of Sensitivity and Distances. Sec. 3.3.5 shows the need to talk about sensitivities of multivariate mappings with respect to a certain argument. We can foresee the need to argue about sensitivities with respect to certain combinations of inputs, and even not knowing precisely which sensitivities are needed in compositions. To keep track of the sensitivities of the workflow components with respect to their different inputs and outputs, we have to add more structure to both distances and sensitivities themselves.

3.3.6.1 Generalized Distances and Sensitivities.

Definition 3.28 (Set of distances). A set V with the following additional algebraic structure is a suitable set of distances:

- V is a commutative monoid, i.e. there is an operation $+$: $V \times V \rightarrow V$ and an element $0 \in V$, such that
 - $(u + v) + w = u + (v + w)$ for all $u, v, w \in V$;
 - $u + v = v + u$ for all $u, v \in V$;
 - $u + 0 = u$ for all $u \in V$.
- V is a partial order, i.e. there is a relation $\leq \subseteq V \times V$, such that
 - $u \leq u$ for all $u \in V$;
 - if $u \leq v$ and $v \leq u$ then $u = v$ for all $u, v \in V$;
 - $u \leq v$ and $v \leq w$ imply $u \leq w$ for all $u, v, w \in V$.
- Addition and order are compatible in the sense that 0 is the least element and all other elements may be considered “positive”. Formally,
 - $0 \leq u$ for all $u \in V$;
 - If $u \leq v$ then $u + w \leq v + w$ for all $u, v, w \in V$.

Obviously, \mathbb{R}_+ with addition and ordering satisfies this definition. For any n , the set of n -tuples \mathbb{R}_+^n also satisfies it, if both addition and ordering are defined componentwise. We expect to such distances with several components to have a number of uses. We use the set of distances for generalizing the notion of a metric space.

Definition 3.29 (Generalized metric space). A generalized metric space is a triple (X, d_X, V_X) , where X is a set, V_X is a set of distances, and the metric $d_X : X \times X \rightarrow V_X$ satisfies

- for all $x, y \in X, d_X(x, y) = 0$ iff $x = y$;

- for all $x, y \in X$, $d_X(x, y) = d_X(y, x)$;
- for all $x, y, z \in Z$, $d_X(x, z) \leq d_X(x, y) + d_X(y, z)$.

We see that this definition coincides with Def. 3.1, except that we are using an arbitrary set of distances V_X instead of \mathbb{R}_+ . The sensitivities of mappings are also generalized. In the following, let $V_X \xrightarrow{m} V_Y$ denote the set of *order-preserving* functions from V_X to V_Y .

Definition 3.30 (Generalized sensitivity). Let (X, d_X, V_X) and (Y, d_Y, V_Y) be generalized metric spaces. Let $f : X \rightarrow Y$ and $c : V_X \xrightarrow{m} V_Y$. We say that the *sensitivity* of f is (at most) c if for all $x, x' \in X$, $d_Y(f(x), f(x')) \leq c(d_X(x, x'))$.

Def. 3.30 indeed generalizes Def. 3.5. Let $c_0 \in \mathbb{R}_+$ be the (ordinary) sensitivity of some mapping $f : X \rightarrow Y$ between (ordinary) metric spaces X and Y with distances d_X and d_Y . If we consider them as generalized metric spaces (X, d_X, \mathbb{R}_+) and (Y, d_Y, \mathbb{R}_+) , then the sensitivity of f is $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, where $c(v) = c_0 \cdot v$ for all $v \in \mathbb{R}_+$.

Proposition 3.1 can be generalized:

Proposition 3.32. Let (X, d_X, V_X) , (Y, d_Y, V_Y) and (Z, d_Z, V_Z) be generalized metric spaces. Let $f : X \rightarrow Y$ have the sensitivity c and $f' : Y \rightarrow Z$ the sensitivity c' . Then $f' \circ f : X \rightarrow Z$ has sensitivity $c' \circ c$.

Proof. Let $x, x' \in X$. Then

$$d_Z(f'(f(x)), f'(f(x'))) \leq c'(d_Y(f(x), f(x'))) \leq c'(c(d_X(x, x'))) . \quad \square$$

In the next section we will see that the generalizations of both Prop. 3.3 and Prop. 3.4 are easy consequences of Prop. 3.32.

For the rest of this section, we make the following notational convention. We will denote generalized metric spaces with capital letters: X, Y, Z , etc. For a generalized metric space X , we let $\text{Carr}(X)$ denote the ‘‘carrier set’’, i.e. the first component of X , $\text{SoD}(X)$ the set of distances used in X , and d_X the metric, mapping pairs of elements of $\text{Carr}(X)$ to $\text{SoD}(X)$.

Before moving on to constructions of generalized metric spaces with more structure, let us mention a quite trivial one. Let X and X' be generalized metric spaces, such that $\text{Carr}(X) = \text{Carr}(X')$. Let X'' be defined by

- $\text{Carr}(X'') = \text{Carr}(X)$;
- $\text{SoD}(X'') = \text{SoD}(X) \times \text{SoD}(X')$, where addition and order are defined componentwise;
- $d_{X''}(x_1, x_2) = (d_X(x_1, x_2), d_{X'}(x_1, x_2))$.

Then X'' is also a generalized metric space. Its metric captures the information present in the metrics of both X and X' .

Generalized distances and metric spaces have been studied in the past, even though the goals have been different. It turns out that any topology over a set arises from a generalized metric over it [51]. This is not the case for ordinary metrics.

3.3.6.2 Limits of Generalized Metric Spaces. Limits in *categories* are certain unique (up to isomorphism) objects, through which certain arrows factorize. They generalize products, as well as inductive datatypes. We start with perhaps the simplest construction of a more complex generalized metric space from simpler ones.

Proposition 3.33. Let X and Y be generalized metric spaces. Then

$$X \times Y := (\text{Carr}(X) \times \text{Carr}(Y), d_{X \times Y}, \text{SoD}(X) \times \text{SoD}(Y))$$

is also a generalized metric space, where the ordering and addition on $\text{SoD}(X) \times \text{SoD}(Y)$ is defined componentwise, and $d_{X \times Y}((x_1, y_1), (x_2, y_2)) = (d_X(x_1, x_2), d_Y(y_1, y_2))$ for all $x_1, x_2 \in \text{Carr}(X)$ and $y_1, y_2 \in \text{Carr}(Y)$.

Proof. We have to verify that $\text{SoD}(X) \times \text{SoD}(Y)$ is a set of distances and $d_{X \times Y}$ is a metric. Both checks are trivial, following directly from the definitions. \square

This simple way of building a product generalized metric space would not have been available if we had considered only ordinary metrics. Proposition 3.33 presents the “correct” way of defining products in the following sense. Consider the category **GM**, defined as follows:

- The objects of **GM** are generalized metric spaces.
- An arrow f from a generalized metric space X to a generalized metric space Y is a pair (k, c) , where $k : \text{Carr}(X) \rightarrow \text{Carr}(Y)$, $c : \text{SoD}(X) \xrightarrow{m} \text{SoD}(Y)$, and k is c -sensitive. We denote the components of an arrow $f = (k, c)$ by $f[f]$ and $c[f]$.

For each object, there exists an identity arrow: the identity map together with identity sensitivity. The composition of arrows is defined componentwise; according to Prop. 3.32, this again gives us an arrow. The category axioms are obviously satisfied.

Proposition 3.34. *The category **GM** has products, they generalize the binary product construction given in Prop. 3.33. In more details: for an index set I , and a set of generalized metric spaces $(X_i)_{i \in I}$, there exists a generalized metric space Y , together with arrows $\pi_i : Y \rightarrow X_i$ for all $i \in I$, such that for any generalized metric space Z together with arrows $f_i : Z \rightarrow X_i$ there exists exactly one arrow $g : Z \rightarrow Y$, such that $f_i = \pi_i \circ g$.*

Proof. Let $\text{Carr}(Y) = \prod_{i \in I} \text{Carr}(X_i)$, $\text{SoD}(Y) = \prod_{i \in I} \text{SoD}(X_i)$ and $d_Y((x_i)_{i \in I}, (x'_i)_{i \in I}) = (d_{X_i}(x_i, x'_i))_{i \in I}$, where \prod denotes the Cartesian product of sets. Define $f[\pi_j]((x_i)_{i \in I}) = x_j$ and $c[\pi_j]((v_i)_{i \in I}) = v_j$, where $x_i \in \text{Carr}(X_i)$ and $v_i \in \text{SoD}(X_i)$.

For the given generalized metric space Z and arrows f_i , define $f[g](z) = (f[f_i](z))_{i \in I} \in \text{Carr}(Y)$ and $c[g](v) = (c[f_i](v))_{i \in I} \in \text{SoD}(Y)$. Clearly, $f_i = \pi_i \circ g$. For any other g' , some of these equalities must not hold. Indeed, if there is some $j \in I$, such that for some $z \in \text{Carr}(Z)$ we have $f[g](z)$ and $f[g'](z)$ differing in j -th component, or for some $v \in \text{SoD}(Z)$ we have $c[g](v)$ and $c[g'](v)$ differing in the j -th component, then $f_j \neq \pi_j \circ g'$. \square

Typically, Y is denoted by $\prod_{i \in I} X_i$ and g by $\langle f_i \rangle_{i \in I}$. For arrows $f : X \rightarrow Y$ and $f' : X' \rightarrow Y'$, one denotes $f \times f' := \langle f \circ \pi_1, f' \circ \pi_2 \rangle : X \times X' \rightarrow Y \times Y'$, where $\pi_1 : X \times X' \rightarrow X$ and $\pi_2 : X \times X' \rightarrow X'$ are given by the product construction.

We can now state the generalizations of Prop. 3.3 and Prop. 3.4.

Proposition 3.35. *Let X, Y, Z, W be generalized metric spaces. Consider the arrows $f : X \rightarrow Y$, $a : Y \rightarrow Z$, $g : X \times Z \rightarrow W$. Then there exists an arrow $h : X \rightarrow W$ with $f[h](x) = f[g](x, f[a](f[f](x)))$ for all $x \in \text{Carr}(X)$ and $c[h](v) = c[g](v, c[a](c[f](v)))$ for all $v \in \text{SoD}(X)$.*

Proof. Take $h = g \circ (id_X \times (a \circ f))$. \square

Prop. 3.35 is specialized to Prop. 3.3 by taking $\text{SoD}(X) = \text{SoD}(Y) = \text{SoD}(Z) = \text{SoD}(W) = \mathbb{R}_+$, $c[f](v) = c \cdot v$, $c[a](v) = v$, and $c[g](v, v') = c' \cdot v + v'$ for all $v, v' \in \mathbb{R}_+$.

Prop. 3.4 does not simplify that much when going to generalized metric spaces. Nevertheless, we can state it as follows. Given a partially ordered set V and a subset $V' \subseteq V$, we let $\text{UB}\{V'\} \subseteq V$ denote the set of all upper bounds of V' : elements that are greater or equal to all elements of V' . A mapping c between sets of distances is *superadditive* if $c(v_1 + v_2) \geq c(v_1) + c(v_2)$ for all v_1, v_2 in the domain of c .

Proposition 3.36. *Let X, Y, Z, W, f, a, g be as in Prop. 3.35. Let ρ and σ be equivalence relations on X that are independent. Let $x \rho x'$ imply $f[f](x) = f[f](x')$, and $x \sigma x'$ imply $f[g](x, z) = f[g](x', z)$ for all $x, x' \in \text{Carr}(X)$ and $z \in \text{Carr}(Z)$. Let $c : \text{SoD}(X) \xrightarrow{m} \text{SoD}(W)$ be any mapping that satisfies the following conditions:*

- for all $v \in \text{SoD}(X)$, $c(v) \in \text{UB}\{c[g](v, 0), c[g](0, c[a](c[f](v)))\}$;
- c is superadditive.

Then there exists an arrow $h : X \rightarrow W$ with $f[h](x) = f[g](x, f[a](f[f](x)))$ for all $x \in \text{Carr}(X)$ and $c[h] = c$.

Proof. Let $x, x' \in \text{Carr}(X)$. Let m and x_0, \dots, x_{2m} be given by the definition of independence of ρ and σ . Similarly to the proof of Prop. 3.4, we get

$$d_W(f[h](x), f[h](x')) \leq \left(\sum_{i=1}^m d_W(f[g](x_{2i-2}, f[a](f[f](x_{2i-2}))), f[g](x_{2i-1}, f[a](f[f](x_{2i-2}))) \right) + \left(\sum_{i=1}^m d_W(f[g](x_{2i}, f[a](f[f](x_{2i-1}))), f[g](x_{2i}, f[a](f[f](x_{2i}))) \right) .$$

Let $v_i = d_X(x_{i-1}, x_i)$. We have

$$d_W(f[g](x_{2i-2}, f[a](f[f](x_{2i-2}))), f[g](x_{2i-1}, f[a](f[f](x_{2i-2})))) \leq c[g](v_{2i-1}, 0) \leq c(v_{2i-1})$$

$$d_W(f[g](x_{2i}, f[a](f[f](x_{2i-1}))), f[g](x_{2i}, f[a](f[f](x_{2i})))) \leq c[g](0, c[a](c[f](v_{2i}))) \leq c(v_{2i})$$

By superadditivity of c ,

$$d_W(f[h](x), f[h](x')) \leq \sum_{i=1}^{2m} c(v_i) \leq c\left(\sum_{i=1}^{2m} v_i\right) = c(d_X(x, x')) . \quad \square$$

Products are one kind of limits. In order to have canonical methods for constructing complex data structures, we are interested also in the existence of other limits. In particular, we are interested in the existence of *equalizers*, because they, together with products, allow to construct all other limits. But as we see below, **GM** does not have equalizers.

Definition 3.31 (Equalizer). Let **C** be a category, X, Y two objects in it, and $f_1, f_2 : X \rightarrow Y$ two arrows between these objects. An *equalizer* of f_1, f_2 is an object Z in **C**, together with an arrow $g : Z \rightarrow X$, such that $f_1 \circ g = f_2 \circ g$. Moreover, for any object Z' and arrow $g' : Z' \rightarrow X$, if $f_1 \circ g' = f_2 \circ g'$, then there exists a unique arrow $h : Z' \rightarrow Z$, such that $g' = g \circ h$.

Consider now the category **GM**, two generalized metric spaces X and Y , and two arrows f_1, f_2 between them. An equalizer $g : Z \rightarrow X$ has to “make equal” both $f[f_1]$ and $f[f_2]$, as well as $c[f_1]$ and $c[f_2]$. It also has to be the largest among those that “make them equal”.

In the category of sets, an equalizer is the subset of X where f_1 and f_2 agree. In **GM**, we must have agreement in both $\text{Carr}(X)$ and $\text{SoD}(X)$. Hence we must have

$$\text{Carr}(Z) \subseteq \{x \in \text{Carr}(X) \mid f[f_1](x) = f[f_2](x)\}$$

$$\text{SoD}(Z) \subseteq \{v \in \text{SoD}(X) \mid c[f_1](v) = c[f_2](v)\} .$$

Also, $\text{SoD}(Z)$ has to be closed with respect to addition.

The following example demonstrates two arrows (with the same domain and codomain) in **GM** that have no equalizer.

Example 3.4. Consider the following generalized metric spaces X and Y , as well as arrows $f_1, f_2 : X \rightarrow Y$:

- $\text{Carr}(X) = \{1_1, 1_2, 2, 3\}$ and $\text{Carr}(Y) = \{1, 2, 3\}$;
- $\text{SoD}(X) = \{0, 1, 2, 3\}$, where ordering is defined by $0 \leq 1 \leq 2 \leq 3$ and addition is defined by $0 + v = v$, and $u + v = 3$ if $u \neq 0 \neq v$;
- $\text{SoD}(Y) = \{0, A, B, C_1, C_2, D\}$, where $0 \leq A \leq B \leq C_1 \leq D$, $B \leq C_2 \leq D$ and the addition is given by the table below.

- d_X and d_Y are given by the tables below.
- $f[f_1] = f[f_2] = \{1_1 \mapsto 1, 1_2 \mapsto 1, 2 \mapsto 2, 3 \mapsto 3\}$
- $c[f_1]$ and $c[f_2]$ are given in the tables below.

+	A	B	C_1	C_2	D
A	B	D	D	D	D
B	D	D	D	D	D
C_1	D	D	D	D	D
C_2	D	D	D	D	D
D	D	D	D	D	D

d_X	1 ₁	1 ₂	2	3
1 ₁	0	1	1	2
1 ₂	1	0	1	2
2	1	1	0	1
3	2	2	1	0
v	0	1	2	3
$c[f_1](v)$	0	A	C_1	D

d_Y	1	2	3	
1	0	A	B	
2	A	0	A	
3	B	A	0	
v	0	1	2	3
$c[f_2](v)$	0	A	C_2	D

If Z and g were an equalizer for f_1, f_2 , then $\text{Carr}(Z) \subseteq \{1_1, 1_2, 2, 3\}$ and $\text{SoD}(Z) \subseteq \{0, 1, 3\}$, while both $f[g]$ and $c[g]$ are identity mappings. Taking $\text{Carr}(Z) = \{1_1, 1_2, 2\}$ and $\text{SoD}(Z) = \{0, 1, 3\}$ would “make f_1 and f_2 equal”. Alternatively, taking $\text{Carr}(Z) = \{2, 3\}$ and $\text{SoD}(Z) = \{0, 1, 3\}$ would “make f_1 and f_2 equal”. But these two options are not isomorphic and there is no option that is larger than both of them.

3.3.6.3 Colimits of Generalized Metric Spaces. Sums are dual for products, and necessary for defining more complex algebraic data structures. A natural definition for binary sums is given by the following proposition.

Proposition 3.37. *Let X and Y be generalized metric spaces. Define Z as follows:*

- $\text{Carr}(Z) = \text{Carr}(X) \uplus \text{Carr}(Y)$, where \uplus denotes the disjoint union, typically defined as $A \uplus B = (\{0\} \times A) \cup (\{1\} \times B)$.
- $\text{SoD}(Z) = \{0, \infty\} \uplus (\text{SoD}(X) \setminus \{0\}) \uplus (\text{SoD}(Y) \setminus \{0\})$, with the order and addition defined as follows:
 - 0 is the smallest and ∞ the largest element. For $v_1, v_2 \in \text{SoD}(X) \setminus \{0\}$ [resp. $\text{SoD}(Y) \setminus \{0\}$], the ordering is the same as in $\text{SoD}(X)$ [resp. $\text{SoD}(Y)$]. If $v_1 \in \text{SoD}(X) \setminus \{0\}$ and $v_2 \in \text{SoD}(Y) \setminus \{0\}$, then v_1 and v_2 are incomparable.
 - $v + 0 = v$ and $v + \infty = \infty$ for all $v \in \text{SoD}(Z)$. If v_1 and v_2 both belong to either $\text{SoD}(X)$ or $\text{SoD}(Y)$, then they are added as in this set. If $v_1 \in \text{SoD}(X) \setminus \{0\}$ and $v_2 \in \text{SoD}(Y) \setminus \{0\}$, then $v_1 + v_2 = \infty$.
- For $x_1, x_2 \in \text{Carr}(X)$, $d_Z(x_1, x_2) = d_X(x_1, x_2)$. Similarly, $d_Z(y_1, y_2) = d_Y(y_1, y_2)$ for $y_1, y_2 \in \text{Carr}(Y)$. If $x \in \text{Carr}(X)$ and $y \in \text{Carr}(Y)$, then $d_Z(x, y) = \infty$.

Then Z is a generalized metric space.

Proof. Similarly to the proof of Prop. 3.33, we have to check that $\text{SoD}(Z)$ is a set of distances and d_Z is a metric. Both checks are trivial. □

We denote Z by $X + Y$. Unfortunately, $X + Y$ is not the coproduct of X and Y in the category **GM**. For (Z, ι_1, ι_2) , where $\iota_1 : X \rightarrow Z$ and $\iota_2 : Y \rightarrow Z$, to be the coproduct of X and Y , there must exist a unique arrow $g : Z \rightarrow W$ for any generalized metric space W and arrows $f_1 : X \rightarrow W$ and $f_2 : Y \rightarrow W$, such that $f_1 = g \circ \iota_1$ and $f_2 = g \circ \iota_2$. Our construction does not fix the value of $c[g](\infty)$ in any way, or even guarantee its existence. Hence, depending on (W, f_1, f_2) , the arrow g might not exist at all, or it might not be unique.

We can get back the coproducts if we put extra restrictions on distances and sensitivities. Namely, let \mathbf{GM}_0^∞ be a subcategory of **GM** containing only such objects and arrows, where

- the sets of distances contain the largest element (denoted ∞);
- the sensitivities preserve both 0 and ∞ .

Note that \mathbf{GM}_0^∞ is closed with respect to products. In \mathbf{GM}_0^∞ , we can define the set of distances of the sum as the disjoint union of the sets of distances of the components, where we do not duplicate the distances 0 and ∞ .

Having sums and products, we can define more complex data structures in \mathbf{GM}_0^∞ , automatically obtaining a canonical set of distances for it. E.g. for a generalized metric space X we can define the generalized metric space $\text{List}(X)$ with the following underlying set and the following set of distances:

- $\text{Carr}(\text{List}(X)) = \text{Carr}(X)^*$;
- $\text{SoD}(\text{List}(X)) = \{0, \infty\} \uplus (\text{SoD}(X)^* \setminus (0^* \cup \infty^*))$.

However, the resulting distance $d_{\text{List}(X)}$ is not particularly interesting. If two lists of elements of $\text{Carr}(X)$ have the same length n , then their distance is defined componentwise, being an element of $\text{SoD}(X)^n$ (with 0^n being identified with the least element 0). If the lists have different lengths, then their distance is ∞ .

We saw that the category \mathbf{GM} had no equalizers. Interestingly, it has coequalizers. The subcategory \mathbf{GM}_0^∞ is closed with respect to the construction we give below, meaning that \mathbf{GM}_0^∞ has both coproducts and colimits. Hence \mathbf{GM}_0^∞ is *cocomplete* — it has all colimits. It remains to be seen what implications this has with respect to the existence of natural definitions of different kinds of algebraic data structures.

Definition 3.32 (Coequalizer). Let \mathbf{C} be a category, X, Y two objects in it, and $f_1, f_2 : X \rightarrow Y$ two arrows between these objects. A *coequalizer* of f_1, f_2 is an object Z in \mathbf{C} , together with an arrow $g : Y \rightarrow Z$, such that $g \circ f_1 = g \circ f_2$. Moreover, for any object Z' and arrow $g' : Y \rightarrow Z'$, if $g' \circ f_1 = g' \circ f_2$, then there exists a unique arrow $h : Z \rightarrow Z'$, such that $g' = h \circ g$.

In the category of sets, the equalizer is a factor set of Y , namely $Z = Y/\rho$, where ρ is the finest equivalence relation that relates $f(x), g(x)$ for all $x \in X$. In the category \mathbf{GM} , a coequalizer needs to make equal not only $f[f_1]$ and $f[f_2]$, but also $c[f_1]$ and $c[f_2]$. To show how this is possible, let us discuss congruences (i.e. structure-preserving equivalences) on sets of distances.

Definition 3.33 (Congruence on a set of distances). Let V be a set of distances. We say that $\sigma \subseteq V \times V$ is a congruence on V , if

- σ is an equivalence relation, i.e. it is reflexive, symmetric and transitive;
- σ is a congruence on the commutative monoid $(V, +, 0)$, i.e. $x \sigma y$ and $w \sigma z$ imply $x + w \sigma y + z$;
- σ preserves ordering, meaning that for any v_0, v_1, \dots, v_{2m} satisfying

$$v_0 \leq v_1 \sigma v_2 \leq v_3 \sigma v_4 \leq \dots \sigma v_{2m} = v_0, \quad (13)$$

the relations $v_i \sigma v_j$ hold for any $i, j \in \{0, \dots, 2m\}$.

Proposition 3.38. Let V be a set of distances and σ a congruence on it. Let $V/\sigma = \{[v]_\sigma \mid v \in V\}$, where $[v]_\sigma = \{u \in V \mid u \sigma v\}$. Define the following algebraic structure on V/σ :

- *Addition:* $[u]_\sigma + [v]_\sigma = [u + v]_\sigma$;
- *Ordering:* $[u]_\sigma \leq [v]_\sigma$ iff there exist $u_0, v_0, \dots, u_m, v_m \in V$, such that

$$u \sigma u_0 \leq v_0 \sigma u_1 \leq v_1 \sigma u_2 \leq v_2 \sigma \dots \leq v_m \sigma v; \quad (14)$$

Then V/σ is a set of distances.

Proof. We have to verify a number of things. First: V/σ is a commutative monoid. This follows from σ being a congruence on $(V, +, 0)$.

Second, V/σ is a partially ordered set. Reflexivity of \leq is immediate. Transitivity is an easy consequence of (14). Consider whether \leq is antisymmetric. Suppose $[u]_\sigma \leq [v]_\sigma$ and $[v]_\sigma \leq [u]_\sigma$. By the definition of ordering, there exist $u_0, v_0, \dots, u_m, v_m$ and $v'_0, u'_0, \dots, v'_n, u'_n$, such that

$$u \sigma u_0 \leq v_0 \sigma u_1 \leq v_1 \sigma \dots \leq v_m \sigma v \sigma v'_1 \leq u'_1 \sigma v'_2 \leq u'_2 \sigma \dots \leq u'_n \sigma u .$$

By the order-preservation of σ we obtain $u \sigma v$ and hence $[u]_\sigma = [v]_\sigma$.

Third, addition and ordering are compatible. We have $[0]_\sigma \leq [v]_\sigma$ because $0 \leq v$. Suppose $[u]_\sigma \leq [v]_\sigma$ and $w \in V$. The relation (14), together with σ being a monoid congruence and the compatibility of addition and ordering on V implies

$$u + w \sigma u_0 + w \leq v_0 + w \sigma u_1 + w \leq v_1 + w \sigma u_2 + w \leq v_2 + w \sigma \cdots \leq v_m + w \sigma v + w,$$

hence $[u + w]_\sigma \leq [v + w]_\sigma$. \square

For equivalences over sets, there exists the smallest equivalence relation that relates certain pairs of values. It turns out that for congruences over distances, the situation is similar. Namely, if σ_i are congruences on a set of distances V , where $i \in I$ for some index set I , then $\sigma = \bigcap_{i \in I} \sigma_i$ is also a congruence on V . Clearly, σ is an equivalence relation and a congruence on the commutative monoid $(V, +, 0)$. Also, if σ preserves ordering, i.e. (13) holds for σ , then it also holds for each σ_i , hence all v_j -s are related by each σ_i , hence they are also related by σ . Now, given pairs $\{(u_j, v_j)\}_{j \in J}$, the smallest congruence σ on V that satisfies $u_j \sigma v_j$ for each $j \in J$, is simply the intersection of all congruences that satisfy these requirements.

In the following, given $U \subseteq V \times V$, let $\text{cl } U$ denote the smallest equivalence relation (if V is an arbitrary set) or the smallest congruence (if V is a set of distances) that contains U . We can now describe the coequalizers.

Proposition 3.39. *The categories \mathbf{GM} and \mathbf{GM}_0^∞ have coequalizers.*

Proof. Let X, Y be two generalized metric spaces and $f_1, f_2 : X \rightarrow Y$ two arrows. Define the equivalence relations ρ over $\text{Carr}(Y)$ and σ over $\text{SoD}(Y)$ as follows:

$$\begin{aligned} \rho &= \text{cl} \{ (f[f_1](x), f[f_2](x)) \mid x \in \text{Carr}(X) \} \\ \sigma &= \text{cl} \{ (c[f_1](u), c[f_2](u)) \mid u \in \text{SoD}(X) \} \cup \{ (d_Y(y_1, y_2), d_Y(y'_1, y'_2)) \mid y_1 \rho y'_1, y_2 \rho y'_2 \} . \end{aligned} \quad (15)$$

Construct a generalized metric space Z as follows:

$$\begin{aligned} \text{Carr}(Z) &= \text{Carr}(Y)/\rho \\ \text{SoD}(Z) &= \text{SoD}(Y)/\sigma \\ d_Z([y_1]_\rho, [y_2]_\rho) &= [d_Y(y_1, y_2)]_\sigma . \end{aligned}$$

The metric d_Z is well-defined because $d_Y(y_1, y_2) \sigma d_Y(y'_1, y'_2)$ whenever $y_1 \rho y'_1$ and $y_2 \rho y'_2$.

Let $g : Y \rightarrow Z$ be such, that $f[g]$ is the natural projection from $\text{Carr}(Y)$ to $\text{Carr}(Y)/\rho$, and $c[g]$ is the natural projection from $\text{SoD}(Y)$ to $\text{SoD}(Y)/\sigma$. Then clearly $g \circ f_1 = g \circ f_2$ due to the definition of ρ and σ .

Suppose that there is an arrow $g' : Y \rightarrow Z'$, such that $g' \circ f_1 = g' \circ f_2$. We can then factor it through g . Define $h : Z \rightarrow Z'$ as follows:

$$\begin{aligned} f[h]([y]_\rho) &= f[g'](y) \\ c[h]([v]_\sigma) &= c[g'](v) . \end{aligned}$$

Then both $f[h]$ and $c[h]$ are well-defined. Indeed, if there exist $y \rho y'$, such that $f[g'](y) \neq f[g'](y')$, then by definition of ρ there must exist y_0, \dots, y_n , such that $y = y_0, y_n = y'$ and for each $i \in \{1, \dots, n\}$ there exists some $x_i \in \text{Carr}(X)$, such that $\{y_{i-1}, y_i\} = \{f_1(x_i), f_2(x_i)\}$. For some i , $f[g'](y_{i-1}) \neq f[g'](y_i)$ and hence also $f[g \circ f_1](x_i) \neq f[g \circ f_2](x_i)$.

Consider the kernel of $c[g']$; it is a congruence on the set of distances $\text{SoD}(Y)$. Clearly, it must contain all pairs named in (15). Hence the kernel is a superset of σ and there are no $v \sigma v'$, such that $c[g'](v) \neq c[g'](v')$.

For any other $h' : Z \rightarrow Z'$, we have $h' \circ g \neq h \circ g$, because $f[g]$ and $c[g]$ are both onto. Hence h is unique and Z, g are a coequalizer.

Finally note that in category \mathbf{GM}_0^∞ , the given construction gives us arrows that preserve the distances 0 and ∞ . Hence we get coequalizers also in \mathbf{GM}_0^∞ . \square

3.3.6.4 Earth Mover's Distance for Probability Distributions. To analyze workflows, we have to understand two things: what happens with distances and sensitivities when we form products, and what happens with then when we form probability distributions. Products have been discussed previously, let us now discuss probabilities.

Let A be an (ordinary) metric space with d_A being the metric on it. Intuitively, the distance d_A is generalized to distributions $\mathcal{D}(A)$ by first thinking about a probability distribution as “dirt” piled over A , with the probability of any particular $a \in A$ showing how much dirt has been put on top of it. The distance between two probability distributions indicates the minimum effort it takes to turn the first pile into the second by moving the dirt around on top of elements of A . Formally:

Proposition 3.40 (Earth mover's distance). *Let $d_A : A \times A \rightarrow \mathbb{R}_+$ be a metric on the set A . Define $\widehat{d}_A : \mathcal{D}(A) \times \mathcal{D}(A) \rightarrow \mathbb{R}_+$ as follows:*

$$\widehat{d}_A(\chi_1, \chi_2) = \inf \left\{ \sum_{a_1, a_2 \in A} \psi(a_1, a_2) \cdot d_A(a_1, a_2) \mid \psi \in \chi_1 \boxtimes \chi_2 \right\} . \quad (16)$$

Then \widehat{d}_A is a metric on $\mathcal{D}(A)$.

Here ψ can be seen as the “plan for moving the dirt from” in order to make one pile similar to the other one.

Proof. We have to show that \widehat{d}_A satisfies the axioms of the metric. First, if $\chi_1 = \chi_2$ then $\widehat{d}_A(\chi_1, \chi_2) = 0$, because then the distribution

$$\psi(a_1, a_2) = \begin{cases} \chi(a_1), & \text{if } a_1 = a_2 \\ 0, & \text{if } a_1 \neq a_2 \end{cases}$$

belongs to $\chi_1 \boxtimes \chi_2$ and all summands in the sum in (16) are 0. *Vice versa*, if $\chi_1 \neq \chi_2$ then any $\psi \in \chi_1 \boxtimes \chi_2$ must put non-zero weight on some pair (a_1, a_2) , where $a_1 \neq a_2$.

Second, \widehat{d}_A is obviously symmetric, because d_A is symmetric.

Third, \widehat{d}_A satisfies the triangle inequality. For any $\psi \in \mathcal{D}(A \times A)$, let $D(\psi) = \sum_{a_1, a_2} \psi(a_1, a_2) \cdot d_A(a_1, a_2)$, where the summation is over all pairs of elements of A . We have defined $\widehat{d}_A(\chi_1, \chi_2)$ as the infimum of $D(\psi)$ over all $\psi \in \chi_1 \boxtimes \chi_2$. Let now $\chi_1, \chi_2, \chi_3 \in \mathcal{D}(A)$. Let $\psi_1 \in \chi_1 \boxtimes \chi_2$ and $\psi_2 \in \chi_2 \boxtimes \chi_3$. Define $\psi_3 \in \mathcal{D}(A \times A)$ by $\psi_3(a_1, a_3) = \sum_{a_2 \in A} \psi_1(a_1, a_2) \cdot \psi_2(a_2, a_3) / \chi_2(a_2)$. Then ψ_3 is a probability distribution. Indeed, $\psi_3(a_1, a_3) \geq 0$ for all $a_1, a_3 \in A$ and

$$\begin{aligned} \sum_{a_1, a_3} \psi_3(a_1, a_3) &= \sum_{a_1, a_2, a_3} \frac{\psi_1(a_1, a_2) \cdot \psi_2(a_2, a_3)}{\chi_2(a_2)} = \sum_{a_1, a_2} \frac{\psi_1(a_1, a_2)}{\chi_2(a_2)} \cdot \sum_{a_3} \psi_2(a_2, a_3) = \\ &= \sum_{a_1, a_2} \frac{\psi_1(a_1, a_2)}{\chi_2(a_2)} \cdot \chi_2(a_2) = \sum_{a_1, a_2} \psi_1(a_1, a_2) = 1 . \end{aligned}$$

Moreover we have

$$\sum_{a_1} \psi_3(a_1, a_3) = \sum_{a_1, a_2} \frac{\psi_1(a_1, a_2) \cdot \psi_2(a_2, a_3)}{\chi_2(a_2)} = \sum_{a_2} \frac{\psi_2(a_2, a_3)}{\chi_2(a_2)} \cdot \sum_{a_1} \psi_1(a_1, a_2) = \sum_{a_2} \psi_2(a_2, a_3) = \chi_3(a_3)$$

and similarly $\sum_{a_3} \psi_3(a_1, a_3) = \chi_1(a_1)$. Hence $\psi_3 \in \chi_1 \boxtimes \chi_3$ and $\widehat{d}_A(\chi_1, \chi_3) \leq D(\psi_3)$. The latter can be bounded as

$$\begin{aligned} D(\psi_3) &= \sum_{a_1, a_2, a_3} \frac{\psi_1(a_1, a_2) \cdot \psi_2(a_2, a_3)}{\chi_2(a_2)} \cdot d_A(a_1, a_3) \leq \\ &= \sum_{a_1, a_2, a_3} \frac{\psi_1(a_1, a_2) \cdot \psi_2(a_2, a_3)}{\chi_2(a_2)} \cdot d_A(a_1, a_2) + \sum_{a_1, a_2, a_3} \frac{\psi_1(a_1, a_2) \cdot \psi_2(a_2, a_3)}{\chi_2(a_2)} \cdot d_A(a_2, a_3) \end{aligned}$$

Here the first sum is equal to

$$\sum_{a_1, a_2} \frac{\psi_1(a_1, a_2)}{\chi_2(a_2)} \cdot d_A(a_1, a_2) \cdot \sum_{a_3} \psi_2(a_2, a_3) = \sum_{a_1, a_2} \psi_1(a_1, a_2) \cdot d_A(a_1, a_2) = D(\psi_1)$$

and similarly the second sum is equal to $D(\psi_2)$. We have $D(\psi_3) \leq D(\psi_1) + D(\psi_2)$ and by taking the infima of all three $D(\cdots)$ -s, we obtain the triangle inequality for \widehat{d}_A . \square

It is possible to generalize earth mover's distance to generalized metric spaces, but only if the sets of distances have some extra structure. We need to be able to multiply the distances with non-negative real numbers; this multiplication has to satisfy certain axioms, similar to axioms of vector spaces.

Definition 3.34 (Distances with multiples). Let V be a set of distances. We say that V has multiples, if there exists an operation $\cdot : \mathbb{R}_+ \times V \rightarrow V$, satisfying the following equalities for all $a, b \in \mathbb{R}_+$ and $u, v \in V$:

$$\begin{array}{ll} a \cdot (b \cdot v) = (ab) \cdot v & a \cdot (u + v) = a \cdot u + a \cdot v \\ 0 \cdot v = 0 & (a + b) \cdot v = a \cdot v + b \cdot v \\ 1 \cdot v = v & u \leq v \Rightarrow a \cdot u \leq a \cdot v \end{array}$$

Note that the axioms also imply $a \cdot v \leq b \cdot v$ if $a \leq b$. Indeed, we have $a \cdot v = a \cdot v + 0 \leq a \cdot v + (b - a) \cdot v = b \cdot v$.

For a partially ordered set V , let $\mathcal{F}(V)$ denote the set of all *upwards closed* subsets of V . I.e. $U \subseteq V$ is an element of $\mathcal{F}(V)$ if $u \in U$ and $u \leq v$ imply $v \in U$ for all $u, v \in V$. For an arbitrary $U \subseteq V$ we let $\uparrow U$ denote the *upwards closure* of U , i.e. the smallest upwards closed set that contains U as a subset.

If V is a set of distances, then we can also define the structure of a set of distances on $\mathcal{F}(V)$. Namely, for $U_1, U_2 \in \mathcal{F}(V)$, we define

- $U_1 + U_2 = \uparrow\{v_1 + v_2 \mid v_1 \in U_1, v_2 \in U_2\}$;
- $U_1 \leq U_2$ iff $U_2 \subseteq U_1$;
- $0_{\mathcal{F}(V)} = \uparrow\{0\} = V$.

It is straightforward to verify that the axioms of a set of distances (Def. 3.28) hold. Also, if V has multiples, then $\mathcal{F}(V)$ has multiples, too, defined simply by $a \cdot U = \uparrow\{c \cdot v \mid v \in U\}$.

Proposition 3.41 (Generalized earth mover's distance). *Let X be a generalized metric space, such that $\text{SoD}(X)$ has multiples. Then Y , defined as follows, is also a generalized metric space.*

- $\text{Carr}(Y) = \mathcal{D}(\text{Carr}(X))$;
- $\text{SoD}(Y) = \mathcal{F}(\text{SoD}(X))$;
- $d_Y(\chi_1, \chi_2) = \uparrow\{\sum_{x, x' \in \text{Carr}(X)} \psi(x, x') \cdot d_X(x, x') \mid \psi \in \chi_1 \boxtimes \chi_2\}$ for all $\chi_1, \chi_2 \in \mathcal{D}(\text{Carr}(X))$.

Proof. We have to verify that the axioms of generalized metric spaces hold. This verification is similar to the proof of Prop. 3.40. \square

The notion that we have defined is indeed a generalization of earth mover's distance. We get it back if we require $\text{SoD}(X) = \mathbb{R}_+$. Note that $\mathcal{F}(\mathbb{R}_+) \cong \mathbb{R}_+$.

In the rest of this section (Sec. 3.3.6.4) we are going to study how the Kleisli composition \circ_{KI} interacts with sensitivities. For this purpose, we introduce the following notation:

- X, Y, Z are three sets;
- V_X, V_Y, V_Z are three sets of distances with multiples;
- d_X, d_Y, d_Z are generalized metrics on X, Y, Z , respectively, ranging over V_X, V_Y, V_Z ;
- $d_{\mathcal{D}(Y)} : \mathcal{D}(Y) \times \mathcal{D}(Y) \rightarrow \mathcal{F}(V_Y)$ and $d_{\mathcal{D}(Z)} : \mathcal{D}(Z) \times \mathcal{D}(Z) \rightarrow \mathcal{F}(V_Z)$ are the generalized earth mover's distances built from d_Y and d_Z .

First we consider how the lifting of a mapping with $\mathcal{D}(\cdot)$ affects its sensitivity. If $c : V_Y \xrightarrow{m} V_Z$ then we say that c is *concave* if $\sum_{i=1}^{\infty} \lambda_i c(v_i) \leq c(\sum_{i=1}^{\infty} \lambda_i v_i)$ for all $\lambda_1, \lambda_2, \dots \in \mathbb{R}_+$ and $v_1, v_2, \dots \in V_Y$ satisfying $\sum_{i=1}^{\infty} \lambda_i = 1$. Note that concavity generalizes superadditivity.

Proposition 3.42. *Let $f : Y \rightarrow Z$ be a mapping between two metric spaces and $c : V_Y \xrightarrow{m} V_Z$ its sensitivity, such that c is concave. Then the sensitivity of the mapping $\mathcal{D}(f) : \mathcal{D}(Y) \rightarrow \mathcal{D}(Z)$ is $\mathcal{F}(c) : \mathcal{F}(V_Y) \rightarrow \mathcal{F}(V_Z)$.*

Before giving the proof, let us recall the meaning of $\mathcal{D}(f)$ and $\mathcal{F}(c)$. These mappings are given by

$$\begin{aligned}\mathcal{D}(f)(\chi) &= \{z \mapsto \sum_{y \in f^{-1}(z)} \chi(y) \mid z \in Z\} \\ \mathcal{F}(c)(U) &= \uparrow\{c(v) \mid v \in U\} .\end{aligned}$$

Proof. Let $\chi, \chi' \in \mathcal{D}(Y)$. Then

$$\begin{aligned}d_{\mathcal{D}(Z)}(\mathcal{D}(f)(\chi), \mathcal{D}(f)(\chi')) &= \uparrow\left\{ \sum_{z, z' \in Z} \psi(z, z') \cdot d_Z(z, z') \mid \psi \in \mathcal{D}(f)(\chi) \boxtimes \mathcal{D}(f)(\chi') \right\} = \\ &\uparrow\left\{ \sum_{y, y' \in Y} \bar{\psi}(y, y') \cdot d_Z(f(y), f(y')) \mid \bar{\psi} \in \chi \boxtimes \chi' \right\} \leq \uparrow\left\{ \sum_{y, y' \in Y} \bar{\psi}(y, y') \cdot c(d_Y(y, y')) \mid \bar{\psi} \in \chi \boxtimes \chi' \right\} \leq \\ &\uparrow\left\{ c\left(\sum_{y, y' \in Y} \bar{\psi}(y, y') \cdot d_Y(y, y') \right) \mid \bar{\psi} \in \chi \boxtimes \chi' \right\} = \uparrow\{c(v) \mid v \in d_{\mathcal{D}(Y)}(\chi, \chi')\} = \mathcal{F}(c)(d_{\mathcal{D}(Y)}(\chi, \chi')) . \quad \square\end{aligned}$$

We continue with studying the sensitivity of $g \circ_{\text{Kl}} f$, where $f : X \rightarrow \mathcal{D}(Y)$ and $g : Y \rightarrow \mathcal{D}(Z)$. By definition, $g \circ_{\text{Kl}} f = \mu \circ \mathcal{D}(g) \circ f$, where $\mu : \mathcal{D}(\mathcal{D}(Z)) \rightarrow \mathcal{D}(Z)$ is the *monadic multiplication* defined by

$$\mu(\Phi) = \{z \mapsto \sum_{\chi \in \mathcal{D}(Z)} \Phi(\chi) \cdot \chi(z) \mid z \in Z\} .$$

Note that $\mathcal{F}(\cdot)$ is also a monad and its monadic multiplication is $\cup : \mathcal{F}(\mathcal{F}(V)) \rightarrow \mathcal{F}(V)$.

Proposition 3.43. *The sensitivity of $\mu : \mathcal{D}(\mathcal{D}(Z)) \rightarrow \mathcal{D}(Z)$ is $\cup : \mathcal{F}(\mathcal{F}(V_Z)) \rightarrow \mathcal{F}(V_Z)$.*

Proof. Let $\Phi, \Phi' \in \mathcal{D}(\mathcal{D}(Z))$. We have

$$\begin{aligned}\cup d_{\mathcal{D}(\mathcal{D}(Z))}(\Phi, \Phi') &= \cup \uparrow\left\{ \sum_{\chi, \chi' \in \mathcal{D}(Z)} \Psi(\chi, \chi') \cdot d_{\mathcal{D}(Z)}(\chi, \chi') \mid \Psi \in \Phi \boxtimes \Phi' \right\} = \\ &\cup \uparrow\left\{ \sum_{\chi, \chi' \in \mathcal{D}(Z)} \uparrow\left\{ \sum_{z, z' \in Z} \Psi(\chi, \chi') \cdot \psi(z, z') \cdot d_Z(z, z') \mid \psi \in \chi \boxtimes \chi' \right\} \mid \Psi \in \Phi \boxtimes \Phi' \right\} \quad (17)\end{aligned}$$

and

$$d_{\mathcal{D}(Z)}(\mu(\Phi), \mu(\Phi')) = \uparrow\left\{ \sum_{z, z' \in Z} \psi(z, z') \cdot d_Z(z, z') \mid \psi \in \mu(\Phi) \boxtimes \mu(\Phi') \right\} . \quad (18)$$

Consider what it means for some $v \in V_Z$ to be an element of (17). There must exist some $\Psi \in \Phi \boxtimes \Phi'$, such that

$$v \in \sum_{\chi, \chi' \in \mathcal{D}(Z)} \uparrow\left\{ \sum_{z, z' \in Z} \Psi(\chi, \chi') \cdot \psi(z, z') \cdot d_Z(z, z') \mid \psi \in \chi \boxtimes \chi' \right\} . \quad (19)$$

This happens if for each $\chi, \chi' \in \mathcal{D}(Z)$, we can pick an element from the set in (19), such that they sum up into something less than or equal to v . Picking an element from this set is equivalent to picking a suitable ψ . Hence we have that v is an element of (17), if there exists $\mathbf{F} : \mathcal{D}(Z)^2 \rightarrow \mathcal{D}(Z \times Z)$, such that $\mathbf{F}(\chi, \chi')$ is an element of $\chi \boxtimes \chi'$ for all $\chi, \chi' \in \mathcal{D}(Z)$ and

$$v \geq \sum_{\chi, \chi' \in \mathcal{D}(Z)} \sum_{z, z' \in Z} \Psi(\chi, \chi') \cdot \mathbf{F}(\chi, \chi')(z, z') \cdot d_Z(z, z') \quad (20)$$

Given such Ψ and \mathbf{F} , define $\psi \in \mathcal{D}(Z \times Z)$ as follows:

$$\psi(z, z') = \sum_{\chi, \chi' \in \mathcal{D}(Z)} \Psi(\chi, \chi') \cdot \mathbf{F}(\chi, \chi')(z, z') .$$

This is indeed a probability distribution, as we can easily verify:

$$\sum_{z, z' \in Z} \psi(z, z') = \sum_{\chi, \chi' \in \mathcal{D}(Z)} \Psi(\chi, \chi') \cdot \sum_{z, z' \in Z} \mathbf{F}(\chi, \chi')(z, z') = \sum_{\chi, \chi' \in \mathcal{D}(Z)} \Psi(\chi, \chi') \cdot 1 = 1 .$$

Consider now the first projection of ψ . Let $z \in Z$.

$$\sum_{z' \in Z} \psi(z, z') = \sum_{\chi, \chi' \in \mathcal{D}(Z)} \Psi(\chi, \chi') \cdot \sum_{z' \in Z} \mathbf{F}(\chi, \chi')(z, z') = \sum_{\chi, \chi' \in \mathcal{D}(Z)} \Psi(\chi, \chi') \cdot \chi(z) = \sum_{\chi \in \mathcal{D}(Z)} \Phi(\chi) \cdot \chi(z) = \mu(\Phi)(z)$$

and symmetrically, the second projection of ψ is equal to $\mu(\Phi')$. Hence $\psi \in \mu(\Phi) \boxtimes \mu(\Phi')$. Comparing (18) and (20), we conclude that v is an element of (18).

Thus $\bigcup d_{\mathcal{D}(\mathcal{D}(Z))}(\Phi, \Phi') \subseteq d_{\mathcal{D}(Z)}(\mu(\Phi), \mu(\Phi'))$, i.e. $d_{\mathcal{D}(Z)}(\mu(\Phi), \mu(\Phi')) \leq \bigcup d_{\mathcal{D}(\mathcal{D}(Z))}(\Phi, \Phi')$ according to the ordering on $\mathcal{F}(V_Z)$. \square

As a simple corollary of Prop. 3.32, Prop. 3.42 and Prop. 3.43 we now get:

Proposition 3.44. *Let $f : X \rightarrow \mathcal{D}(Y)$ and $g : Y \rightarrow \mathcal{D}(Z)$ have sensitivities $c : V_X \xrightarrow{m} \mathcal{F}(V_Y)$ and $c' : V_Y \xrightarrow{m} \mathcal{F}(V_Z)$, respectively, where c' is concave. Then the sensitivity of $g \circ_{\text{Kl}} f$ is $c' \circ_{\text{Kl}} c$.*

Proof. According to the propositions proved before, the sensitivity of $g \circ_{\text{Kl}} f$ is $\bigcup \circ \mathcal{F}(c') \circ c$, which is the Kleisli composition of c and c' . \square

We are left to investigate how the concavity and Kleisli composition interact.

Proposition 3.45. *If $c : V_X \xrightarrow{m} \mathcal{F}(V_Y)$ and $c' : V_Y \xrightarrow{m} \mathcal{F}(V_Z)$ are concave, then $c' \circ_{\text{Kl}} c : V_X \xrightarrow{m} \mathcal{F}(V_Z)$ is also concave.*

Proof. Indeed, let $v_1, v_2, \dots \in V_X$, $\lambda_1, \lambda_2, \dots \in \mathbb{R}_+$, $\sum_{i=1}^{\infty} \lambda_i = 1$. Denote $\mathbb{C} = c(\sum_{i=1}^{\infty} \lambda_i v_i)$ and $\mathbb{C}' = \sum_{i=1}^{\infty} \lambda_i c(v_i)$. By concavity of c , we have $\mathbb{C} \geq \mathbb{C}'$ or $\mathbb{C} \subseteq \mathbb{C}'$. We get (recall that “ \subseteq ” is “ \geq ”)

$$\begin{aligned} (c' \circ_{\text{Kl}} c)\left(\sum_{i=1}^{\infty} \lambda_i v_i\right) &= \bigcup_{u \in \mathbb{C}} c'(u) \subseteq \bigcup_{u \in \mathbb{C}'} c'(u) = \bigcup_{u_1 \in c(v_1)} \bigcup_{u_2 \in c(v_2)} \dots c'\left(\sum_{i=1}^{\infty} \lambda_i u_i\right) \subseteq \\ &\bigcup_{u_1 \in c(v_1)} \bigcup_{u_2 \in c(v_2)} \dots \sum_{i=1}^{\infty} \lambda_i c'(u_i) = \sum_{i=1}^{\infty} \lambda_i \bigcup_{u_i \in c(v_i)} c'(u_i) = \sum_{i=1}^{\infty} \lambda_i (c' \circ_{\text{Kl}} c)(v_i) . \quad \square \end{aligned}$$

3.3.6.5 Set-of-Sets Distance for Probability Distributions. The simple statements of the results in Sec. 3.3.6.4 show that earth mover’s distance is a natural extension of generalized metrics and sensitivities to probability distributions. Unfortunately, the metrics (over probability distributions) relevant for differential privacy are not earth mover’s distances. Intuitively, the earth mover’s distance characterizes the “average” distance between two probability distributions, while differential privacy is concerned with the “worst-case” distance.

For a partially ordered set V , let $\mathcal{I}(V)$ denote the set of all its *downwards closed* non-empty subsets. For a probability distribution $\chi \in \mathcal{D}(X)$, let $\text{supp } \chi \subseteq X$ denote its *support*, i.e. $\text{supp } \chi = \{x \in X \mid \chi(x) > 0\}$. If $f : X \rightarrow \mathcal{D}(Y)$ is a mapping, then let $\bar{f} : \mathcal{D}(X) \rightarrow \mathcal{D}(Y)$ denote its lifting to $\mathcal{D}(X)$; formally, $\bar{f} = \mu \circ \mathcal{D}(f)$. The following proposition gives a suitable lifting of an arbitrary generalized metric to probability distributions and its relation to differential privacy.

Proposition 3.46 (Set-of-sets distance). *Let (X, d_X, V_X) be a generalized metric space. Then $(\mathcal{D}(X), d_{\mathcal{D}(X)}, \mathcal{F}(I(V_X)))$ is also a generalized metric space, where*

$$d_{\mathcal{D}(X)}(\chi_1, \chi_2) = \uparrow\{\downarrow\{d_X(x, x') \mid (x, x') \in \text{supp } \psi\} \mid \psi \in \chi_1 \boxtimes \chi_2\}$$

for all $\chi_1, \chi_2 \in \mathcal{D}(X)$. Moreover, if f is a mapping from a generalized metric space (X, d_X, V_X) to an (ordinary) metric space $(\mathcal{D}(Y), d_{\text{dp}}, \mathbb{R}_+)$ with the sensitivity $\varepsilon : V_X \xrightarrow{m} \mathbb{R}_+$, then the sensitivity of \bar{f} is

$$E(\mathcal{V}) = \min_{U \in \mathcal{V}} \max_{v \in U} \varepsilon(v)$$

for all $\mathcal{V} \in \mathcal{F}(I(V_X))$.

We see that compared to Prop. 3.41, we are now taking the set of all distances, instead of finding their linear combination. Downwards closure corresponds to the ordering of such subsets for the purposes of computing the sensitivity — we are only interested in the maximal elements of each such set. Similarly to $\mathcal{F}(V)$, the set $I(V)$ can also be given the structure of a set of distances, if V has it: addition is again defined pointwise and $0_{I(V)} = \{0\}$. A bit differently, the ordering \leq for $I(V)$ coincides with the subset relation \subseteq .

Proof. It is straightforward to verify that $d_{\mathcal{D}(X)}$ is a generalized metric; the proof follows along the lines of the proof of Prop. 3.40. For investigating the sensitivity of \bar{f} , let $\chi, \chi' \in \mathcal{D}(X)$, $\psi \in \chi \boxtimes \chi'$ and $y \in Y$. Then

$$\begin{aligned} \Pr[\bar{f}(\chi) = y] &= \sum_{x \in X} \chi(x) \cdot \Pr[f(x) = y] = \sum_{x, x' \in X} \psi(x, x') \cdot \Pr[f(x) = y] \leq \\ &\sum_{x, x' \in X} \psi(x, x') \cdot e^{\varepsilon(d_X(x, x'))} \Pr[f(x') = y] \leq \sum_{x, x' \in X} \psi(x, x') \cdot e^{\max_{x \in \text{supp } \psi(\cdot, x')} \varepsilon(d_X(x, x'))} \Pr[f(x') = y] = \\ &\sum_{x' \in X} \chi'(x') \cdot e^{\max_{x \in \text{supp } \psi(\cdot, x')} \varepsilon(d_X(x, x'))} \Pr[f(x') = y] \leq e^{\max_{x, x' \in \text{supp } \psi} \varepsilon(d_X(x, x'))} \cdot \sum_{x' \in X} \chi'(x') \cdot \Pr[f(x') = y] = \\ &e^{\max_{x, x' \in \text{supp } \psi} \varepsilon(d_X(x, x'))} \cdot \Pr[\bar{f}(\chi') = y], \end{aligned}$$

where $\text{supp } \psi(\cdot, x')$ denotes the set of all $x \in X$, such that $\psi(x, x') > 0$. We obtain

$$d_{\text{dp}}(\bar{f}(\chi), \bar{f}(\chi')) = \max_{y \in Y} \left| \ln \frac{\Pr[\bar{f}(\chi') = y]}{\Pr[\bar{f}(\chi) = y]} \right| \leq \min_{\psi \in \chi \boxtimes \chi'} \max_{x, x' \in \text{supp } \psi} \varepsilon(d_X(x, x')) = \min_{U \in d_{\mathcal{D}(X)}(\chi, \chi')} \max_{v \in U} \varepsilon(v) . \quad \square$$

We can again show that the operations with probability distribution monad are matched in the set-of-sets monad. There are analogues to Prop. 3.42, Prop. 3.43 and Prop. 3.44. Let $X, Y, Z, V_X, V_Y, V_Z, d_X, d_Y, d_Z$ be as in Sec. 3.3.6.4. Let $d_{\mathcal{D}(Y)} : \mathcal{D}(Y) \times \mathcal{D}(Y) \rightarrow \mathcal{F}(I(V_Y))$ and $d_{\mathcal{D}(Z)} : \mathcal{D}(Z) \times \mathcal{D}(Z) \rightarrow \mathcal{F}(I(V_Z))$ be set-of-sets distances on $\mathcal{D}(Y)$ and $\mathcal{D}(Z)$.

Proposition 3.47. *Let $f : Y \rightarrow Z$ be a mapping between two metric spaces and $c : V_Y \xrightarrow{m} V_Z$ its sensitivity. Then the sensitivity of the mapping $\mathcal{D}(f) : \mathcal{D}(Y) \rightarrow \mathcal{D}(Z)$ is $\mathcal{F}(I(c))$.*

Let $U \in \mathcal{F}(I(V_Y))$. By definition,

$$\mathcal{F}(I(c))(U) = \uparrow\{\downarrow\{c(v) \mid v \in \mathbf{v}\} \mid \mathbf{v} \in U\} .$$

Proof. Let $\chi, \chi' \in \mathcal{D}(Y)$. Then

$$\begin{aligned} d_{\mathcal{D}(Z)}(\mathcal{D}(f)(\chi), \mathcal{D}(f)(\chi')) &= \uparrow\{\downarrow\{d_Z(z, z') \mid (z, z') \in \text{supp } \psi\} \mid \psi \in \mathcal{D}(f)(\chi) \boxtimes \mathcal{D}(f)(\chi')\} = \\ \uparrow\{\downarrow\{d_Z(f(y), f(y')) \mid (y, y') \in \text{supp } \psi^\circ\} \mid \psi^\circ \in \chi \boxtimes \chi'\} &\leq \uparrow\{\downarrow\{c(d_Y(y, y')) \mid (y, y') \in \text{supp } \psi^\circ\} \mid \psi^\circ \in \chi \boxtimes \chi'\} = \\ \mathcal{F}(I(c))(\uparrow\{\downarrow\{d_Y(y, y') \mid (y, y') \in \text{supp } \psi^\circ\} \mid \psi^\circ \in \chi \boxtimes \chi'\}) &= \mathcal{F}(I(c))(d_{\mathcal{D}(Y)}(\chi, \chi')) . \quad \square \end{aligned}$$

Before studying the sensitivity of $\mu : \mathcal{D}(\mathcal{D}(Z)) \rightarrow \mathcal{D}(Z)$, let us discuss, what is the monadic multiplication for $\mathcal{F}(I(\cdot))$. It must map from $\mathcal{F}(I(\mathcal{F}(I(V_Z))))$ to $\mathcal{F}(I(V_Z))$. For an arbitrary partially ordered set V , consider the mapping $\delta_V : I(\mathcal{F}(V)) \rightarrow \mathcal{F}(I(V))$, defined as follows:

$$\delta_V(\{U_1, U_2, \dots\}) = \uparrow\{\downarrow\{v_1, v_2, \dots\} \mid v_1 \in U_1, v_2 \in U_2, \dots\}$$

Here $U_1, U_2, \dots \in \mathcal{F}(V)$. As the argument of δ_V is downwards closed, we have that if some U_i is an element of this argument, then all supersets of U_i are also elements. We can think of the argument of δ_V as a formal “maximum of minima”. The result of δ_V is a formal “minimum of maxima” that has the same value. I.e. the elements of V are considered as uninterpreted values. We take the minimum of each U_i ; it is natural to consider the sets U_i upwards closed, because the minimum is not affected by extra larger elements. We then take the maximum of all minima to obtain the argument to δ_V . In the result of δ_V , we first take the maxima; these are taken over sets consisting of a single element from each U_i . Finally, we take the minimum of the maxima of all such sets.

The monadic multiplication $\mathbf{M} : \mathcal{F}(I(\mathcal{F}(I(V_Z)))) \rightarrow \mathcal{F}(I(V_Z))$ is

$$\mathbf{M} = \mathcal{F}(\bigcup) \circ \bigcup \circ \mathcal{F}(\delta_{I(V_Z)}) . \quad (21)$$

Let us explain that construction. We have $\mathcal{F}(\delta_{I(V_Z)}) : \mathcal{F}(I(\mathcal{F}(I(V_Z)))) \rightarrow \mathcal{F}(\mathcal{F}(I(I(V_Z))))$, i.e. the first step in the construction of \mathbf{M} “swaps the two middle layers”. The following \bigcup collapses the two $\mathcal{F}(\cdot)$ -layers into one, and the final $\mathcal{F}(\bigcup)$ does the same with the two $I(\cdot)$ -layers.

Proposition 3.48. *The sensitivity of $\mu : \mathcal{D}(\mathcal{D}(Z)) \rightarrow \mathcal{D}(Z)$ is $\mathbf{M} : \mathcal{F}(I(\mathcal{F}(I(V_Z)))) \rightarrow \mathcal{F}(I(V_Z))$.*

Proof. Let \mathbf{F} be the following set of functions $f : \mathcal{D}(Z) \times \mathcal{D}(Z) \rightarrow \mathcal{D}(Z \times Z)$:

$$\mathbf{F} = \{f : \mathcal{D}(Z) \times \mathcal{D}(Z) \rightarrow \mathcal{D}(Z \times Z) \mid \forall(\chi, \chi') : f(\chi, \chi') \in \chi \boxtimes \chi'\} .$$

Let $\Phi, \Phi' \in \mathcal{D}(\mathcal{D}(Z))$. We have

$$\begin{aligned} \mathbf{M}(d_{\mathcal{D}(\mathcal{D}(Z))}(\Phi, \Phi')) &= \mathbf{M}(\uparrow\{\downarrow\{d_{\mathcal{D}(Z)}(\chi, \chi') \mid (\chi, \chi') \in \text{supp } \Psi\} \mid \Psi \in \Phi \boxtimes \Phi'\}) = \\ &= \mathbf{M}(\uparrow\{\downarrow\{\uparrow\{\downarrow\{d_Z(z, z') \mid (z, z') \in \text{supp } \psi\} \mid \psi \in \chi \boxtimes \chi'\} \mid (\chi, \chi') \in \text{supp } \Psi\} \mid \Psi \in \Phi \boxtimes \Phi'\}) = \\ &= (\mathcal{F}(\bigcup) \circ \bigcup)(\uparrow\{\uparrow\{\downarrow\{d_Z(z, z') \mid (z, z') \in \text{supp } f(\chi, \chi')\} \mid (\chi, \chi') \in \text{supp } \Psi\} \mid f \in \mathbf{F}\} \mid \Psi \in \Phi \boxtimes \Phi') = \\ &= \uparrow\{\downarrow\{d_Z(z, z') \mid (\chi, \chi') \in \text{supp } \Psi, (z, z') \in \text{supp } f(\chi, \chi')\} \mid f \in \mathbf{F}, \Psi \in \Phi \boxtimes \Phi'\} \quad (22) \end{aligned}$$

and

$$d_{\mathcal{D}(Z)}(\mu(\Phi), \mu(\Phi')) = \uparrow\{\downarrow\{d_Z(z, z') \mid (z, z') \in \text{supp } \psi\} \mid \psi \in \mu(\Phi) \boxtimes \mu(\Phi')\} . \quad (23)$$

Consider what it means for some $U \in I(V_Z)$ to be an element of (22). There must exist $\Psi \in \Phi \boxtimes \Phi'$ and $f \in \mathbf{F}$ so, that

$$\downarrow\{d_Z(z, z') \mid (\chi, \chi') \in \text{supp } \Psi, (z, z') \in \text{supp } f(\chi, \chi')\} \subseteq U . \quad (24)$$

Let $\psi \in \mathcal{D}(Z \times Z)$ be the following probability distribution:

$$\psi(z, z') = \sum_{\chi, \chi' \in \mathcal{D}(Z)} \Psi(\chi, \chi') \cdot f(\chi, \chi')(z, z') . \quad (25)$$

This is indeed a probability distribution, one can verify it in the same way as in the proof of Prop. 3.43. We can continue in the same way as in the proof of Prop. 3.43 and find that $\psi \in \mu(\Phi) \boxtimes \mu(\Phi')$.

We have

$$\text{supp } \psi \subseteq \bigcup_{(\chi, \chi') \in \text{supp } \Psi} \text{supp } f(\chi, \chi') . \quad (26)$$

Indeed, let $(z, z') \in \text{supp}(\psi)$. By (25), there exist $(\chi, \chi') \in \text{supp } \Psi$, such that $f(\chi, \chi')(z, z') > 0$. This is exactly what is required by the right hand side of (26).

Applying d_Z to both sides of (26) and taking the downwards closure, we obtain

$$\downarrow\{d_Z(z, z') \mid (z, z') \in \text{supp } \psi\} \subseteq \downarrow\{d_Z(z, z') \mid (\chi, \chi') \in \text{supp } \Psi, (z, z') \in \text{supp } f(\chi, \chi')\} .$$

Together with (23) and (24), this implies that U is an element of (23). Thus $\mathbf{M}(d_{\mathcal{D}(\mathcal{D}(Z))}(\Phi, \Phi')) \subseteq d_{\mathcal{D}(Z)}(\mu(\Phi), \mu(\Phi'))$, i.e. $d_{\mathcal{D}(Z)}(\mu(\Phi), \mu(\Phi')) \leq \mathbf{M}(\cup d_{\mathcal{D}(\mathcal{D}(Z))}(\Phi, \Phi'))$ according to the ordering on $\mathcal{F}(\mathcal{I}(V_Z))$. \square

An immediate corollary is

Proposition 3.49. *Let $f : X \rightarrow \mathcal{D}(Y)$ and $g : Y \rightarrow \mathcal{D}(Z)$ have sensitivities $c : V_X \xrightarrow{m} \mathcal{F}(\mathcal{I}(V_Y))$ and $c' : V_Y \xrightarrow{m} \mathcal{F}(\mathcal{I}(V_Z))$, respectively. Then the sensitivity of $g \circ_{\text{Kl}} f$ is $c' \circ_{\text{Kl}} c$.*

Proof. Same as Prop. 3.44. \square

3.3.6.6 Normal Distances. The distance defined in Prop. 3.46 is useful if we have a mapping to $\mathcal{D}(Y)$, where we are interested in the distance d_{dp} . This typically happens at the last stages of a workflow. In other positions, we have a component that implements some mapping $f : X \rightarrow \mathcal{D}(Y)$ with the sensitivity c for some generalized metrics on X and $\mathcal{D}(Y)$. The sensitivity of lifted mapping $\bar{f} : \mathcal{D}(X) \rightarrow \mathcal{D}(Y)$ is somehow related to c ; the exact relationship depends on the generalized metric we consider on $\mathcal{D}(X)$. To analyze a workflow, we need to compose f in parallel with an identity function (actually: the monadic unit) $\eta_Z : Z \rightarrow \mathcal{D}(Z)$, and find the sensitivity of $\bar{f} \times \eta_Z : \mathcal{D}(X \times Z) \rightarrow \mathcal{D}(Y \times Z)$. Given these sensitivities, Prop. 3.32 allows us to find the sensitivity of the entire workflow. We can find this sensitivity if the distances satisfy some mild conditions.

Definition 3.35 (Normal metric). Let $d_{\mathcal{D}(X)}$ be a generalized metric on $\mathcal{D}(X)$, using the set of distances $V_{\mathcal{D}(X)}$. We say that $d_{\mathcal{D}(X)}$ is a *normal metric*, if there exists a set of distances $V_{\mathcal{D}(X)}^b$ and a mapping $d_{\mathcal{D}(X)}^b : \mathcal{D}(X \times X) \rightarrow V_{\mathcal{D}(X)}^b$, such that

$$\begin{aligned} V_{\mathcal{D}(X)} &= \mathcal{F}(V_{\mathcal{D}(X)}^b) \\ d_{\mathcal{D}(X)}(\chi, \chi') &= \uparrow\{d_{\mathcal{D}(X)}^b(\psi) \mid \psi \in \chi \boxtimes \chi'\} \\ d_{\mathcal{D}(X)}^b(\psi_3) &\leq d_{\mathcal{D}(X)}^b(\psi_1) + d_{\mathcal{D}(X)}^b(\psi_2), \end{aligned}$$

where the last inequality holds for all $\psi_1, \psi_2 \in \mathcal{D}(X \times X)$ satisfying $\psi_1 \downarrow_2 = \psi_2 \downarrow_1 =: \phi$ and ψ_3 is defined by $\psi_3(x_1, x_3) = \sum_{x_2 \in X} \psi_1(x_1, x_2) \psi_2(x_2, x_3) / \phi(x_2)$.

If $d_{\mathcal{D}(X)}$ is a normal metric then we say that $(\mathcal{D}(X), d_{\mathcal{D}(X)}, V_{\mathcal{D}(X)})$ is a *normal metric space*.

Clearly, earth mover's distance (Prop 3.41) is a normal distance. Also, the differential privacy distance d_{dp} is normal. Indeed, for any set X , define $d_{\text{dp}}^b : \mathcal{D}(X \times X) \rightarrow \mathbb{R}_+$ simply by $d_{\text{dp}}^b(\psi) = d_{\text{dp}}(\psi \downarrow_1, \psi \downarrow_2)$. The codomain of d_{dp} is $\mathbb{R}_+ \cong \mathcal{F}(\mathbb{R}_+)$ and $d_{\text{dp}}(\chi, \chi') = \inf_{\psi \in \chi \boxtimes \chi'} d_{\text{dp}}^b(\psi)$.

Proposition 3.50 (Product of normal metric spaces). *Let $(\mathcal{D}(X), d_{\mathcal{D}(X)}, \mathcal{F}(V_{\mathcal{D}(X)}^b))$ and $(\mathcal{D}(Y), d_{\mathcal{D}(Y)}, \mathcal{F}(V_{\mathcal{D}(Y)}^b))$ be normal metric spaces. Then $(\mathcal{D}(X \times Y), d_{\mathcal{D}(X \times Y)}, \mathcal{F}(V_{\mathcal{D}(X)}^b \times V_{\mathcal{D}(Y)}^b))$ is also a normal metric space, where*

$$d_{\mathcal{D}(X \times Y)}^b(\psi) = (d_{\mathcal{D}(X)}^b(\psi \downarrow_{1,3}), d_{\mathcal{D}(Y)}^b(\psi \downarrow_{2,4})) .$$

Here $d_{\mathcal{D}(X \times Y)}^b : \mathcal{D}(X \times Y \times X \times Y) \rightarrow V_{\mathcal{D}(X)}^b \times V_{\mathcal{D}(Y)}^b$. I.e. the argument of $d_{\mathcal{D}(X \times Y)}^b$ is a probability distribution over quadruples ψ . In the definition, we project it once to its first and third component (giving an element of $\mathcal{D}(X \times X)$), and once to its second and fourth component (giving an element of $\mathcal{D}(Y \times Y)$).

Proof. It is straightforward to verify that $d_{\mathcal{D}(X \times Y)}$ satisfies the axioms of a generalized metric. \square

For the purposes of the proof of the next proposition, let us define how to apply the function $f : X \rightarrow \mathcal{D}(Z)$ to a probability distribution $\phi \in \mathcal{D}(X \times X)$. There is no single way to apply it, and as a result we obtain a set of probability distributions $\mathcal{A}_f(\phi) \subseteq \mathcal{D}(Z \times Z)$. The elements of this set are obtained by selecting a probability distribution $\psi_{x,x'} \in f(x) \boxtimes f(x')$ for each pair $x, x' \in X$, and then summing them up according to ϕ . Formally,

$$\mathcal{A}_f(\phi) = \left\{ \{(z, z') \mapsto \sum_{x,x' \in X} \phi(x, x') \cdot \psi_{x,x'}(z, z') \mid z, z' \in Z\} \mid x, x' \in X, \psi_{x,x'} \in f(x) \boxtimes f(x') \right\} .$$

Suppose $\chi, \chi' \in \mathcal{D}(X)$ and $\phi \in \chi \boxtimes \chi'$. Then $\mathcal{A}_f(\phi) \subseteq \bar{f}(\chi) \boxtimes \bar{f}(\chi')$. The opposite also holds: if $\psi \in \bar{f}(\chi) \boxtimes \bar{f}(\chi')$, then there exists some $\phi \in \chi \boxtimes \chi'$, such that $\psi \in \mathcal{A}_f(\phi)$. To see this, consider the function $g : X \rightarrow \mathcal{D}(X \times Z)$, defined for all $x, x' \in X$ and $z \in Z$ by

$$\Pr[g(x) = (x', z)] = \begin{cases} \Pr[f(x) = z], & \text{if } x = x' \\ 0, & \text{otherwise.} \end{cases}$$

Clearly, for each $\psi \in \bar{f}(\chi) \boxtimes \bar{f}(\chi') \subseteq \mathcal{D}(Z \times Z)$ there exists some $\psi^\circ \in \bar{g}(\chi) \boxtimes \bar{g}(\chi') \subseteq \mathcal{D}(X \times Z \times X \times Z)$, such that $\psi = \psi^\circ \downarrow_{2,4}$. Now take $\phi = \psi^\circ \downarrow_{1,3}$.

Proposition 3.51 (Parallel composition in normal metric spaces). *Let (X, d_X, V_X) be a generalized metric space. Let $(\mathcal{D}(Y), d_{\mathcal{D}(Y)}, \mathcal{F}(V_Y))$ and $(\mathcal{D}(Z), d_{\mathcal{D}(Z)}, \mathcal{F}(V_Z))$ be normal metric spaces. Let $f : X \rightarrow \mathcal{D}(Z)$ have sensitivity $c : V_X \xrightarrow{m} \mathcal{F}(V_Z)$. Let $(\mathcal{D}(X), d_{\mathcal{D}(X)}, \mathcal{F}(V'_X))$ be a normal metric space. Let the sensitivity of $\bar{f} : \mathcal{D}(X) \rightarrow \mathcal{D}(Z)$ be $\bar{c}(U) = \bigcup_{v \in U} C(v)$ for $U \in \mathcal{F}(V'_X)$ and $C : V'_X \rightarrow \mathcal{F}(V_Z)$. Then the sensitivity of $f \times \eta_Y : \mathcal{D}(X \times Y) \rightarrow \mathcal{D}(Z \times Y)$ is*

$$c'(\mathcal{V}) = \bigcup_{(v,u) \in \mathcal{V}} C(v) \times \uparrow\{u\}$$

for any $\mathcal{V} \in \mathcal{F}(V'_X \times V_Y)$.

Proof. Let $\phi, \phi' \in \mathcal{D}(X \times Y)$ and $(v, u) \in d_{\mathcal{D}(X \times Y)}(\phi, \phi')$. It is sufficient for us to show that $(v^\#, u^\#) \in d_{\mathcal{D}(Z \times Y)}(\bar{f} \times \eta_Y(\phi), \bar{f} \times \eta_Y(\phi'))$ for any $v^\# \in C(v)$ and $u^\# \geq u$. We are thus looking for a $\psi^\# \in \bar{f} \times \eta_Y(\phi) \boxtimes \bar{f} \times \eta_Y(\phi')$, such that $(v^\#, u^\#) \geq d_{\mathcal{D}(Z \times Y)}^\flat(\psi^\#) \in V_Z \times V_Y$.

Consider the set of probability distributions $\bigcup\{\mathcal{A}_f(\rho) \mid \rho \in \phi \downarrow_1 \boxtimes \phi' \downarrow_1\}$. It contains all distributions $\rho^\# \in \mathcal{D}(Z \times Z)$ that are used to define $d_{\mathcal{D}(Z)}(\bar{f}(\phi \downarrow_1), \bar{f}(\phi' \downarrow_1))$. It hence also contains a distribution $\rho^\#$, such that $v^\# \in d_{\mathcal{D}(Z)}^\flat(\rho^\#)$.

We can similarly consider the set of probability distributions $\bigcup\{\mathcal{A}_{f \times \eta_Y}(\psi) \mid \psi \in \phi \boxtimes \phi'\}$. The distributions $\rho^\#$ are obtained from these distributions $\psi^\#$ by projecting onto the first and third components. At the same time, the distributions in $\phi \downarrow_2 \boxtimes \phi' \downarrow_2$ are obtained from the distributions $\psi^\#$ by projecting onto the second and fourth components. There exists a distribution $\psi^\#$ in this set, such that $(\downarrow_{1,3}\psi^\#) = \rho^\#$ and $d_{\mathcal{D}(Y)}^\flat(\psi^\# \downarrow_{2,4}) = u$. This is the distribution $\psi^\#$ that we were looking for. \square

3.3.6.7 Distances and Sensitivities for Streams of Data. For a set X let $\mathcal{S}(X)$ denote the set of *streams* over X , i.e. the set of infinite sequences over X . If we define the lifting of a mapping $f : X \rightarrow Y$ to $\mathcal{S}(f) : \mathcal{S}(X) \rightarrow \mathcal{S}(Y)$ pointwise, then $\mathcal{S}(\cdot)$ is a functor. If (X, d_X, V_X) is a generalized metric space, then $(\mathcal{S}(X), d_{\mathcal{S}(X)}, \mathcal{S}(V_X))$ is also a generalized metric space, where $d_{\mathcal{S}(X)}$ is defined pointwise (i.e. both arguments of $d_{\mathcal{S}(X)}$ are assumed to *move at the same pace*).

A stream $\mathbf{x} \in \mathcal{S}(X)$ can be decomposed to its *head* — an element of X —, and its *tail* — a stream over X again. We write $\mathbf{x} = x : \mathbf{x}'$ to denote that x is the head of \mathbf{x} and \mathbf{x}' is the tail of \mathbf{x} .

If $f : X \rightarrow \mathcal{D}(Y)$ is ε -differentially private, then the differential privacy level of $\mathcal{S}(f) : \mathcal{S}(X) \rightarrow \mathcal{S}(\mathcal{D}(Y)) \subseteq \mathcal{D}(\mathcal{S}(Y))$ depends on how we assume the adversary to differentiate between different streams of values from Y . However, it is natural to assume that the level of distinction between $\mathcal{S}(f)(\mathbf{x}_1)$ and

$\mathcal{S}(f)(\mathbf{x}_2)$ is defined solely in terms of the distances between the corresponding elements in the resulting streams. This justifies using $\mathcal{S}(V_X)$ as the set of distances for $\mathcal{S}(X)$.

A (synchronous) stream processor, transforming a stream $\mathbf{x} \in \mathcal{S}(X)$ to $\mathbf{y} \in \mathcal{S}(Y)$ consists of a set of states S , the initial state $s_0 \in S$, and a mapping $f : S \times X \rightarrow S \times Y$. The stream processor works by reading the next element of \mathbf{x} , applying f to it and the current state, receiving back the next element of \mathbf{y} and the next state. We call it synchronous because it produces exactly one element of \mathbf{y} for each element of \mathbf{x} that it reads. This restriction is natural because we think of the streams we are working with to be indexed by time moments. Formally, given S , s_0 and f , we define $\widetilde{f}^{[s_0]} : \mathcal{S}(X) \rightarrow \mathcal{S}(Y)$ by

$$\widetilde{f}^{[s_0]}(x : \mathbf{x}) = \mathbf{let} (y, s_1) = f(s_0, x) \mathbf{in} y : \widetilde{f}^{[s_1]}(\mathbf{x}) .$$

If (X, d_X, V_X) and (Y, d_Y, V_Y) are generalized metric spaces, then we can also define the structure of a generalized metric space on their Cartesian product. It results in $(X \times Y, d_{X \times Y}, V_X \times V_Y)$, where the ordering and addition on $V_X \times V_Y$ is defined pointwise, and $d_{X \times Y}$ simply applies both d_X and d_Y . If the sensitivity of the mapping f defining a stream processor is $c_f : V_S \times V_X \rightarrow V_S \times V_Y$, then the sensitivity of $\widetilde{f}^{[s_0]}$ is simply $c_f^{[0]}$.

In practice, it will be more likely that the type of our stream processors is $f : S \times X \rightarrow \mathcal{D}(S \times Y)$, i.e. the stream processor is probabilistic and defines a function $\widetilde{f}^{[s_0]} : \mathcal{S}(X) \rightarrow \mathcal{D}(\mathcal{S}(Y))$. If this is the case, and the sensitivity of f is given by the mapping $c_f : V_S \times V_X \rightarrow \mathcal{F}(\mathcal{I}(V_S \times V_Y))$, then the sensitivity c' of $\widetilde{f}^{[s_0]}$ is the following. Define the mapping $c^\# : V_S \times \mathcal{S}(V_X) \rightarrow \mathcal{F}(\mathcal{I}(\mathcal{S}(V_Y)))$ through the following co-inductive construction. In order to compute $c^\#(w_0, x : \mathbf{x})$,

1. let $\mathbf{X} = c(w_0, x) \in \mathcal{F}(\mathcal{I}(V_S \times V_Y))$;
2. let $\mathfrak{X} = \mathcal{F}(\mathcal{I}(g))(\mathbf{X}) \in \mathcal{F}(\mathcal{I}(V_Y \times \mathcal{F}(\mathcal{I}(\mathcal{S}(V_Y)))))$, where $g : V_S \times V_Y \rightarrow V_Y \times \mathcal{F}(\mathcal{I}(\mathcal{S}(V_Y)))$ is defined by $g(v_s, v_y) = (v_y, c^\#(v_s, \mathbf{x}))$;
3. let $\mathfrak{Y} = \mathcal{F}(\mathcal{I}(h))(\mathfrak{X}) \in \mathcal{F}(\mathcal{I}(\mathcal{F}(\mathcal{I}(\mathcal{S}(V_Y)))))$, where $h : V_Y \times \mathcal{F}(\mathcal{I}(\mathcal{S}(V_Y))) \rightarrow \mathcal{F}(\mathcal{I}(\mathcal{S}(V_Y)))$ is defined by $h(v_y, \mathbf{Y}) = \mathcal{F}(\mathcal{I}(v_y : \cdot))(\mathbf{Y})$, i.e. h prepends v_y to all streams in \mathbf{Y} ;
4. return $\mu(\mathfrak{Y})$, where μ is the monad multiplication for $\mathcal{F}(\mathcal{I}(\cdot))$.

Then $c' = c^\#(0, \cdot)$. The construction of $c^\#$ is well-founded because it is recursively applied only to components of its arguments.

If c is sufficiently simple (e.g. it is guaranteed to return only finitely generated sets of finitely generated sets), then the construction of c' can be executed. In this case, we have a means to analyse workflows processing streams, where each stream processor f only has state with finitely many variables taking values in \mathbb{R} . The size of the entire stream, however, is not bounded.

3.3.7 Derivative Sensitivity for Row Multiplicities. In this section, we will study the following problem. Let $X = (\mathbb{R}^n, d)$ be a metric space, with d being the ℓ_1 -distance (Definition 3.15). Given a function $f : X \rightarrow \mathbb{R}$, and the desired level of differential privacy (DP), find a noise distribution $\eta : X \rightarrow \mathcal{D}(\mathbb{R})$, such that the probabilistic mapping $g(\vec{x}) := f(\vec{x}) + \eta(\vec{x})$ has that level of DP, and η does not add overly much noise.

In Sec. 3.3.1.2, we discussed that the noise magnitude depends on function sensitivity, i.e. how much a change in the function input affects the function output. There exist different flavours of function sensitivity. The *global* sensitivity is defined for all inputs in the function domain.

Definition 3.36 (global sensitivity). For $f : X \rightarrow Y$, the *global sensitivity* of f is

$$GS_f = \max_{x, x' \in X} \frac{d_Y(f(x), f(x'))}{d_X(x, x')} .$$

Note that Def. 3.36 is equivalent to Def. 3.5.

Since noise is always added to the output of a query that is applied to a particular state of the database, and some state may require less noise than the other, the noise magnitude may depend on the data to which the function is applied. The *local* sensitivity of a function depends on the actual data instance.

Definition 3.37 (local sensitivity). For $f : X \rightarrow Y$, an integer-valued metric $d_X : X \times X \rightarrow \mathbb{N}$, and $x \in X$, the *local sensitivity* of f at x is

$$LS_f(x) = \max_{x' \in X: d_X(x, x')=1} d_Y(f(x), f(x')) .$$

In this work, we choose local smoothed sensitivity [52] as a suitable basis for choosing the magnitude of the noise. We apply smoothing not to the local sensitivity (although that would work, too), but to its analogue for continuous functions, which we define here and call *derivative sensitivity*.

Derivative sensitivity has to be applied to the semantics of SQL queries. Hence we define a continuous semantics for it, which matches the natural one if applied to databases where records have integer multiplicities. We also show how to compute the smoothed derivative sensitivity of a SQL query.

Let us recall some notions from Sec. 3.3.1. For a set X , let $\mathcal{D}(X)$ denote the set of all probability distributions over it. An element $\chi \in \mathcal{D}(X)$ is a mapping from X to $[0, 1]$, and we use the notation $\chi(x)$ to denote the probability weight that χ assigns to $x \in X$. For distributions $\chi, \chi' \in \mathcal{D}(X)$, their *differential privacy distance* (or *DP-distance*) is defined by

$$d_{dp}(\chi, \chi') = \sup_{x \in X} \left| \frac{\ln \chi(x)}{\ln \chi'(x)} \right| .$$

Given a mapping f from one metric space (X, d_X) to another metric space (Y, d_Y) , we say that the *sensitivity* of f is (at most) $c \in \mathbb{R}_+$, if for all $x, x' \in X$, the inequality $d_Y(f(x), f(x')) \leq c \cdot d_X(x, x')$ holds. A mapping $f : X \rightarrow \mathcal{D}(Y)$ (where X is a metric space) is ϵ -*differentially private*, if it is ϵ -sensitive from (X, d_X) to $(\mathcal{D}(Y), d_{dp})$.

In Section 3.3.6, we studied the possibility of distances being not non-negative real numbers, but something more general. One relevant example is the set $\mathcal{F}(\mathbb{R}_+ \times \mathbb{R}_+)$, the set of all *upwards closed* sets of pairs of non-negative real numbers. The upwards closedness means that if $Z \in \mathcal{F}(\mathbb{R}_+ \times \mathbb{R}_+)$, $(\epsilon, \delta) \in Z$, $\epsilon' \geq \epsilon$, and $\delta' \geq \delta$, then $(\epsilon', \delta') \in Z$ as well. One can define ordering and addition on the elements of $\mathcal{F}(\mathbb{R}_+ \times \mathbb{R}_+)$, allowing its elements to be treated as distances and the triangle inequalities to be stated.

This example is significant in defining (ϵ, δ) -differential privacy. Let us recall the definition of the latter.

Definition 3.38 ([36]). Let X be a metric space and $f : X \rightarrow \mathcal{D}(Y)$. The mapping f is (ϵ, δ) -*differentially private* if for all $Y' \subseteq Y$, and for all x, x' , where $d_X(x, x') = 1$, the following inequality holds:

$$\Pr[f(x) \in Y'] \leq e^\epsilon \Pr[f(x') \in Y'] + \delta . \quad (27)$$

First, if $\chi, \chi' \in \mathcal{D}(X)$, then we define the distance

$$d_{DP}(\chi, \chi') = \bigcap_{Y' \subseteq X} \left\{ (\epsilon, \delta) \mid \begin{array}{l} \Pr[x \in Y' \mid x \leftarrow \chi] \leq e^\epsilon (\Pr[x \in Y' \mid x \leftarrow \chi'] + \delta) \\ \Pr[x \in Y' \mid x \leftarrow \chi'] \leq e^\epsilon (\Pr[x \in Y' \mid x \leftarrow \chi] + \delta) \end{array} \right\} \quad (28)$$

Clearly, $d_{DP}(\chi, \chi') \in \mathcal{F}(\mathbb{R}_+ \times \mathbb{R}_+)$. Now, a mapping $f : X \rightarrow \mathcal{D}(Y)$ from a metric space X is (ϵ, δ) -differentially private, if it is $\uparrow\{(\epsilon, \delta)\}$ -sensitive for the distance d_{DP} being used on $\mathcal{D}(Y)$.

Let us also state how ordering and addition on $\mathcal{F}(\mathbb{R}_+ \times \mathbb{R}_+)$ is defined. Let $Z_1, Z_2 \in \mathcal{F}(\mathbb{R}_+ \times \mathbb{R}_+)$. We have $Z_1 \leq Z_2$ iff $Z_2 \subseteq Z_1$. In this way, the entire set $\mathbb{R}_+ \times \mathbb{R}_+$ is the least element, corresponding to two distributions being equal. Indeed, this equality is expressed by $(0, 0) \in \mathbb{R}_+ \times \mathbb{R}_+$. Such definition of ordering is the standard one. The addition is also standard:

$$Z_1 + Z_2 = \{(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2) \mid (\epsilon_1, \delta_1) \in Z_1, (\epsilon_2, \delta_2) \in Z_2\} .$$

If Z_1 and Z_2 are upwards closed, then so is $Z_1 + Z_2$ due to the continuousness of \mathbb{R} . It is also easy to see that the operation $+$ is associative, commutative, has the zero element $\mathbb{R}_+ \times \mathbb{R}_+$ and is compatible with ordering (meaning that $Z_1 \leq Z_2$ implies $Z_1 + Z_3 \leq Z_2 + Z_3$ for any Z_3).

The following proposition shows that d_{DP} satisfies the triangle inequality.

Proposition 3.52. Let $\chi_1, \chi_2, \chi_3 \in \mathcal{D}(X)$. Then $d_{\text{DP}}(\chi_1, \chi_3) \leq d_{\text{DP}}(\chi_1, \chi_2) + d_{\text{DP}}(\chi_2, \chi_3)$.

Proof. Let $(\epsilon_1, \delta_1) \in d_{\text{DP}}(\chi_1, \chi_2)$ and $(\epsilon_2, \delta_2) \in d_{\text{DP}}(\chi_2, \chi_3)$. According to the definition of $+$, the pair $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ is a member of $d_{\text{DP}}(\chi_1, \chi_2) + d_{\text{DP}}(\chi_2, \chi_3)$. We have to show that it is also a member of $d_{\text{DP}}(\chi_1, \chi_3)$ according to (28). Let $X' \subseteq X$. Then

$$\begin{aligned} \Pr[x \in X' | x \leftarrow \chi_1] &\leq e^{\epsilon_1}(\Pr[x \in X' | x \leftarrow \chi_2] + \delta_1) \leq e^{\epsilon_1}((e^{\epsilon_2}(\Pr[x \in X' | x \leftarrow \chi_3] + \delta_2) + \delta_1) = \\ &e^{\epsilon_1 + \epsilon_2} \Pr[x \in X' | x \leftarrow \chi_3] + e^{\epsilon_1 + \epsilon_2} \delta_2 + e^{\epsilon_1} \delta_1 \leq e^{\epsilon_1 + \epsilon_2} (\Pr[x \in X' | x \leftarrow \chi_3] + \delta_2 + \delta_1) \end{aligned}$$

as necessary. \square

Let us show the relationship between d_{DP} and (ϵ, δ) -differential privacy. Note that Definition 3.38 is symmetric with respect to x and x' , because their role may be swapped.

The formula (27) differs from (28) in one important aspect. Namely, in (28), the quantity δ is multiplied with e^ϵ , while in (27), it is not. While the difference of the factor e^ϵ seems small in first glance, it is not if we start considering “group privacy”, i.e. distances (in X) different from 1. Let $d_X(x, x') = L$. If $f : X \rightarrow \mathcal{D}(Y)$ is $\uparrow\{(\epsilon, \delta)\}$ -sensitive with respect to the distance d_{DP} on $\mathcal{D}(Y)$, then we know that $(L\epsilon, L\delta) \in d_{\text{DP}}(f(x), f(x'))$. But if f is (ϵ, δ) -differentially private, then we only get

$$\Pr[f(x) \in Y'] \leq e^{L\epsilon} \Pr[f(x') \in Y'] + \frac{e^{L\epsilon} - 1}{e^\epsilon - 1} \delta$$

from Definition 3.38.

It is not difficult to show that if we do not multiply δ with e^ϵ , then d_{DP} is no longer a distance; in particular, it would not satisfy the triangle inequality. For example, let us pick

$$\chi_1 = \text{Ber}(0.01) \quad \chi_2 = \text{Ber}(0.03) \quad \chi_3 = \text{Ber}(0.07) \quad \epsilon = \ln 2 \quad \delta = 0.01 .$$

Here $\text{Ber}(p)$ is the *Bernoulli distribution*. It returns 1 with probability p and 0 with probability $1 - p$. We have $(\epsilon, \delta) \in d_{\text{DP}}(\chi_1, \chi_2)$ and also $(\epsilon, \delta) \in d_{\text{DP}}(\chi_2, \chi_3)$, but not $(2\epsilon, 2\delta) \in d_{\text{DP}}(\chi_1, \chi_3)$. Indeed,

$$\begin{aligned} \Pr[x = 1 | x \leftarrow \chi_2] &= 0.03 = 2 \cdot 0.01 + 0.01 = e^\epsilon \cdot \Pr[x = 1 | x \leftarrow \chi_1] + \delta \\ \Pr[x = 1 | x \leftarrow \chi_3] &= 0.07 = 2 \cdot 0.03 + 0.01 = e^\epsilon \cdot \Pr[x = 1 | x \leftarrow \chi_2] + \delta \\ \Pr[x = 1 | x \leftarrow \chi_3] &= 0.07 > 4 \cdot 0.01 + 0.02 = e^{2\epsilon} \cdot \Pr[x = 1 | x \leftarrow \chi_1] + 2\delta . \end{aligned}$$

3.3.7.1 Sensitivity for Continuous Functions. For differentiable functions, the notions of sensitivity and derivatives are very tightly related.

Definition 3.39. Let $f : X \rightarrow \mathbb{R}$. The *derivative sensitivity* of f is the following mapping from X to \mathbb{R}_+ , where \mathbb{R}_+ denotes the set of all non-negative real numbers:

$$\text{DS}_f(\vec{x}) = \max_i \left| \frac{\partial f}{\partial x_i}(\vec{x}) \right| .$$

Here x_i denotes the i -th component of the vector of variables \vec{x} .

For these functions, the derivative sensitivity can be used to obtain results very similar to [52], where local sensitivity was used instead. As we discuss below, derivative sensitivity is usually simpler to use, as long as the mapping f is differentiable.

Definition 3.40 ([52]). Let $p : X \rightarrow \mathbb{R}$ and $\beta \in \mathbb{R}$. The mapping p is β -smooth, if $p(\vec{x}) \leq e^{\beta \cdot d(\vec{x}, \vec{x}')} \cdot p(\vec{x}')$ for all $\vec{x}, \vec{x}' \in X$.

As we see below, crucial in selecting the amount of noise to be added to $f(\vec{x})$ is the knowledge of a β -smooth upper bound on the derivative sensitivity of f . We let c denote such a bound. We consider the same noise distributions as in [52]. For a parameter $\gamma \in \mathbb{R}_+$, $\gamma > 1$, the *generalized Cauchy distribution* $GenCauchy(\gamma) \in \mathcal{D}(\mathbb{R})$ is given by the proportionality

$$GenCauchy(\gamma)(x) \propto \frac{1}{1 + |x|^\gamma}$$

(“usual” Cauchy distribution is obtained for $\gamma = 2$). Noise distributed by generalized Cauchy distribution, weighed by a smooth upper bound on the derivative sensitivity of f , allows us to achieve ϵ -DP.

Theorem 3.53. *Let $\gamma, b, \beta \in \mathbb{R}_+$, $\gamma > 1$. Let $\epsilon = (\gamma + 1)(b + \beta)$. Let η be a random variable distributed according to $GenCauchy(\gamma)$. Let c be a β -smooth upper bound on DS_f for a function $f : X \rightarrow \mathbb{R}$. Then $g(\vec{x}) : f(\vec{x}) + \frac{c(\vec{x})}{b} \cdot \eta$ is ϵ -differentially private.*

Proof. Let $\eta \sim GenCauchy(\gamma)$. The generalized Cauchy distribution is relatively stable under shifts and stretchings, satisfying the following inequalities for all $a_1, a_2, c_1, c_2 \in \mathbb{R}$, which we state here without proof [52]:

$$\begin{aligned} d_{dp}(a_1 + c_1 \cdot \eta, a_2 + c_1 \cdot \eta) &\leq (\gamma + 1) \cdot \left| \frac{a_2 - a_1}{c_1} \right| \\ d_{dp}(c_1 \cdot \eta, c_2 \cdot \eta) &\leq (\gamma + 1) \cdot \left| \ln \frac{c_2}{c_1} \right|. \end{aligned}$$

The combination of these two inequalities gives

$$d_{dp}(a_1 + c_1 \cdot \eta, a_2 + c_2 \cdot \eta) \leq (\gamma + 1) \cdot \left(\frac{|a_2 - a_1|}{\max\{|c_1|, |c_2|\}} + \left| \ln \frac{c_2}{c_1} \right| \right). \quad (29)$$

Let $\vec{x}, \vec{x}' \in X$. Suppose that they differ only in the i_0 -th coordinate. W.l.o.g. assume that $x'_{i_0} \geq x_{i_0}$. Denote $L = x'_{i_0} - x_{i_0}$. We have to show that $d_{dp}(g(\vec{x}'), g(\vec{x})) \leq \epsilon L = (\gamma + 1)(b + \beta)L$.

We can substitute the definition of g into the left side of the desired inequality above, and using the inequality (29) and the definition of smoothness, obtain

$$\begin{aligned} d_{dp}(g(\vec{x}), g(\vec{x}')) &= d_{dp}\left(f(\vec{x}) + \frac{c(\vec{x})}{b} \cdot \eta, f(\vec{x}') + \frac{c(\vec{x}')}{b} \cdot \eta\right) \leq \\ &(\gamma + 1) \cdot \left(b \cdot \frac{|f(\vec{x}') - f(\vec{x})|}{|c(\vec{x})|} + \left| \ln \frac{c(\vec{x}')}{c(\vec{x})} \right| \right) \leq (\gamma + 1) \cdot \left(b \cdot \frac{|f(\vec{x}') - f(\vec{x})|}{|c(\vec{x})|} + \beta L \right) \end{aligned}$$

Unfortunately, we cannot directly bound $|f(\vec{x}') - f(\vec{x})|/|c(\vec{x})|$ with L . Instead, if we let $\vec{y}[v]$ denote the tuple \vec{x} , where the i_0 -th component is replaced with v , then we can only claim (using the mean value theorem), that there exists some v_0 in the segment (x_{i_0}, x'_{i_0}) satisfying

$$|f(\vec{x}') - f(\vec{x})| = \left| \frac{\partial f}{\partial x_{i_0}}(\vec{y}[v_0]) \right| \cdot (x'_{i_0} - x_{i_0}) \leq |c(\vec{y}[v_0])| \cdot L,$$

where the last inequality is due to c being an upper bound on the derivative sensitivity of f . However, by using this claim many times, we obtain the necessary inequality as follows. Let $n \in \mathbb{N}$ be arbitrary. Let $v_0 = x_{i_0}$, $v_n = x'_{i_0}$ and $v_i = ((n - i) \cdot x_{i_0} + i \cdot x'_{i_0})/n$, i.e. the values v_0, \dots, v_n are evenly distributed from x_{i_0} to x'_{i_0} . Again, the mean value theorem implies that there exist t_1, \dots, t_n with $v_{i-1} \leq t_i \leq v_i$, satisfying

$$|f(\vec{y}[v_i]) - f(\vec{y}[v_{i-1}])| = \left| \frac{\partial f}{\partial x_{i_0}}(\vec{y}[t_i]) \right| \cdot (v_i - v_{i-1}) \leq |c(\vec{y}[t_i])| \cdot \frac{L}{n} \leq e^{\beta L/n} \cdot |c(\vec{y}[v_{i-1}])| \cdot \frac{L}{n}$$

for all $i \in \{1, \dots, n\}$. Here the last inequality follows from the β -smoothness of c . We can use these claims together with the triangle inequality and obtain

$$d_{\text{dp}}(g(\vec{x}), g(\vec{x}')) \leq \sum_{i=1}^n d_{\text{dp}}(g(\vec{y}[v_{i-1}]), g(\vec{y}[v_i])) = \sum_{i=1}^n d_{\text{dp}}\left(f(\vec{y}[v_{i-1}]) + \frac{c(\vec{y}[v_{i-1}])}{b} \cdot \eta, f(\vec{y}[v_i]) + \frac{c(\vec{y}[v_i])}{b} \cdot \eta\right) \leq$$

$$(\gamma + 1) \sum_{i=1}^n \left(b \cdot \frac{|f(\vec{y}[v_i]) - f(\vec{y}[v_{i-1}])|}{|c(\vec{y}[v_{i-1}])|} + \frac{\beta L}{n} \right) \leq (\gamma + 1) \sum_{i=1}^n \left(b \cdot e^{\beta L/n} \cdot \frac{L}{n} + \frac{\beta L}{n} \right) = (\gamma + 1)(be^{\beta L/n} + \beta)L .$$

This inequality holds for any $n \in \mathbb{N}$. If $n \rightarrow \infty$ then $e^{\beta L/n} \rightarrow 1$ and we obtain the inequality that we had to show.

If \vec{x} and \vec{x}' differ in more than one coordinate, then we can transform \vec{x} to \vec{x}' by changing one coordinate at a time, and using the triangle inequality. \square

Generalized Cauchy distributions have heavy tails, meaning that using the mechanism of Theorem 3.53 will in general introduce a lot of noise. Less noise is possible if we resort to (ϵ, δ) -DP. In this case, we may add noise from the usual Laplacian distribution, which has light tails. The following theorem makes use of the distribution $Lap(1) \in \mathcal{D}(\mathbb{R})$, defined by $Lap(1)(x) \propto e^{-|x|}$. We first have to state the results about the self-similarity of $Lap(1)$ under shifting and stretching.

Lemma 3.54. *Let $\eta \sim Lap(1)$. Let $a_1, a_2 \in \mathbb{R}$, $c_1, c_2 \in \mathbb{R}_+$, $c_1 \leq c_2$. Define $\beta = \ln(c_2/c_1)$ and let $\epsilon \geq \beta$. Let $\delta \geq e^{-\epsilon - (\epsilon + \beta)/(e^\beta - 1)}$. Then the following holds.*

$$\left(\frac{|a_2 - a_1|}{c_1}, 0 \right) \in d_{\text{DP}}(a_1 + c_1 \cdot \eta, a_2 + c_1 \cdot \eta)$$

$$(\epsilon, \delta) \in d_{\text{DP}}(c_1 \cdot \eta, c_2 \cdot \eta) .$$

Proof. The probability density functions (PDF) and the cumulative density functions (CDF) of the distributions named above are the following:

$$\text{PDF}_{c_1 \cdot \eta}(x) = \frac{1}{2c_1} e^{-|x|/c_1} \qquad \text{PDF}_{a_1 + c_1 \cdot \eta}(x) = \frac{1}{2c_1} e^{-|x - a_1|/c_1}$$

$$\text{PDF}_{c_2 \cdot \eta}(x) = \frac{1}{2c_2} e^{-|x|/c_2} \qquad \text{PDF}_{a_2 + c_1 \cdot \eta}(x) = \frac{1}{2c_1} e^{-|x - a_2|/c_1}$$

and

$$\text{CDF}_{c_1 \cdot \eta}(x) = \begin{cases} e^{x/c_1}/2, & \text{if } x < 0 \\ 1 - e^{-x/c_1}/2, & \text{if } x \geq 0 \end{cases} \qquad \text{CDF}_{c_2 \cdot \eta}(x) = \begin{cases} e^{x/c_2}/2, & \text{if } x < 0 \\ 1 - e^{-x/c_2}/2, & \text{if } x \geq 0 \end{cases}$$

The first claim of the lemma is shown by

$$\max_{x \in \mathbb{R}} \left| \ln \frac{\text{PDF}_{a_1 + c_1 \cdot \eta}(x)}{\text{PDF}_{a_2 + c_1 \cdot \eta}(x)} \right| = \max_{x \in \mathbb{R}} \left| \ln \frac{e^{-|x - a_1|/c_1}}{e^{-|x - a_2|/c_1}} \right| \leq \ln e^{|a_1 - a_2|/c_1} = \frac{|a_2 - a_1|}{c_1},$$

showing that $\frac{|a_2 - a_1|}{c_1} \geq d_{\text{dp}}(a_1 + c_1 \cdot \eta, a_2 + c_1 \cdot \eta)$. To show the second claim, consider the following function f :

$$f(x) = \left| \ln \frac{\text{PDF}_{c_1 \cdot \eta}(x)}{\text{PDF}_{c_2 \cdot \eta}(x)} \right|.$$

We are interested in the set of x -s that satisfy $f(x) \leq \epsilon$. We have

$$f(x) = \left| \ln \left(\frac{c_2}{c_1} \cdot e^{|x|/c_2 - |x|/c_1} \right) \right| = \left| \beta - \frac{c_2 - c_1}{c_1 c_2} |x| \right| .$$

The condition $f(x) \leq \epsilon$ is equivalent to

$$|x| \leq (\epsilon + \beta) \frac{c_1 c_2}{c_2 - c_1} . \quad (30)$$

To obtain the distance $\uparrow\{(\epsilon, \delta)\}$, it is sufficient to take δ equal to $e^{-\epsilon}$ times the probability that either x , when sampled according to either $c_2 \cdot \eta$ or $c_1 \cdot \eta$, does not satisfy (30). This probability is larger for $c_2 \cdot \eta$ because $c_2 \geq c_1$. Let us compute this probability.

$$\Pr[x < 0 \wedge f(x) > \epsilon | x \leftarrow c_2 \cdot \eta] = \frac{1}{2} e^{(\epsilon+\beta) \frac{c_1}{c_2-c_1}} = \frac{1}{2} e^{(\epsilon+\beta) \frac{1}{\beta-1}} .$$

The probability that we are looking for is twice the quantity above. Multiplying it with $e^{-\epsilon}$ gives us the statement of the lemma. \square

The next lemma provides a more coarse, but simpler upper bound for the DP-distance between stretched versions of the Laplace distribution.

Lemma 3.55. *Let $\eta \sim \text{Lap}(1)$. Let $c_1, c_2 \in \mathbb{R}_+$, $c_1 \leq c_2$. Define $\beta = \ln(c_2/c_1)$. Let $\epsilon \geq \beta$. Let $k = 1 + \epsilon/\beta$. Then $(\epsilon, e^{-k}) \in d_{\text{DP}}(c_1 \cdot \eta, c_2 \cdot \eta)$.*

Proof. Let $\delta = e^{-\epsilon - \frac{\epsilon+\beta}{\beta-1}}$. By the previous lemma, $(\epsilon, \delta) \in d_{\text{DP}}(c_1 \cdot \eta, c_2 \cdot \eta)$. We will now show that $e^{-k} \geq \delta$.

Indeed,

$$\begin{aligned} \delta \leq e^{-k} &\Leftrightarrow e^{-\epsilon - \frac{\epsilon+\beta}{\beta-1}} \leq e^{-k} \Leftrightarrow -\epsilon - \frac{\epsilon + \beta}{\beta - 1} \leq -k \Leftrightarrow (k - 1)\beta + \frac{k\beta}{\beta - 1} \geq k \Leftrightarrow \frac{1}{\beta - 1} \geq \frac{1}{\beta} - \frac{k - 1}{k} \Leftrightarrow \\ &\frac{1}{\beta - 1} \geq \frac{1}{\beta} - \frac{1}{2} \Leftrightarrow \frac{1}{\beta} \leq \frac{1}{2} + \frac{1}{\beta - 1} \Leftrightarrow \beta \geq \frac{2(e^\beta - 1)}{e^\beta + 1} \Leftrightarrow \beta + \frac{4}{e^\beta + 1} \geq 2, \end{aligned}$$

where the “ \Leftrightarrow ” claim holds because $k \geq 2$. Consider now the function $f(x) = x + 4/(e^x + 1)$. We have $f(0) = 2$. Also, f is a monotone function. Hence the claim $\beta + 4/(e^\beta + 1) \geq 2$ holds. \square

We can now proceed to state the differential privacy theorem for the added Laplace noise. We start with a technical notion that we use in the proof, and which can also be used to state the additional conditions when we get a tighter bound for the (ϵ, δ) -differential privacy of the noised function.

Definition 3.41. A path in X is a continuous function $h : [0, 1] \rightarrow X$. The path h is *shortest*, if for all $x_1, x_2, x_3 \in [0, 1]$, $x_1 \leq x_2 \leq x_3$, the equality $d_X(h(x_1), h(x_3)) = d_X(h(x_1), h(x_2)) + d_X(h(x_2), h(x_3))$ holds.

Theorem 3.56. *Let $b, \beta, \epsilon \in \mathbb{R}_+$, $b > 0$, $b + \beta \leq \epsilon$. Define $k = 1 + (\epsilon - b)/\beta$. Let $\delta = e^{-k}$. Let η be a random variable distributed according to $\text{Lap}(1)$. Let c be a β -smooth upper bound on DS_f for a function $f : X \rightarrow \mathbb{R}$. Define $g(\vec{x}) := f(\vec{x}) + \frac{c(\vec{x})}{b} \cdot \eta$. Then*

- for any $\vec{x}_1, \vec{x}_2 \in X$, $(\epsilon \cdot L, 2\delta) \in d_{\text{DP}}(g(\vec{x}_1), g(\vec{x}_2))$, where $L = d_X(\vec{x}_1, \vec{x}_2)$;
- (in particular,) g is $(\epsilon, 2e^\epsilon \delta)$ -differentially private.

If, additionally for any two points $\vec{x}_1, \vec{x}_2 \in X$ there exists a shortest path h in X , such that c is monotonic along that path, then the factor “2” in previous statements can be removed..

Proof. Let b, β, ϵ be as in the statement of the theorem. Let $\eta \sim \text{Lap}(1)$ and $\vec{x}_1, \vec{x}_2 \in X$. Let $L = d_X(\vec{x}_1, \vec{x}_2)$. Let h be a shortest path from \vec{x}_1 to \vec{x}_2 . Let \vec{x}_μ be a point on the path h , such that $c(\vec{x}_\mu) = \max_{t \in [0, 1]} c(h(t))$. Let $L_1 = d_X(\vec{x}_1, \vec{x}_\mu)$ and $L_2 = d_X(\vec{x}_\mu, \vec{x}_2)$. Note that $L = L_1 + L_2$. Define the following probability distributions:

$$\begin{aligned} \chi_1 &= f(\vec{x}_1) + \frac{c(\vec{x}_1)}{b} \cdot \eta & \chi_2 &= f(\vec{x}_1) + \frac{c(\vec{x}_\mu)}{b} \cdot \eta \\ \chi_3 &= f(\vec{x}_2) + \frac{c(\vec{x}_\mu)}{b} \cdot \eta & \chi_4 &= f(\vec{x}_2) + \frac{c(\vec{x}_2)}{b} \cdot \eta . \end{aligned}$$

We want to show that $(\epsilon L, 2\delta) \in d_{\text{DP}}(\chi_1, \chi_4)$.

By Lemma 3.54, $(b \cdot |f(\vec{x}_2) - f(\vec{x}_1)|/c(\vec{x}_\mu), 0) \in d_{\text{DP}}(\chi_2, \chi_3)$. The difference between $f(\vec{x}_2)$ and $f(\vec{x}_1)$ can be upper-bounded as follows:

$$|f(\vec{x}_2) - f(\vec{x}_1)| = \left| \int_h f'(h) ds \right| \leq \left| \int_0^1 f'(h(t)) \|h'(t)\| dt \right| \leq \left| \int_0^1 c(h(t)) \|h'(t)\| dt \right| \leq \int_0^1 c(\vec{x}_\mu) \|h'(t)\| dt = c(\vec{x}_\mu) \cdot d_X(h(1) - h(0)) = L \cdot c(\vec{x}_\mu),$$

hence $(bL, 0) \in d_{\text{DP}}(\chi_2, \chi_3)$. Note that, in $\int_h f'(h) ds$, the integral is over the path h , and the derivative f' is w.r.t. distance d_X along the path h . This derivative exists for a differentiable f , which follows from the definition of DS_f . More generally, if X is a *Banach space*, its existence follows from the existence of the Fréchet derivative of f (we will discuss Banach spaces and Fréchet derivatives in more details in Sec. 3.3.8.2). The ds is an infinitesimal distance (according to d_X) along the path h , and the h in $f'(h)$ denotes the point on the path h .

We will now compare χ_1 and χ_2 . We use the previous lemma with the following instantiations:

Quantity in Lemma 3.55	Instantiation
c_1	$c(\vec{x}_1)$
c_2	$c(\vec{x}_\mu)$
β	$\ln(c(\vec{x}_\mu)/c(\vec{x}_1))$
ϵ	$(\epsilon - b)L_1$
k	$1 + (\epsilon - b)L_1 / \ln(c(\vec{x}_\mu)/c(\vec{x}_1))$

Lemma 3.55 required that $\epsilon \geq \beta$. This condition is translated to

$$(\epsilon - b)L_1 \geq \ln \frac{c(\vec{x}_\mu)}{c(\vec{x}_1)} .$$

Let us verify that it holds:

$$\ln \frac{c(\vec{x}_\mu)}{c(\vec{x}_1)} \leq \beta \cdot d_X(\vec{x}_1, \vec{x}_\mu) \leq \beta L_1 \leq (\epsilon - b)L_1 \quad \blacksquare .$$

We also lower-bound the value of k from Lemma 3.55, in order to simplify its expression:

$$1 + \frac{(\epsilon - b)L_1}{\ln \frac{c(\vec{x}_\mu)}{c(\vec{x}_1)}} \geq 1 + \frac{(\epsilon - b)L_1}{\beta L_1} = k .$$

We obtain

$$((\epsilon - b)L_1, e^{-k}) \in d_{\text{DP}}(\chi_1, \chi_2) .$$

Similarly, we can obtain

$$((\epsilon - b)L_2, e^{-k}) \in d_{\text{DP}}(\chi_3, \chi_4) .$$

Using the triangle inequality, we can combine $d_{\text{DP}}(\chi_i, \chi_{i+1})$ for $i \in \{1, 2, 3\}$:

$$(\epsilon L, 2e^{-k}) \in d_{\text{DP}}(\chi_1, \chi_4)$$

as required. If $d_X(\vec{x}_1, \vec{x}_2) = 1$, then $(\epsilon, 2e^{-k}) \in d_{\text{DP}}(g(\vec{x}_1), g(\vec{x}_2))$, i.e. g is $(\epsilon, 2\delta)$ -differentially private.

If c is monotone along the path h , then the point \vec{x}_μ coincides with either \vec{x}_1 or \vec{x}_2 . W.l.o.g. assume $\vec{x}_\mu = \vec{x}_1$. Then $\chi_1 = \chi_2$ and $(0, 0) \in d_{\text{DP}}(\chi_1, \chi_2)$. The triangle inequality now gives $(\epsilon L, e^{-k}) \in d_{\text{DP}}(\chi_1, \chi_4)$. \square

Even though we obtained a result about the (ϵ, δ) -differential privacy level of g , the guarantee given by Theorem 3.56 is of a rather different kind. Indeed, it gives a reasonable ‘‘group privacy’’ also when the arguments to f are far apart.

$$E ::= T \mid \pi_S(E) \mid \sigma_\theta(E) \mid E \times E \mid E \cup E \mid E \cap E \mid E - E \mid \rho_{S \rightarrow S'}(E) \mid \text{Dis}(E)$$

$$\theta ::= 0 \mid 1 \mid P(\vec{s}) \mid \theta \wedge \theta \mid \neg\theta$$

Here $E \in \mathbf{RExp}$, $\theta \in \mathbf{BExp}$, $T \in \mathbf{R}$, $S, S' \in \mathbf{N}^{*\neq}$ and $|S| = |S'|$, $P \in \mathbf{P}$, $\vec{s} \in \mathbf{N}^*$.

Figure 37: Syntax of Relational Algebra

3.3.7.2 Continuous Semantics for Relational Algebra Queries. SQL is used to specify queries against relational databases, i.e. databases made up of several tables, where each table is a multiset of records of certain type. We are going to apply the theory described above to the differential privacy of SQL queries. For this to be possible, we need to give a semantics for SQL that is based on the multiplicities of records being real numbers. When applied to “normal” databases, i.e. to databases where each row in each table has an integer multiplicity, then our semantics will coincide with the standard semantics of SQL.

The semantics of SQL queries has recently been studied in [53], where its expressivity has been shown to be the same as a relational algebra fragment. The fragment consists of a following components:

Names. These are used to name the columns of tables/relations. Let \mathbf{N} be the set of all names. This set is assumed to be countably infinite.

Values. These inhabit the cells in relations. To simplify the treatment, all values are taken to have the same type. Let \mathbf{C} be the set of all values.

Basic relations. These are the names of the tables in the database. Let \mathbf{R} be the (finite) set of these names.

Predicate symbols. These are used to state boolean conditions on top of values. Let \mathbf{P} be the set of these symbols. We assume that a symbol for equality checking is an element of \mathbf{P} .

On top of these sets, the syntax of relational algebra is built. The syntax defines two categories: relational algebra expressions (the set \mathbf{RExp}) and boolean expressions (the set \mathbf{BExp}) given in Fig. 37. Given a set X , we let X^* denote all sequences of the elements of X . We also let $X^{*\neq} \subset X^*$ denote all such sequences where all elements are different. If $Y \in X^{*\neq}$, then we may also interpret Y as a subset of X , consisting of all elements that occur in it.

We see that a relational algebra expression can be a table in the database, a projection, a filter, a combination of results of two expressions (cartesian product, union, intersection, difference), renaming of columns. Finally, if E is an expression then $\text{Dis}(E)$ is also an expression, corresponding to the removal of repeated rows from the result of E . It corresponds to the `SELECT DISTINCT` keyword in SQL.

The semantics of relational algebra is defined as follows. For each predicate symbol $P \in \mathbf{P}$, let there be an associated semantic function $\llbracket P \rrbracket : \mathbf{C}^* \rightarrow \mathbb{B}$, where $\mathbb{B} = \{0, 1\}$. Let there be a mapping $\ell : \mathbf{R} \rightarrow \mathbf{N}^{*\neq}$, giving the schema of each basic relation. For a schema $S \in \mathbf{N}^{*\neq}$, the set of relations matching S is

$$\mathbf{Rels}(S) = \{\mathbf{r} : (S \rightarrow \mathbf{C}) \rightarrow \mathbb{R}_+ \mid |\text{supp } \mathbf{r}| < \infty\},$$

where $\text{supp } \mathbf{r} \subseteq (S \rightarrow \mathbf{C})$ denotes the *support* of \mathbf{r} — the set of all such tuples $X : S \rightarrow \mathbf{C}$ where $\mathbf{r}(X) \neq 0$. We see that \mathbf{r} gives the multiplicity of each possible row in a table with schema S , and the multiplicities can be fractional numbers. A database D maps each table name $T \in \mathbf{R}$ to a relation $D^T \in \mathbf{Rels}(\ell(T))$.

We extend the schema ℓ to all relational algebra expressions in the manner given in Figure 38. In this extension, the *type* $\tau(\theta) \subset \mathbf{N}$ of a boolean expression θ is the set of all names that occur in it. We consider a relational algebra expression E syntactically valid only if $\ell(E)$ is defined, and define the semantics only for syntactically valid expressions.

To define the semantics of removing repeated rows, we have to explain what happens with their multiplicity. We do not have to fix this in a unique manner, but can choose the behavior of the `DISTINCT`-keyword in a way that is most beneficial to our analysis, given that it behaves in the traditional way on

$\ell(\pi_S(E)) = S$	undefined, if $S \not\subseteq \ell(E)$
$\ell(\sigma_\theta(E)) = \ell(E)$	undefined, if $\tau(\theta) \not\subseteq \ell(E)$
$\ell(E_1 \times E_2) = \ell(E_1) \parallel \ell(E_2)$	undefined, if $\ell(E_1) \cap \ell(E_2) \neq \emptyset$
$\ell(E_1 \cup E_2) = \ell(E_1)$	undefined, if $\ell(E_1) \neq \ell(E_2)$
$\ell(E_1 \cap E_2) = \ell(E_1)$	undefined, if $\ell(E_1) \neq \ell(E_2)$
$\ell(E_1 - E_2) = \ell(E_1)$	undefined, if $\ell(E_1) \neq \ell(E_2)$
$\ell(\rho_{S \rightarrow S'}(E)) = S'$	undefined, if $\ell(E) \neq S$
$\ell(\text{Dis}(E)) = \ell(E)$	

Figure 38: Schemas of Relational Expressions

$$\begin{aligned}
\llbracket 0 \rrbracket(r) &= 0 \\
\llbracket 1 \rrbracket(r) &= 1 \\
\llbracket P(s_1, \dots, s_k) \rrbracket(r) &= \llbracket P \rrbracket(r(s_1), \dots, r(s_k)) \\
\llbracket \theta_1 \wedge \theta_2 \rrbracket(r) &= \min\{\llbracket \theta_1 \rrbracket(r), \llbracket \theta_2 \rrbracket(r)\} \\
\llbracket \neg\theta \rrbracket(r) &= 1 - \llbracket \theta \rrbracket(r)
\end{aligned}$$

Figure 39: Semantics of Boolean Expressions

$$\begin{aligned}
\llbracket T \rrbracket_D(r) &= D^T(r) \\
\llbracket \pi_S(E) \rrbracket_D(r) &= \sum_{\substack{r': \ell(E) \rightarrow \mathbf{C} \\ \forall s \in S: r(s) = r'(s)}} \llbracket E \rrbracket_D(r') \\
\llbracket \sigma_\theta(E) \rrbracket_D(r) &= \llbracket E \rrbracket_D(r) \cdot \llbracket \theta \rrbracket(r) \\
\llbracket E_1 \times E_2 \rrbracket_D(r) &= \llbracket E_1 \rrbracket_D(r|_{\ell(E_1)}) \cdot \llbracket E_2 \rrbracket_D(r|_{\ell(E_2)}) \\
\llbracket E_1 \cup E_2 \rrbracket_D(r) &= \llbracket E_1 \rrbracket_D(r) + \llbracket E_2 \rrbracket_D(r) \\
\llbracket E_1 \cap E_2 \rrbracket_D(r) &= \min\{\llbracket E_1 \rrbracket_D(r), \llbracket E_2 \rrbracket_D(r)\} \\
\llbracket E_1 - E_2 \rrbracket_D(r) &= \max\{0, \llbracket E_1 \rrbracket_D(r) - \llbracket E_2 \rrbracket_D(r)\} \\
\llbracket \rho_{s_1 \dots s_k \rightarrow s'_1 \dots s'_k}(E) \rrbracket_D(r) &= \llbracket E \rrbracket_D(\{ \forall i: s_i \mapsto r(s'_i) \}) \\
\llbracket \text{Dis}(E) \rrbracket_D(r) &= \text{Sgm}(\llbracket E \rrbracket_D(r))
\end{aligned}$$

Figure 40: Semantics of Expressions of Relational Algebra

integral multiplicities. Hence we state that there is a function $\text{Sgm} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which is continuous, differentiable, and *monotonic*. Moreover, it satisfies $\text{Sgm}(0) = 0$ and $\text{Sgm}(x) = 1$ for all $x \geq 1$.

The semantics $\llbracket \theta \rrbracket$ of a boolean expression θ is a mapping from a row (with the type $S \rightarrow \mathbf{C}$, where $S \supseteq \tau(\theta)$) to a boolean value. It is given in Figure 39.

The semantics $\llbracket E \rrbracket_D$ (relative to a database D) of a relational algebra expression E is a relation in $\mathbf{Rels}(\ell(E))$, i.e. a mapping from rows to their multiplicities, where a row r has the type $\ell(E) \rightarrow \mathbf{C}$. It is given in Figure 40.

The purpose of adding the noise to the result of a query E is to hide small changes in the database that is the subject of the query. Hence the semantics of E has to be seen as a function that takes a

database as an input. We have

$$\llbracket E \rrbracket : \prod_{T \in \mathbf{R}} \mathbf{Rels}(\ell(T)) \rightarrow \mathbf{Rels}(\ell(E)),$$

where the domain of this function is the type of databases. The notation $\llbracket E \rrbracket_D$ means the application of the function $\llbracket E \rrbracket$ to the database D . We consider $\llbracket E \rrbracket$ as a multivariate function (also returning many values), with each variable taking values in \mathbb{R}_+ . A variable corresponds to a possible row in one of the tables, i.e. it is identified by

1. the name of the table $T \in \mathbf{R}$;
2. the values in the row: $\mathbf{r} : \ell(T) \rightarrow \mathbf{C}$.

Let $x_{T;\mathbf{r}}$ be the variable that corresponds to the potential row \mathbf{r} in the table T . When describing a database, the value of $x_{T;\mathbf{r}}$ is the multiplicity of the row \mathbf{r} in the table T .

SQL rewriting. There may be a number of different ways to extend an integral semantics of a SQL query to a fractional one. All these fractional extensions are equally valid as the basis for computing the magnitude of noise to be added to the query result in order to obtain a certain level of differential privacy. They may lead to the selection of different amounts of noise, or its addition in different places, resulting in different accuracy, but the privacy guarantee remains, because all different fractional semantics agree on integral points.

A database engine typically rewrites SQL queries in order to make their execution more efficient. The engine aims to preserve the integral semantics of the query, but does not care about the fractional semantics. Hence the fractional semantics may change during rewriting. Nevertheless, the discussion in the previous paragraph shows that the privacy analysis performed on the initial query remains valid also for the rewritten query. Our extension of the semantics hence does not bring about a need to change the query execution strategies of database engines.

3.3.7.3 Smooth Derivative Sensitivity for Relational Algebra Queries. A relational algebra expression returns relations, but our noise addition mechanisms naturally work with functions that return a numeric value. We note that this restriction is similar to other differential privacy mechanisms for database queries, e.g. PINQ [54]. These numbers could be the values of particular cells in the returned relation, or some summary of the relation. Inspired by the SQL-statements turning up in scenarios investigated by Brandeis CRTs, we focus on counting the rows of a relation. Let

$$\text{Count}_E(D) = \llbracket \pi_\emptyset(E) \rrbracket_D(\emptyset)$$

be the number of rows in the result of the query E (as a function of the database instance D). Let the derivative sensitivity of this be

$$\text{DSC}_E(D) = \text{DS}_{\text{Count}_E}(D)$$

Our methods for finding smooth derivative sensitivity of Count_E -queries cannot currently handle all of the relational algebra specified in Figure 37, but they can cover expressions E corresponding to the SQL-queries of the form

$$\text{SELECT } A_1, \dots, A_k \text{ FROM } T_1, \dots, T_n \text{ WHERE } \textit{condition} \tag{31}$$

where *condition* may be any boolean expression in Figure 37 and where the tables may appear in the list T_1, \dots, T_n several times.

Finding the derivative. As defined previously, Let $x_{T;\mathbf{r}}$ is the multiplicity of the row \mathbf{r} in the table T of the database D . Considering different cases of E , we find the derivative of the count function with respect to $x_{T;\mathbf{r}}$. The simplest case is, when E is just a table. Then

$$\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{T'}(D)) = \begin{cases} 1, & \text{if } T = T' \\ 0, & \text{if } T \neq T' \end{cases}.$$

The derivative of the count of the projection, renaming, and filtering operations are trivial to inductively specify:

$$\begin{aligned}\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\pi_S(E)}(D)) &= \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_E(D)) \\ \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\rho_{s_1 \dots s_k \rightarrow s'_1 \dots s'_k}}(E)(D)) &= \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_E(D)) .\end{aligned}$$

To compute the derivative of the count of a filtering operation, we extend the definition of $\llbracket \theta \rrbracket$ so that it can also be applied to rows that do not contain all the columns referenced in θ . In this case, the unbound variables are assumed to be existentially quantified, i.e.

$$\llbracket \theta \rrbracket(\mathbf{r}) = \sup_{\mathbf{r}' \text{ extends } \mathbf{r}} \llbracket \theta \rrbracket(\mathbf{r}'),$$

where $\mathbf{r}' : \tau(\theta) \rightarrow \mathbf{C}$ ranges over all possible rows that define the attributes used by θ . The derivative of the count of the filtering is the derivative of the count of the original expression, moderated by the truth value of θ on the row \mathbf{r} .

$$\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(T)}(D)) = \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_T(D)) \cdot \llbracket \theta \rrbracket(\mathbf{r}) .$$

The derivative of the count of a Cartesian product is defined as the derivative of a product function:

$$\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{E_1 \times E_2}(D)) = \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{E_1}(D)) \cdot \text{Count}_{E_2}(D) + \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{E_2}(D)) \cdot \text{Count}_{E_1}(D)$$

We are interested in the case where filtering is applied to the product (and projection is applied afterwards, but its derivative was trivial). In this case, the derivative can be upper-bounded as follows:

$$\begin{aligned}\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(E_1 \times E_2)}(D)) &\leq \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(E_1) \times \sigma_\theta(E_2)}(D)) = \\ &\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(E_1)}(D)) \cdot \text{Count}_{\sigma_\theta(E_2)}(D) + \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(E_2)}(D)) \cdot \text{Count}_{\sigma_\theta(E_1)}(D) \leq \\ &\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{E_1}(D)) \cdot \text{Count}_{\sigma_\theta(E_2)}(D) + \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{E_2}(D)) \cdot \text{Count}_{\sigma_\theta(E_1)}(D)\end{aligned}$$

If the Cartesian product is applied to any number of expressions E_j , we get

$$\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\prod_j E_j}(D)) = \sum_i \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{E_i}(D)) \cdot \text{Count}_{\prod_{j \neq i} E_j}(D),$$

which can be upper-bounded as

$$\begin{aligned}\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(\prod_j E_j)}(D)) &\leq \\ \sum_i \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(E_i)}(D)) \cdot \text{Count}_{\sigma_\theta(\prod_{j \neq i} E_j)}(D) &\leq \sum_i \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{E_i}(D)) \cdot \text{Count}_{\sigma_\theta(\prod_{j \neq i} E_j)}(D) \quad (32)\end{aligned}$$

If each subquery E_i is a table ($E_i = T_i$) and all T_i are distinct then only one of the summands can be nonzero—the one where $T = T_i$. If the T_i are not distinct then there would be one summand for each copy of the table T :

$$\begin{aligned}\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_\theta(\prod_j T_j)}(D)) &\leq \\ \sum_{i: T_i=T} \frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{T_i}(D)) \cdot \text{Count}_{\sigma_{\theta[\mathbf{r}/T_i]}(\prod_{j \neq i} T_j)}(D) &= \sum_{i: T_i=T} \text{Count}_{\sigma_{\theta[\mathbf{r}/T_i]}(\prod_{j \neq i} T_j)}(D) .\end{aligned}$$

Here we can be more precise than in the computation (32). As we know that \mathbf{r} corresponds to a row in the relational algebra expression T_i , we can specialize the filter θ , making the references to the columns of T_i to be equal to the values of the attributes in \mathbf{r} . We denote the specialization with $\theta[\mathbf{r}/T_i]$. It does not quite fit the syntax in Figure 37, but its semantics can be straightforwardly defined.

Finding the derivative sensitivity. Now the derivative sensitivity is found by taking the maximum of all the derivatives (which are non-negative for the current case, so it is not necessary to take the absolute value).

$$\begin{aligned} \text{DSC}_{\sigma_{\theta}(\prod_j T_j)}(D) &= \max_{T, \mathbf{r}} \frac{\partial}{\partial x_{T, \mathbf{r}}} (\text{Count}_{\sigma_{\theta}(\prod_j T_j)}(D)) \\ \text{DSC}_{\sigma_{\theta}(\prod_j T_j)}(D) &\leq \max_{T, \mathbf{r}} \sum_{i: T_i=T} \text{Count}_{\sigma_{\theta(\mathbf{r}/T_i)}(\prod_{j \neq i} T_j)}(D) \end{aligned}$$

Patterns. Above we showed how to compute an upper bound of $\text{Count}_{\sigma_{\theta(\mathbf{r}/T_i)}(\prod_{j \neq i} T_j)}(D)$ for each possible row \mathbf{r} separately. However, the number of possible rows may be very large or even infinite, making this approach impractical. To make it more practical, we split the space of possible input rows into “equivalence classes” such that for each row \mathbf{r} in an equivalence class X , the derivative $\text{Count}_{\sigma_{\theta(\mathbf{r}/T_i)}(\prod_{j \neq i} T_j)}(D)$ is the same. Sometimes, we may use an upper bound instead of the actual derivative and require only the upper bounds to be the same within an equivalence class. This reduces the number of equivalence classes. Each equivalence class can contain rows from only one table.

When computing an upper bound of

$$\text{DSC}_{\sigma_{\theta}(\prod_j T_j)}(D) \leq \max_{T, \mathbf{r}} \sum_{i: T_i=T} \text{Count}_{\sigma_{\theta(\mathbf{r}/T_i)}(\prod_{j \neq i} T_j)}(D)$$

the quantity $\text{Count}_{\sigma_{\theta(\mathbf{r}/T_i)}(\prod_{j \neq i} T_j)}(D)$ will then depend on both \mathbf{r} and i instead of only i (which would make the number of different variants finite). The number of possible \mathbf{r} can be very large. This is where the equivalence classes become useful. Suppose that $\theta = \theta_1 \wedge \dots \wedge \theta_n$ and some of the θ_k represent an equality constraint that equates two columns (these are often used for joining two tables). We take the transitive closure \mathcal{T} of this set of equality constraints. When going through the elements $\tilde{\mathbf{r}}$ of $\prod_{j \neq i} T_j$, we use \mathcal{T} to determine which cells of \mathbf{r} are equal to some of the cells of $\tilde{\mathbf{r}}$. These cells can be bound to concrete values in the expression θ . The values in other cells of \mathbf{r} should range over all elements of \mathbf{C} ; instead of enumerating all their possible values, we existentially quantify these cells.

Such \mathbf{r} is represented as a pattern $\mathbf{p} \in \ell(T) \rightarrow \mathbf{C} \cup \{*\}$, where some cells have concrete values but the rest (which are existentially quantified) are represented as $*$. For fixed i , the cells where the $*$ occur are the same for all obtained patterns. Thus the subsets of $\mathbf{C}^{\ell(T_i)}$ represented by different obtained patterns are disjoint and we can just count the number of times each pattern is obtained. If the tables T_i are all different then number of summands in

$$\max_{T, \mathbf{r}} \sum_{i: T_i=T} \text{Count}_{\sigma_{\theta}(\prod_{j \neq i} T_j)}(D)$$

is 1 and we can simply take the maximum of the counts of the patterns.

However, if the number of summands is greater than 1 then the $*$ do not necessarily occur in the same places in all the patterns. Then we need to unify the patterns. For example, if from one summand we obtain a pattern

$$(*, 100, *)$$

with count 5 and from another summand the pattern

$$(1000, *, *)$$

with count 10, then we unify the two patterns and get

$$(1000, 100, *)$$

with count $5 + 10 = 15$. In general, we must unify many such patterns. Then we must consider the unifications of all the subsets of patterns because the unification of a larger subset might not match anything. To find all those unifications, we use a divide-and-conquer approach, splitting the set of

patterns in half, unifying (and sorting) each half recursively, and then merging the results, similarly to merge sort.

After unification, we have a set of pairs (\mathbf{p}, d) , each denoting that for all rows \mathbf{r} matched by the pattern \mathbf{p} but not by a more specific pattern in the obtained pattern set, there is an upper bound for the derivative:

$$\frac{\partial}{\partial x_{T;\mathbf{r}}}(\text{Count}_{\sigma_{\theta}(\prod_j T_j)}(D)) \leq d .$$

Also, each row \mathbf{r} for which the derivative is nonzero, is matched by at least one obtained pattern, exactly one of these patterns is the most specific (the one that is the unification of all obtained patterns matching \mathbf{r}) and this has the maximum d among the matching patterns.

Now, to compute an upper bound of $\text{DSC}_{\sigma_{\theta}(\prod_j T_j)}(D)$, we simply have to take the maximum of the obtained d :

$$\text{DSC}_{\sigma_{\theta}(\prod_j T_j)}(D) \leq \max_{\mathbf{p}} d = \max_{\mathbf{p}} \sum_{i: R(\mathbf{p}) \subseteq C^{\ell(T_i)}} \text{Count}_{\sigma_{\theta[\mathbf{p}/T_i]}(\prod_{j \neq i} T_j)}(D)$$

Here \mathbf{p} ranges over the set of obtained patterns, $R(\mathbf{p})$ is the set of rows matched by the pattern \mathbf{p} . The summation here is the one performed during the unification of patterns, each summand denotes the count of one of the original patterns before unification. $\theta[\mathbf{p}/T_i]$ denotes that the columns of T_i (except those for which there is $*$ in the pattern) in θ there are bound to the values in \mathbf{p} (i.e. replaced by constants).

Let us now describe in more details how pattern unification works.

Unification of patterns. Let us consider a table T used in the query. It may be used more than once, i.e. $T = T_i$ may hold for more than one i . Let $\mathcal{I} = \{i \mid T = T_i\}$. The set of possible rows in T is $\mathcal{R} = \mathbf{C}^{\ell(T)}$. This is the set of rows whose addition or removal is considered when determining the sensitivity of a query (as opposed to the set of rows D^T in the table T in the actual database D). The set of possible patterns is $\mathcal{P} = (\mathbf{C} \cup \{*\})^{\ell(T)} \cup \{\perp\}$. We have added the null pattern \perp that does not match any row. The set of rows matched by a pattern $\mathbf{p} \in \mathcal{P}$, is $\rho(\mathbf{p}) = \{\mathbf{r} \in \mathcal{R} \mid \forall s \in \ell(T), \mathbf{p}(s) = \mathbf{r}(s) \vee \mathbf{p}(s) = *\}$ if $\mathbf{p} \neq \perp$, and $\rho(\perp) = \emptyset$. The set of rows matched by a pattern set $P \subseteq \mathcal{P}$ is $\bigcup_{\mathbf{p} \in P} \rho(\mathbf{p})$. The *intersection* of two patterns $\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{P}$ is a pattern $\mathbf{p}_1 \cap \mathbf{p}_2 \in \mathcal{P}$ such that $\rho(\mathbf{p}_1) \cap \rho(\mathbf{p}_2) = \rho(\mathbf{p}_1 \cap \mathbf{p}_2)$. It is easy to see that $\mathbf{p}_1 \cap \mathbf{p}_2$ exists and is unique. A set of patterns $P \subseteq \mathcal{P}$ is *disjoint* if $\forall \mathbf{p}_1, \mathbf{p}_2 \in P, \mathbf{p}_1 \neq \mathbf{p}_2 \Rightarrow \rho(\mathbf{p}_1) \cap \rho(\mathbf{p}_2) = \emptyset$.

Consider a fixed $i \in \mathcal{I}$. Before, we went through $\sigma_{\theta_i}(\prod_{j \neq i} T_j)$ and obtained a disjoint set of patterns P_i with a count (multiplicity) for each pattern. Let us call a set of patterns with a (possibly fractional) multiplicity (in \mathbb{R}_+) for each pattern, where the multiplicity $M(\mathbf{p})$ of a pattern \mathbf{p} is monotonically decreasing in the set of matched rows of the pattern ($\rho(\mathbf{p}_1) \subseteq \rho(\mathbf{p}_2) \Rightarrow M(\mathbf{p}_1) \geq M(\mathbf{p}_2)$), a *pattern map*. We would like to take the union $\bigcup_{i \in \mathcal{I}} P_i$ but this is not necessarily disjoint. Thus a row may be matched by more than one pattern, making it difficult to find the represented row(s) with the highest multiplicity.

We need a more general class of pattern sets which can represent $\bigcup_{i \in \mathcal{I}} \rho(P_i)$ but each row would be represented by only one pattern in the set. Let P be a pattern set and let a row \mathbf{r} be represented by the most specific pattern $\mathbf{p} \in P$ that matches \mathbf{r} , i.e. $\rho(\mathbf{p}) \subseteq \rho(\mathbf{p}')$ for all $\mathbf{p}' \in P$ that match \mathbf{r} . Then $\rho(\mathbf{p}) = \bigcap \{\rho(\mathbf{p}') \mid \mathbf{p}' \in P, \mathbf{r} \in \rho(\mathbf{p}')\}$ and $\mathbf{p} = \bigcap \{\mathbf{p}' \mid \mathbf{p}' \in P, \mathbf{r} \in \rho(\mathbf{p}')\}$. Thus \mathbf{p} is unique. To ensure that \mathbf{p} exists, we require that P be closed under intersection. We call such a P a *closed pattern set*. A disjoint pattern set can be converted to a closed pattern set by adding the null pattern.

Suppose now that we have two closed pattern sets $P_1, P_2 \in \mathcal{P}$. How do we find a closed pattern set $P \in \mathcal{P}$ such that $\rho(P) = \rho(P_1) \cup \rho(P_2)$? The naive method is to compute P as $P_1 \cup P_2 \cup \{\mathbf{p}_1 \cap \mathbf{p}_2 \mid \mathbf{p}_1 \in P_1, \mathbf{p}_2 \in P_2\}$. This has complexity $\Omega(|P_1| \cdot |P_2|)$. We will instead use a method similar to the merge of merge sort, having complexity $O(|P_1| + |P_2| + |P|)$ if the length of patterns $|\ell(T)|$ is considered constant (it may be exponential in $\ell(T)$, at most $O(3^{\ell(T)})$, due to the three recursive MergeCPS calls when computing B_1, B_2, B_3 ; but $\ell(T)$ is usually small in practice and the worst case is probably not reached very easily, unless $B_1 = B_2 = B_3$). $|P|$ is also $\Theta(|P_1| \cdot |P_2|)$ in the worst case but often is much less in practice. This method assumes that P_1 and P_2 are input lexicographically sorted (\perp is considered as the smallest element, and $*$ is larger than any element of \mathbf{C}). Then P is also output lexicographically sorted. Let \top be the largest element of \mathcal{P} , i.e. $\top = (*, \dots, *)$. Then $P \cup \{\top\} = \{\mathbf{p}_1 \cap \mathbf{p}_2 \mid \mathbf{p}_1 \in P_1 \cup \{\top\}, \mathbf{p}_2 \in P_2 \cup \{\top\}\}$.

- **function** MergeCPS(P_1, P_2):
 - **Input:** P_1 and P_2 are lexicographically sorted lists (without repetitions) of patterns, which are also closed pattern sets when viewed as sets with added \perp . All patterns in P_1 and P_2 must have the same length.
 - **Output:** A lexicographically sorted list (without repetitions) of patterns P , which is also a closed pattern set when viewed as a set with added \perp . Also $P = \{\mathbf{p}_1 \cap \mathbf{p}_2 \mid \mathbf{p}_1 \in P_1, \mathbf{p}_2 \in P_2\}$ holds when P, P_1, P_2 are viewed as sets.
 - If P_1 or P_2 (or both) is an empty list then **return** an empty list.
 - If the length of the patterns in the input is zero then **return** a list containing only the empty pattern (the pattern $()$ with length zero).
 - Let A be the sorted list of elements that occur as the first component of some pattern in P_1 or P_2 . Always include $*$ in A even if does not occur as the first component of any pattern.
 - For all $i \in \{1, 2\}, a \in A$, let P_{ia} be the lexicographically sorted (possibly empty) list of the patterns in P_i that begin with a , with the first component (a) removed.
 - Let $Q_* = \text{MergeCPS}(P_{1*}, P_{2*})$.
 - For all $a \in A \setminus \{*\}$, let $Q_a = \text{SimpleMerge}(B_1, B_2, B_3)$, where $B_1 = \text{MergeCPS}(P_{1a}, P_{2a})$, $B_2 = \text{MergeCPS}(P_{1*}, P_{2a})$, $B_3 = \text{MergeCPS}(P_{1a}, P_{2*})$.
 - For all $a \in A$, let Q'_a be the list obtained from Q_a by prepending a to each pattern of Q_a .
 - Let P be the concatenation of the lists Q'_a over $a \in A$ in ordered by a .
 - **Return** P .
- **function** SimpleMerge(B_1, B_2, B_3) merges three sorted lists (without repetitions) and removes the repetitions of elements.

Algorithm 3: Merging closed pattern sets

Algorithm 3 computes $\{\mathbf{p}_1 \cap \mathbf{p}_2 \mid \mathbf{p}_1 \in P_1, \mathbf{p}_2 \in P_2\}$, i.e. it merges closed pattern sets. We actually needed to merge pattern maps, i.e. for each pattern $\mathbf{p} = \mathbf{p}_1 \cap \mathbf{p}_2$ in the merged pattern set we also need to compute the multiplicity $M(\mathbf{p}) = M(\mathbf{p}_1) + M(\mathbf{p}_2)$. It is easy to see that the monotonicity of the multiplicity is maintained by the merge (because $\mathbf{p} = \mathbf{p}_1 \cap \mathbf{p}_2 \wedge \mathbf{p}' = \mathbf{p}'_1 \cap \mathbf{p}'_2 \wedge \rho(\mathbf{p}) \subseteq \rho(\mathbf{p}') \Rightarrow \rho(\mathbf{p}_1) \subseteq \rho(\mathbf{p}'_1) \wedge \rho(\mathbf{p}_2) \subseteq \rho(\mathbf{p}'_2)$). The monotonicity ensures that the most specific pattern which represents a row gives the maximum multiplicity. To compute the multiplicities, we need to slightly augment Algorithm 3. In SimpleMerge, when merging two equal patterns, we retain the copy with the maximum multiplicity (i.e. the one obtained from the most specific patterns). In MergeCPS, multiplicities are added when merging patterns with length zero. Elsewhere, multiplicity is always attached to each pattern and it is not changed when the first component of a pattern is removed or added back.

Now we can merge two closed pattern maps. If we need to merge more than two (i.e. $|I| > 2$), we split the set of pattern maps in half, merge each half recursively and then merge the results, similarly to merge sort. The initial pattern maps need to be lexicographically sorted. If they are not then we can consider each pattern (with multiplicity) in the pattern maps as a separate pattern map and merge the many small pattern maps recursively. This is the approach taken in the implementation.

Smoothing the derivative sensitivity. After computing the derivative sensitivity, we need to find a smooth upper bound of this. We smoothe $\text{Count}_{\sigma_{\theta[\mathbf{p}/T_i]}(\prod_{j \neq i} T_j)}(D)$ separately for each \mathbf{p}, i . This gives us a smooth upper bound for the whole expression because the smoothing distributes over maximum and sum. In the following, we write simply θ instead of $\theta[\mathbf{p}/T_i]$. Let

$$A^{(k)}(D) = \max_{d(D, D')=k} \text{Count}_{\sigma_{\theta}(\prod_{j \neq i} T_j)}(D') . \quad (33)$$

This is similar to the *sensitivity at distance k* from [52]. In our case, k need not be an integer. We can actually generalize $d(D, D')$ even more. Let us assume that the difference in rows for one table can

be more significant than the different of rows in another tables. That is, we can assign a weight G_i to the table T_i , so that

$$d((T_1, T_2, \dots, T_m), (T'_1, T'_2, \dots, T'_m)) = \sum_{i=1}^m G_i \sum_{x \in X_i} |T_i(x) - T'_i(x)| ,$$

$$G_i > 0 .$$

Putting $G_i = 1$ for all i , we get a standard definition of sensitivity at distance k . In this section, we will only need $G_i = 1$. The more general result will be needed in Section 3.3.9.

Because $\text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}$ is monotonically increasing, we may consider only those D' that are obtained by adding rows to D . Note that

$$\text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D) \leq \text{Count}_{\prod_{j \neq i} \sigma_\theta(T_j)}(D) = \prod_{j \neq i} \text{Count}_{\sigma_\theta(T_j)}(D) .$$

Suppose we add k_j rows to T_j for all j , so that $\sum_j G_j k_j = k$. Then

$$\text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D') \leq \prod_{j \neq i} \text{Count}_{\sigma_\theta(T_j)}(D') \leq \prod_{j \neq i} (\text{Count}_{\sigma_\theta(T_j)}(D) + k_j) .$$

We can get a better bound by noting that every extra row counted in $\text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D')$ compared to $\text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D)$ is also an extra row in $\text{Count}_{\prod_{j \neq i} \sigma_\theta(T_j)}(D')$ compared to $\text{Count}_{\prod_{j \neq i} \sigma_\theta(T_j)}(D)$. Thus

$$\begin{aligned} & \text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D') - \text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D) \leq \text{Count}_{\prod_{j \neq i} \sigma_\theta(T_j)}(D') - \text{Count}_{\prod_{j \neq i} \sigma_\theta(T_j)}(D) = \\ & = \prod_{j \neq i} \text{Count}_{\sigma_\theta(T_j)}(D') - \prod_{j \neq i} \text{Count}_{\sigma_\theta(T_j)}(D) \leq \prod_{j \neq i} (\text{Count}_{\sigma_\theta(T_j)}(D) + k_j) - \prod_{j \neq i} \text{Count}_{\sigma_\theta(T_j)}(D) \end{aligned}$$

$$\text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D') \leq \text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D) + \prod_{j \neq i} (\text{Count}_{\sigma_\theta(T_j)}(D) + k_j) - \prod_{j \neq i} \text{Count}_{\sigma_\theta(T_j)}(D) .$$

We already described how to compute $\text{Count}_{\sigma_\theta(\prod_{j \neq i} T_j)}(D)$. Now we show how to compute the maximum value of

$$N = \prod_{j \neq i} (\text{Count}_{\sigma_\theta(T_j)}(D) + k_j) - \prod_{j \neq i} \text{Count}_{\sigma_\theta(T_j)}(D) .$$

Let $n_j = \text{Count}_{\sigma_\theta(T_j)}(D)$. Then

$$N = \prod_{j \neq i} (n_j + k_j) - \prod_{j \neq i} n_j .$$

We want to find $\max_{\sum G_j k_j = k} N$. Note that

$$\frac{\partial N}{\partial k_m} = \prod_{j \neq \{i, m\}} (n_j + k_j) = \frac{1}{n_m + k_m} \prod_{j \neq i} (n_j + k_j)$$

$$\frac{\partial N}{\partial (G_m k_m)} = \frac{1}{G_m (n_m + k_m)} \prod_{j \neq i} (n_j + k_j) .$$

The derivative is largest for the m for which $G_m (n_m + k_m)$ is smallest. Let us call the product of G_m and the number of rows in table T_m the modified count of table T_m . Thus we should add rows to the table with the smallest modified count first. After the modified count of this table reaches the modified count of the table with the second smallest modified count, we continue by adding rows to those two tables simultaneously (keeping their modified counts equal), until they reach the modified count of the third smallest table, and so on, until the whole budget of k units of d -distance (where one unit can be used to

change the modified count of a table by 1) has been distributed. This algorithm allows finding an upper bound

$$A(k) = \text{Count}_{\sigma_{\theta}(\prod_{j \neq i} T_j)}(D) + \prod_{j \neq i} (n_j + k_j) - \prod_{j \neq i} n_j$$

of $A^{(k)}(D)$, but we are interested in finding the smooth upper bound

$$\max_k e^{-k\beta} A(k) .$$

Consider first the logarithm of the argument of max:

$$L = -k\beta + \ln A(k)$$

and the derivative of the logarithm

$$\frac{\partial L}{\partial k} = -\beta + \frac{1}{A(k)} \cdot \frac{\partial(A(k))}{\partial k} .$$

The derivative of $A(k)$ is

$$\frac{\partial(A(k))}{\partial k} = \frac{\partial N}{\partial(G_m k_m)} = \frac{1}{G_m(n_m + k_m)} \prod_{j \neq i} (n_j + k_j)$$

where $m = \text{argmin}_j(G_j(n_j + k_j))$. Also

$$A(k) = \prod_{j \neq i} (n_j + k_j) - C$$

where $C = \prod_{j \neq i} n_j - \text{Count}_{\sigma_{\theta}(\prod_{j \neq i} T_j)}(D) \geq 0$. Now

$$\frac{\partial L}{\partial k} = -\beta + \frac{1}{G_m(n_m + k_m)} \cdot \frac{\prod_{j \neq i} (n_j + k_j)}{\prod_{j \neq i} (n_j + k_j) - C} = -\beta + \frac{1}{G_m(n_m + k_m)} \cdot \left(1 + \frac{C}{\prod_{j \neq i} (n_j + k_j) - C}\right)$$

The quantity $\frac{1}{G_m(n_m + k_m)}$ is decreasing in k when $k > 0$. Also, because $C \geq 0$ and when $k > 0$, $C < \prod_{j \neq i} (n_j + k_j)$, the quantity $1 + \frac{C}{\prod_{j \neq i} (n_j + k_j) - C}$ is positive and monotonically decreasing in k when $k > 0$. Thus also $\frac{\partial L}{\partial k}$ is decreasing in k when $k \in (0, \infty)$. Note also that $\lim_{k \rightarrow \infty} \frac{\partial L}{\partial k}(k) = -\beta$. Thus L is either decreasing in $(0, \infty)$, or increasing in $(0, s]$ and decreasing in $[s, \infty)$ for some s . Because L is continuous, we get that L is either decreasing in $[0, \infty)$, or increasing in $[0, s]$ and decreasing in $[s, \infty)$ for some s . The maximum that we want to find is obtained with $k = 0$ in the first case, and with $k = s$ in the second case. We know how to compute $\frac{\partial L}{\partial k}(k)$. First try to compute $\frac{\partial L}{\partial k}(0)$. If it is not defined (a division by zero occurs) then $\lim_{k \rightarrow 0^+} \frac{\partial L}{\partial k}(k) = +\infty$ and we have the second case. Otherwise if $\frac{\partial L}{\partial k}(0) \leq 0$ then we have the first case, and if $\frac{\partial L}{\partial k}(0) > 0$ then we have the second case. In the first case, we have the result. In the second case, we use binary search to find s , which is the value of k where $\frac{\partial L}{\partial k}(k)$ changes from positive to negative.

Let the θ now again denote a general boolean expression, not $\theta[p/T_i]$. Now we have found the $k = s$ that maximizes L and thus also $e^L = e^{-k\beta} A(k)$. This maximum of e^L is a β -smooth upper bound of $\text{Count}_{\sigma_{\theta[p/T_i]}(\prod_{j \neq i} T_j)}(D)$. We compute these upper bounds for each p, i and take the sum and the maximum to get a β -smooth upper bound of $\text{DSC}_{\sigma_{\theta}(\prod_j T_j)}(D)$.

Now we know how to compute a β -smooth upper bound of the derivative sensitivity of queries of the form $\sigma_{\theta}(\prod_j T_j)$. We can slightly generalize it to include a projection and renaming of columns at the top level, as these do not change the count or its derivatives. Thus we can compute a β -smooth upper bound of the derivative sensitivity of a query of the form $\rho_{s_1 \dots s_k \rightarrow s'_1 \dots s'_k}(\pi_S(\sigma_{\theta}(\prod_j T_j)))$ which is the same as that of the query $\sigma_{\theta}(\prod_j T_j)$.

3.3.8 Derivative Sensitivity for Components. Sensitivity w.r.t. component has been defined in Def. 3.19-3.22 of Sec. 3.3.5. These definitions are given in terms of *global* sensitivity, i.e. sensitivity that does not depend on the actual input. In this section, we extend the notion of sensitivity w.r.t. component to local and derivative. We mostly deal with A-sensitivity, and a bit with B-sensitivity (Def. 3.19-3.20). It is more difficult to give intuitive meaning to differential privacy using C- and D-sensitivity (Def. 3.21-3.22), and we have shown in Theorem 3.12 that they equal to A- and B-sensitivities using Hausdorff distances that we will consider in this section.

3.3.8.1 Differential Privacy and Local Sensitivity w.r.t. Component. We want a definition that guarantees indistinguishability for data instances x and x' that differ in component ρ by a certain amount, but are otherwise as similar as possible. The definitions of A-B-C-D-sensitivities w.r.t. components are aimed to match this definition and provide means of achieving it.

In Sec. 3.3.5, for an element $x \in X$, we called an element $x' \in X'$ that is the closest to x a *friend* of x in the set X' (Definition 3.16), denoted $x \overset{X'}{\sim} x'$. Let us use a shorthand notation $x \overset{\rho}{\sim} x'$ for $x \overset{[x']_{\rho}}{\sim} x'$. Using $x \overset{\rho}{\sim} x'$ to express the property “ x and x' as similar as possible while ρ is fixed”, we get the following definitions of differential privacy w.r.t. component.

Definition 3.42 (differential A-privacy w.r.t. component). Let X be a metric space, and $f : X \rightarrow D(Y)$. The mapping f is ε -differentially A-private w.r.t. component ρ if for all $Y' \subseteq Y$ and for all $x, x' \in X$ such that $x \overset{\rho}{\sim} x'$ and $d_X(x, x') = 1$, the following inequality holds:

$$Pr[f(x) \in Y'] \leq e^{\varepsilon} Pr[f(x') \in Y'] .$$

Definition 3.43 (differential B-privacy w.r.t. component). Let X be a metric space, and $f : X \rightarrow D(Y)$. The mapping f is ε -differentially B-private w.r.t. component ρ if for all $Y' \subseteq Y$ and for all $x, x' \in X$ such that $x \overset{\rho}{\sim} x'$ and $d_X([x]_{\rho}, [x']_{\rho}) = 1$, the following inequality holds:

$$Pr[f(x) \in Y'] \leq e^{\varepsilon} Pr[f(x') \in Y'] .$$

The definition of (ε, δ) -differential privacy would be analogous. Both Def. 3.42 and Def. 3.43 are suitable for cases where the data owner is interested in concealing changes in particular rows and columns of his datatables. For example, it could be the case where the table is shared by several owners, and one owner knows only some particular rows or columns in the database, and wants to protect his own data independently of the values belonging to the other data owners. While Def. 3.42 allows us to compute more precise bounds for the added noise, Def. 3.43 is better for capturing the change in a particular class ρ , and corresponds to our intuition of hiding a particular component. If we are dealing with components that are vector coordinates, and the distance corresponds to ℓ_p -norm, we actually have $d_X([x]_{\rho}, [x']_{\rho}) = d_X(x, x')$.

Local sensitivity w.r.t. component. The difference from the global sensitivity is that the local sensitivity may depend on the input. Similarly to *global* sensitivity w.r.t. component (Definitions 3.19-3.22), we list four slightly different definitions of *local* sensitivity w.r.t. component, now allowing sensitivity to depend on the function input.

Definition 3.44 (local A-B-C-D- sensitivity w.r.t. component). Let $c : \mathcal{P}(X) \rightarrow \mathbb{R}^+$, where $\mathcal{P}(X)$ denotes the powerset of X . Mapping $f : X \rightarrow Y$ is called locally c -sensitive w.r.t. the component ρ if for all $X_1, X_2 \in X/\rho$ we have

- A- $\forall x_1 \in X_1, x_2 \in X_2 : x_1 \overset{X_2}{\sim} x_2 \implies d_Y(f(x_1), f(x_2)) \leq c(X_1) \cdot d_X(x_1, x_2);$
- B- $\forall x_1 \in X_1, x_2 \in X_2 : x_1 \overset{X_2}{\sim} x_2 \implies d_Y(f(x_1), f(x_2)) \leq c(X_1) \cdot \tilde{d}_X(X_1, X_2);$
- C- $\forall X'_1 \subseteq X_1, X'_2 \subseteq X_2 : X'_1 \overset{X_2}{\sim} X'_2 \implies \tilde{d}_Y(f(X'_1), f(X'_2)) \leq c(X_1) \cdot \tilde{d}_X(X'_1, X'_2);$
- D- $\forall X'_1 \subseteq X_1, X'_2 \subseteq X_2 : X'_1 \overset{X_2}{\sim} X'_2 \implies \tilde{d}_Y(f(X'_1), f(X'_2)) \leq c(X_1) \cdot \tilde{d}_X(X_1, X_2).$

Alternatively, we could define $c : X \rightarrow \mathbb{R}^+$ on particular points $x_1 \in X_1$. Such definition would give more precise bounds since $c(x)$ needs to give a suitable bound only for x , and not necessarily for all other elements of $[x]_\rho$. On the other hand, $c([x]_\rho)$ is more useful if some data is missing and only the ρ -class of the input is known, e.g. in the case of shared database.

The proofs of local sensitivities w.r.t. composed components are analogous to those of global sensitivity. As an example, we show the composition proof local A-sensitivity, and the other types of local sensitivities are analogous. We see that another reason for taking $c([x]_\rho)$ is composability.

Theorem 3.57. *Let (ρ, σ) be an adjacent component pair (Def. 3.23). Let the mapping $f : X \rightarrow Y$ be c_ρ and c_σ locally A-sensitive w.r.t. ρ and σ respectively. Then, f is $(c_\rho + c_\sigma)$ locally A-sensitive w.r.t. $\rho \sqcap \sigma$.*

Proof. Let $x, x' \in X$ be such that $x \stackrel{[x']_{\rho \sqcap \sigma}}{\sim} x'$. We need to show that $d_Y(f(x), f(x')) \leq (c_\rho + c_\sigma)([x]_{\rho \sqcap \sigma}) \cdot d_X(x, x')$. Since (ρ, σ) is a adjacent pair, there exists $x'' \in X$ such that $x \stackrel{[x']_\rho}{\sim} x''$ and $x' \stackrel{[x']_\sigma}{\sim} x''$. Since $x'' \in [x']_\rho$ and $x'' \in [x']_\sigma$, we have $[x'']_\rho = [x']_\rho$, $[x'']_\sigma = [x']_\sigma$, and $x \stackrel{[x'']_\rho}{\sim} x''$, $x' \stackrel{[x'']_\sigma}{\sim} x''$, so we can apply A-sensitivity.

- $d_Y(f(x), f(x'')) \leq c_\rho([x']_\rho) \cdot d_X(x, x'')$ for $\hat{x} \in \{x, x''\}$;
- $d_Y(f(x''), f(x')) \leq c_\sigma([x']_\sigma) \cdot d_X(x'', x')$ for $\bar{x} \in \{x', x''\}$.

Since $x \stackrel{[x']_\rho}{\sim} x''$, the element x' cannot be closer to x than x'' is, so $d_X(x, x'') \leq d_X(x, x')$.

Since $x' \stackrel{[x']_\sigma}{\sim} x''$, the element x cannot be closer to x' than x'' is, so $d_X(x'', x') \leq d_X(x, x')$.

We now estimate $d_Y(f(x), f(x'))$ from above, using the triangle inequality. The tricky part is that the definition of local sensitivity does not allow to use $c_\sigma([x']_\sigma)$ as a multiplier for $d_X(x'', x')$, and $c_\rho([x']_\rho)$ as a multiplier for $d_X(x'', x)$. However, we would like to get new sensitivity $c_{\rho\sigma}$ that would depend either only on x , or only on x' . Since the intermediate point x'' shares its σ -class with x , we have $c_\sigma([x'']_\sigma) = c_\sigma([x]_\sigma)$. Taking $\hat{x} := x$ and $\bar{x} := x''$ gives us

$$\begin{aligned} d_Y(f(x), f(x')) &\leq d_Y(f(x), f(x'')) + d_Y(f(x''), f(x')) \\ &\leq c_\rho([x]_\rho) \cdot d_X(x, x'') + c_\sigma([x'']_\sigma) \cdot d_X(x'', x') \\ &\leq c_\rho([x]_\rho) \cdot d_X(x, x') + c_\sigma([x'']_\sigma) \cdot d_X(x, x') \\ &= c_\rho([x]_\rho) \cdot d_X(x, x') + c_\sigma([x]_\sigma) \cdot d_X(x, x') \\ &= (c_\rho + c_\sigma)([x]_{\rho \sqcap \sigma}) \cdot d_X(x, x') . \end{aligned}$$

Here the function addition is defined as $(c_\rho + c_\sigma)([x]_{\rho \sqcap \sigma}) := c_\rho([x]_\rho) + c_\sigma([x]_\sigma)$. The definition is correct, since the classes $[x]_\rho$ and $[x]_\sigma$ are uniquely determined by $[x]_{\rho \sqcap \sigma}$. \square

Sensitivity w.r.t. basis change. Let us now assume that the distances in X are ℓ_p -norms, and their generalization to sets is a Hamming distance – we want that the vector distances would be subadditive w.r.t. distances of their coordinates.

Let $X = (X_1, \dots, X_n)$ be a set of vectors $\vec{x} = (x_1, \dots, x_n)$, where the coordinates x_i are independent. Suppose that we know the sensitivity c_i of some function $f : X \rightarrow Y$ w.r.t. each component ρ_i , which is the i -th coordinate of \vec{x} . We want to find the sensitivities of the same function w.r.t. some *other* set of components $\vec{z} = (z_1, \dots, z_m)$, such that x_i is uniquely determined by z_1, \dots, z_m for all $i \in [n]$. We may think that the vector \vec{z} represents m equivalence classes of \vec{x} w.r.t. m different components, such that the m classes uniquely determine \vec{x} . Let $g : Z \rightarrow X$ be such that $\vec{x} = g(\vec{z})$. We are interested in sensitivity of the function $f \circ g$ w.r.t. the component σ_i that corresponds to the i -th coordinate of \vec{z} .

First of all, we may take Theorem 3.31 and extend it to *local* sensitivity.

Theorem 3.58. *Let $g : Z \rightarrow X$ be locally c_g sensitive w.r.t. component σ . Let $f : X \rightarrow Y$ be locally c_f sensitive. Then, the function composition $f \circ g$ is $\lambda Z' \cdot c_g(Z') \cdot c_f(\sup\{c_f(g(z)) \mid z \in Z'\})$ sensitive w.r.t. σ , using any definition of sensitivity (A-, B-, C-, D-).*

Proof. If global sensitivity was considered as in Theorem 3.31, we would get global sensitivity $c_g \cdot c_f$ in all four cases. For local sensitivity, this value would also depend on the particular input. First, assume that only c_g is local, and c_f is global. For $z \in Z$, we would directly get $c_g([z]_\sigma) \cdot c_f$ (A- and B- sensitivity). We would also get $c_g(Z') \cdot c_f$ for a subset of a class $Z' \in Z/\sigma$ (C- and D- sensitivity).

We may actually obtain a better bound than the global sensitivity c_f , since $g(z)$ does not take arbitrary values, and the set of possible inputs for f is constrained. Knowing that $z \in Z'$, the set of possible inputs for c_f is $\{g(z)|z \in Z'\}$. We take sup of c_f applied to this set. \square

A similar bound can be obtained if we treat g as a set functions g_1, \dots, g_m , each g_i mapping z to a particular equivalence class of x . This is stated in the next theorem.

Theorem 3.59. *Let ρ_1, \dots, ρ_n be pairwise adjacent equidistant components of a set X , such that $X/\rho_1 \times \dots \times X/\rho_n \sim X$, and $h : X/\rho_1 \times \dots \times X/\rho_n \rightarrow X$ is the corresponding natural isomorphism. Let $f : X \rightarrow Y$, and let c_i be the local sensitivity of f w.r.t. ρ_i for all $i \in [n]$. Let Z be a set, and let $g_i : Z \rightarrow X/\rho_i$ be functions that satisfy $\forall x \in X \exists z \in Z \forall i \in [n] : [x]_{\rho_i} = g_i(z)$. Let $\sigma_1, \dots, \sigma_m$ be pairwise adjacent equidistant components of Z , and c_j^i the sensitivity of g_i w.r.t. σ_j . Let $g : Z \rightarrow X$ be defined as $g(z) := h(g_1(z), \dots, g_n(z))$. The sensitivity of the function $f \circ g$ w.r.t. the component σ_j is $c(Z') = \sum_{i=1}^n \sup\{c_i(g_i(z))|z \in Z'\} \cdot c_j^i(Z')$.*

Proof. Let local sensitivity w.r.t. component be defined as in Def. 3.44. Since the components ρ_i and σ_i are equidistant, A- and B-sensitivity are equivalent, so without loss of generality, let us consider A-sensitivity. We have

$$x \stackrel{[x']_{\rho_i}}{\sim} x' \implies d_Y(f(x), f(x')) \leq c_i([x]_{\rho_i}) \cdot d_X(x, x') .$$

Since the components ρ_i are pairwise adjacent, for any pair $x, x' \in X$, there is a sequence of $n + 1$ elements x^i starting with $x^1 = x$ and ending with $x^{n+1} = x'$, where $[x^i]_{\rho_i} = [x]_{\rho_i}$, $[x^{i+1}]_{\rho_i} = [x']_{\rho_i}$, and $x^i \stackrel{[x']_{\rho_i}}{\sim} x^{i+1}$ for all $i \in [n]$. Using triangle inequality, for all $x, x' \in X$, we get

$$d_Y(f(x), f(x')) \leq \sum_{i=1}^n d_Y(f(x^i), f(x^{i+1})) \leq \sum_{i=1}^n c_i([x^i]_{\rho_i}) \cdot d_X(x^i, x^{i+1}) .$$

Since the components ρ_i are equidistant, and $x^i \stackrel{[x']_{\rho_i}}{\sim} x^{i+1}$, we have $d_X(x^i, x^{i+1}) = \tilde{d}_X([x]_{\rho_i}, [x']_{\rho_i})$ for all $i \in [n]$. On the other hand, there exist $z, z' \in Z$ such that $g_i(z) = [x]_{\rho_i}$ and $g_i(z') = [x']_{\rho_i}$ for all $i \in [n]$. Together, this gives us

$$\begin{aligned} d_Y(f(x), f(x')) &\leq \sum_{i=1}^n c_i([x]_{\rho_i}) \cdot d_X(x^i, x^{i+1}) \\ &= \sum_{i=1}^n c_i([x]_{\rho_i}) \cdot \tilde{d}_X([x]_{\rho_i}, [x']_{\rho_i}) \\ &= \sum_{i=1}^n c_i([x]_{\rho_i}) \cdot \tilde{d}_X(g_i(z), g_i(z')) . \end{aligned}$$

To proceed further, we need to know the sensitivities of functions g_i w.r.t. distances $\tilde{d}_X(g_i(z), g_i(z'))$. Since we estimate the local sensitivity of $f \circ g$ w.r.t. the component σ_j , we assume $z \stackrel{[z']_{\sigma_j}}{\sim} z'$. Let c_j^i be the local sensitivity of g_i w.r.t. σ_j . By definition of local sensitivity w.r.t. component, we have

$\tilde{d}_X(g_i(z), g_i(z')) \leq c_j^i([z]_{\sigma_j}) \cdot d_Z(z, z')$. We get

$$\begin{aligned} d_Y(f(x), f(x')) &\leq \sum_{i=1}^n c_i([x]_{\rho_i}) \cdot \tilde{d}_X(g_i(z), g_i(z')) \\ &\leq \sum_{i=1}^n c_i([x]_{\rho_i}) \cdot c_j^i([z]_{\sigma_j}) \cdot d_Z(z, z') \\ &\leq \sum_{i=1}^n c_i(g_i(z)) \cdot c_j^i([z]_{\sigma_j}) \cdot d_Z(z, z') . \end{aligned}$$

By definition, $g(z) = h(g_1(z), \dots, g_n(z)) = h([x]_{\rho_1}, \dots, [x]_{\rho_n})$, and since h is the natural isomorphism, we have $h([x]_{\rho_1}, \dots, [x]_{\rho_n}) = x$. Hence, we may substitute $x = g(z)$, $x' = g(z')$, and get

$$d_Y(fg(z), fg(z')) \leq \sum_{i=1}^n c_i(g_i(z)) \cdot c_j^i([z]_{\sigma_j}) \cdot d_Z(z, z') .$$

This still does not look like local sensitivity w.r.t. component, since $c_i(g_i(z))$ depends on z , not $[z]_{\sigma_j}$. The bound is nevertheless valid and can be used to compute noise for differential privacy, but it is not composable. To make it composable, we may define a valid upper bound $\hat{c}_i(Z') = \sup\{c_i(g_i(z)) \mid z \in Z'\}$ for $Z' \in Z/\sigma_j$, that now depends only on the σ_j -class of z . The sensitivity of $f \circ g$ w.r.t. the component σ_j will then be the function $c(Z') = \sum_{i=1}^n \hat{c}_i(Z') \cdot c_j^i(Z')$. \square

The same derivation could be applied to global sensitivity. In that case, the sensitivity would be a constant, and we would get

$$d_Y(fg(z), fg(z')) \leq \left(\sum_{i=1}^n c_i \cdot c_j^i \right) \cdot d_Z(z, z') .$$

Example 3.5. The function g may express transition to polar coordinates: $(x_1, x_2) := (\sqrt{z_1^2 + z_2^2}, \arctan \frac{z_1}{z_2})$. We have $c_1^1 = c_2^1 = 1$, and the functions c_1^2, c_2^2 are more complicated. If we apply to the radius some function f whose sensitivity c_f we already know, we may compute how the output actually depends on the particular coordinates. For If we have $f(x_1, x_2) = 2 \cdot x_1$, then $c_1 = 2$ and $c_2 = 0$. In this case, we do not need to estimate c_1^2 and c_2^2 since they are not used by f . We get that the sensitivity w.r.t. coordinate z_j is $c_1 \cdot c_j^1 = 2$. This means that, if we shift the value of z_1 or z_2 by k , then the final output will differ in at most $2k$. Indeed, we cannot achieve better results for z_1 in this particular case. If $z_2 = 0$, then we have $y = z_1$, so z_1 affects the output as well as y .

Derivative sensitivity w.r.t. component. If we define a multivariate function $f : X_1 \times \dots \times X_n \rightarrow \mathbb{R}$ for some X_1, \dots, X_n , intuitively, we would like the partial derivative $\frac{df_{\vec{x}}}{dX_i}$ (assuming that it is defined in X_i and exists on the point \vec{x}) to be related to the derivative sensitivity w.r.t. component X_i , since this value describes how the function output changes with X_i . This is motivated by Definition 3.39 of derivative sensitivity for row multiplicities, where $X_i = \mathbb{R}$ for all i . Like Definition 3.39, it depends on $\vec{x} \in (X_1, \dots, X_n)$, and, hence, on the other components X_j for $j \neq i$.

In some cases, when only the equivalence class of \vec{x} is known, we may want that the other components should remain “invisible”. This can be necessary e.g. in the case where the party who chooses a suitable ϵ for differential privacy sees only some part of the database. An upper bound that depends only on $x_i \in X_i$ would be $\sup_{\vec{z} \in (X_1, \dots, X_n) : (z_i = x_i)} \left\| \frac{df_{\vec{z}}}{dX_i} \right\|$, defined as ∞ if the supremum does not exist. We see that we would get a more precise bound having an entire vector \vec{x} , and in practice it may be more reasonable to compute $\frac{df_{\vec{x}}}{dX_i}$ for an actual database \vec{x} using privacy-preserving computation methods.

Complementary components. Let us now consider univariate $f : X \rightarrow \mathbb{R}$ (the space X can be multi-dimensional as well). We want to define sensitivity w.r.t. some abstract component ρ of X . If we want to apply the notion of derivative w.r.t. $[x]_\rho$, we first need to rewrite the function f in such a way that, instead of one variable $x \in X$, it takes at least two variables, one of which is $[x]_\rho$, and the other one contains all information that is sufficient to reconstruct x and evaluate f on it. Since we only have ρ , we need to define the other component (let us denote in $\bar{\rho}$) ourselves. It should always be possible to reconstruct x from $[x]_\rho$ and $[x]_{\bar{\rho}}$, and on the other hand, both $[x]_\rho$ and $[x]_{\bar{\rho}}$ should be necessary to reconstruct x , so that sensitivity w.r.t. ρ would not be 0. Taking into account these two properties, we get the following definition.

Definition 3.45 (complementary component). Let ρ be a component of X . The *complementary component* of ρ , denoted $\bar{\rho}$, is the component of X that satisfies $X/(\rho \sqcap \bar{\rho}) \sim X$ and $X/(\rho \sqcup \bar{\rho}) \sim \{X\}$.

In this definition, $\rho \sqcap \bar{\rho}$ is the finest possible component (each element in its own class), and $\rho \sqcup \bar{\rho}$ is the coarsest possible (all elements are in the same class). The component $\bar{\rho}$ that satisfies this definition is in general not unique.

Example 3.6. Let $X = \{a, b, c\}$, and $(a, b) \in \rho$, but $(b, c), (a, c) \notin \rho$. The two possible candidates for $\bar{\rho}$ are:

- σ_1 such that $(a, c) \in \sigma_1, (a, b), (b, c) \notin \sigma_1$;
- σ_2 such that $(b, c) \in \sigma_2, (a, b), (a, c) \notin \sigma_2$.

The complementary component $\bar{\rho}$ is not unique even for the simple case of vectors with ℓ_p norms. If we have $\vec{x} = (x_1, x_2)$, and ρ is such that $[\vec{x}]_\rho = x_1$, then we may choose $[\vec{x}]_{\bar{\rho}} = x_2$ as well as $[\vec{x}]_{\bar{\rho}} = x_1 + x_2$. In the second case, any function will first need to compute $x_2 = (x_1 + x_2) - x_1$. The choice of $\bar{\rho}$ does not affect the actual sensitivity w.r.t. ρ in any way, but formally, the derivatives are different. The problem is that $[\vec{x}]_{\bar{\rho}} = x_1 + x_2$ changes with x_1 , while for $[\vec{x}]_{\bar{\rho}} = x_2$ it remains the same.

We see that the definition of sensitivity w.r.t. ρ is ambiguous, dependent on the particular choice of $\bar{\rho}$. The choice of $\bar{\rho}$ should depend on the distance that defines the relation $x \stackrel{\rho}{\sim} x'$, describing what we actually want to hide. In Example 3.6, $x \stackrel{\rho}{\sim} x'$ holds for $d_X((x_1, x_2), (x'_1, x_2)) = |x_1 - x'_1|$, and in general not for $d_X((x_1, x_1 + x_2), (x'_1, x'_1 + x_2)) = 2|x_1 - x'_1|$. We want that the definition of $\bar{\rho}$ would be somehow related to the distance $d_X(\cdot, \cdot)$. We may remove the property $X/(\rho \sqcup \bar{\rho}) \sim \{X\}$ if we replace it with an alternative property that prevents $X/\bar{\rho} \subseteq X/\rho$.

Definition 3.46 (orthogonal complement). Let ρ be a component of X . The *orthogonal complement* of ρ , denoted $\bar{\rho}$, is a component of X that satisfies the following properties:

- $X/(\rho \sqcap \bar{\rho}) \sim X$;
- $x \stackrel{\rho}{\sim} x' \implies [x]_{\bar{\rho}} = [x']_{\bar{\rho}}$.

The orthogonal complement does not exist in general. It does not exist for the components on the set $X = \{a, b, c\}$ considered in Example 3.6. Moreover, the orthogonal complement is still not unique. However, if there are several orthogonal complements, any of them would give a valid bound for x and x' that satisfy $x \stackrel{\rho}{\sim} x'$.

The finest orthogonal complement can be constructed from ρ as follows. First, define $x \bar{\rho} x'$ and $x' \bar{\rho} x$ for all $x \stackrel{\rho}{\sim} x'$, and then extend the relation to transitive closure. The component $\bar{\rho}$ definitely cannot be finer, since all relations $x \bar{\rho} x'$ are necessary for a component to be an orthogonal complement, so if $X/(\rho \sqcap \bar{\rho}) \sim X$ does not hold, then the orthogonal complement does not exist.

Definition 3.47. Let ρ be a component of X . The *minimal orthogonal complement* of ρ , denoted $\bar{\rho}$, is the component of X that is defined as follows:

- $x \bar{\rho} x$ for $x \in X$;
- $x \bar{\rho} x', x' \bar{\rho} x$ for all $x \stackrel{\rho}{\sim} x'$;

- $x \bar{\rho} x' \wedge x' \bar{\rho} x'' \implies x' \bar{\rho} x''$ for all $x, x', x'' \in X$.

The definition is correct since $\bar{\rho}$ is by construction an equivalence relation. We need to show that it is an orthogonal complement.

Conjecture 3.60. The minimal orthogonal complement $\bar{\rho}$ of ρ is an orthogonal complement of ρ whenever $X/(\rho \sqcap \bar{\rho}) \sim X$.

Proof. The condition $X/(\rho \sqcap \bar{\rho}) \sim X$ is an assumption. Since $x \bar{\rho} x'$ and $x' \bar{\rho} x$ are defined for all $x \stackrel{\rho}{\sim} x'$, we have $x \stackrel{\rho}{\sim} x' \implies [x]_{\bar{\rho}} = [x']_{\bar{\rho}}$. \square

We use the notion of minimal orthogonal complement to define derivative sensitivity w.r.t. component.

Definition 3.48 (A-derivative sensitivity w.r.t. component). Let $f : X \rightarrow \mathbb{R}$. Let ρ be a component of X . Let $\sigma := \bar{\rho}$ be the minimal orthogonal complement of ρ , and ψ the isomorphism $X/(\rho \sqcap \sigma) \rightarrow X$. The *derivative sensitivity w.r.t. component ρ* is the following mapping from X to \mathbb{R}^+ :

$$DS_f^\rho(x) = \left\| \frac{d(f \circ \psi)_{([x]_\rho, [x]_\sigma)}}{d\rho} \right\| .$$

Definition 3.49 (B-derivative sensitivity w.r.t. component). Let $f : X \rightarrow \mathbb{R}$. Let ρ be a component of X . Let $\sigma := \bar{\rho}$ be the minimal orthogonal complement of ρ , and ψ the isomorphism $X/(\rho \sqcap \sigma) \rightarrow X$. The *derivative sensitivity w.r.t. component ρ* is the following mapping from X/ρ to \mathbb{R}^+ :

$$\hat{D}S_f^\rho(X^\rho) = \sup_{x \in X^\rho} \left\| \frac{d(f \circ \psi)_{([x]_\rho, [x]_\sigma)}}{d\rho} \right\| .$$

Intuitively, A- and B-derivative sensitivity w.r.t. component ρ can be used to make a function f A- and B-differentially private w.r.t. ρ , using e.g. the same construction as in Theorem 3.53. For this, we need to define more formally what the derivatives w.r.t. $[x]_\rho$ are. We can use Theorem 3.53 directly only if the set X/ρ is isomorphic to \mathbb{R} . Alternatively, we can generalize Theorem 3.53. We will do it in Sec. 3.3.8.2 where we provide more concrete definitions of components and derivative sensitivity w.r.t. them.

Complementary components that depend on the inputs. As discussed in Example 3.6, the minimal orthogonal complement of ρ on $X = \{a, b, c, \}$, where $(a, b) \in \rho$, $(b, c) \notin \rho$, $(a, c) \notin \rho$ does not exist. The problem is that the two candidate components σ_1 and σ_2 are only partially suitable. For $a \stackrel{\rho}{\sim} c$, we have $[a]_{\sigma_1} = [c]_{\sigma_1}$, and for $b \stackrel{\rho}{\sim} c$, we have $[b]_{\sigma_2} = [c]_{\sigma_2}$, but neither σ_1 nor σ_2 satisfies both. In practice, we could find separately the sensitivity for $a \stackrel{\rho}{\sim} c$ and $b \stackrel{\rho}{\sim} c$ using different definitions of $\bar{\rho}$, and then find an upper bound that satisfies both.

The construction of such set of components is always possible. For all $x, x' \in X$ satisfying $x \stackrel{\rho}{\sim} x'$, define a unique component $\bar{\rho}_{x, x'}$ such that $y \bar{\rho}_{x, x'} y' \iff (x = y \wedge x' = y') \vee (x = y' \wedge x' = y)$. Each $\bar{\rho}_{x, x'}$ puts together only elements that have been in different components of ρ , so we have $\forall i \in X \times X : X/(\rho \sqcap \bar{\rho}_i) \sim X$ and $x \stackrel{\rho}{\sim} x' \implies \exists i \in X \times X : [x]_{\bar{\rho}_i} = [x']_{\bar{\rho}_i}$.

Definition 3.50. Let ρ be a component of X . The *minimal orthogonal complement set* of ρ , denoted $(\bar{\rho}_i)_{i \in X \times X}$, is a set of components of X defined as follows:

- $x \bar{\rho}_i x$ for $x \in X, i \in X \times X$;
- $x \bar{\rho}_{x, x'} x'$ and $x' \bar{\rho}_{x, x'} x$ for all $x \stackrel{\rho}{\sim} x'$;
- $x \bar{\rho}_i x' \wedge x' \bar{\rho}_i x'' \implies x' \bar{\rho}_i x''$ for all $x, x', x'' \in X, i \in X \times X$.

We can now give a generalization of Definition 3.49.

Definition 3.51 (B-derivative sensitivity w.r.t. component (generalization)). Let $f : X \rightarrow \mathbb{R}$. Let ρ be a component of X . Let $(\sigma_i)_{i \in X \times X}$ be the minimal orthogonal complement set of ρ , and ψ_i for $i \in X \times X$ the isomorphism $X/(\rho \sqcap \sigma_i) \rightarrow X$. The *derivative sensitivity w.r.t. component ρ* is the following mapping from X/ρ to \mathbb{R}^+ :

$$\tilde{DS}_f^\rho(X^\rho) = \sup_{i \in X \times X, x \in X^\rho} \left\| \frac{d(f \circ \psi_i)_{([x]_\rho, [x]_\sigma)}}{d\rho} \right\| .$$

While Definition 3.51 looks complicated, the idea is still the same as for generalization of A-sensitivity to B-sensitivity. Namely, given $[x]_\rho$, we compute A-sensitivity for all possible $z \in X$ such that $[z]_\rho = [x]_\rho$ (each z may have its own complementary component σ_z), and take the worst-case (i.e. largest) derivative sensitivity over $z \in X$.

3.3.8.2 Derivative Sensitivity for Banach Spaces. If we want to use derivative sensitivity w.r.t. components in practice, it is not sufficient just to have the definitions. We also need a way to efficiently compute the derivative sensitivity, and convert it to a differential privacy mechanism. In this section, we show how it can be done for *Banach spaces*.

Banach spaces. First, we recall some basics of Banach space theory. Banach spaces do not allow arbitrary metrics and instead require norms but many useful metrics can also be viewed as norms. We will further denote vectors by \vec{x} , and norms by $\|\cdot\|_N$, where N specifies the particular norm. Formal definition of a norms and a seminorm is given in Definition 3.14.

Definition 3.52 (Banach space). A *Banach space* is a vector space with a norm that is *complete* (i.e. each converging sequence has a limit).

Banach spaces combine vector spaces with distances, which are necessary for defining differential privacy. The completeness property allows us to define derivatives. Using the norm of a Banach space, we may generalize the notion of continuous function from real numbers to Banach spaces.

Definition 3.53 (Continuous function in Banach space). Let V and W be Banach spaces, and $U \subset V$ an open subset of V . A function $f : U \rightarrow W$ is called *continuous* if

$$\forall \epsilon > 0 : \exists \delta > 0 : \|x - x'\|_V \leq \delta \Rightarrow \|f(x) - f(x')\|_W \leq \epsilon.$$

The notion of the derivative of a function can be also extended to Banach spaces.

Definition 3.54 (Fréchet derivative). Let V and W be Banach spaces, and $U \subset V$ an open subset of V . A function $f : U \rightarrow W$ is called *Fréchet differentiable at $x \in U$* if there exists a bounded linear operator $df_x : V \rightarrow W$ such that $\lim_{h \rightarrow 0} \frac{\|f(x+h) - f(x) - df_x(h)\|_W}{\|h\|_V} = 0$. Such operator df_x is called Fréchet derivative of f at the point x .

The mean value theorem can be generalized to Banach spaces (to a certain extent).

Theorem 3.61 (Mean value theorem ([55], Chapter XII)). *Let V and W be Banach spaces, and $U \subset V$ an open subset of V . Let $f : U \rightarrow W$, and let $x, x' \in U$. Assume that f is defined and is continuous at each point $(1-t)x + tx'$ for $0 \leq t \leq 1$, and differentiable for $0 < t < 1$. Then there exists $t^* \in (0, 1)$ such that*

$$\|f(x) - f(x')\|_W \leq \|df_z\|_{V \rightarrow W} \|x - x'\|_V$$

for $z = (1 - t^*)x + t^*x'$, where $\|\cdot\|_{V \rightarrow W}$ denotes the norm of operator that maps from V to W .

Derivative sensitivity in Banach spaces. Recall that, for row multiplicities, we defined derivative sensitivity as $DS_f(\vec{x}) = \max_i \left| \frac{\partial f}{\partial x_i}(\vec{x}) \right|$, where x_i denotes the i -th component of the vector of variables \vec{x} (Definition 3.39). We extend it to the case where X is any Banach space:

Definition 3.55. Let X be (an open convex subset of) a Banach space. Let $f : X \rightarrow \mathbb{R}$. Let f be Fréchet differentiable at each point of X . The *derivative sensitivity* of f is the following mapping from X to \mathbb{R}_+ , where \mathbb{R}_+ denotes the set of all non-negative real numbers:

$$DS_f(\vec{x}) = \|df_{\vec{x}}\| .$$

where $df_{\vec{x}}$ is the Fréchet derivative of f at \vec{x} and $\|df_{\vec{x}}\|$ is the operator norm of $df_{\vec{x}}$.

We generalize the results of Sec. 3.3.7 from \mathbb{R} to Banach spaces. First, we extend the definition of smoothness to the case where X is any Banach space:

Definition 3.56. Let $p : X \rightarrow \mathbb{R}$ and $\beta \in \mathbb{R}$. The mapping p is β -smooth, if $p(\vec{x}) \leq e^{\beta \cdot \|\vec{x} - \vec{x}'\|} \cdot p(\vec{x}')$ for all $\vec{x}, \vec{x}' \in X$.

We provide generalizations of Theorem 3.53 and Theorem 3.56 to the case where X is a Banach space:

Theorem 3.62. Let $\gamma, b, \beta \in \mathbb{R}_+$, $\gamma > 1$. Let $\epsilon = (\gamma + 1)(b + \beta)$. Let η be a random variable distributed according to $GenCauchy(\gamma)$. Let c be a β -smooth upper bound on DS_f for a function $f : X \rightarrow \mathbb{R}$. Then $g(\vec{x}) : f(\vec{x}) + \frac{c(\vec{x})}{b} \cdot \eta$ is ϵ -differentially private.

Theorem 3.63. Let $b, \beta, \epsilon \in \mathbb{R}_+$, $b > 0$, $b + \beta \leq \epsilon$. Define $k = 1 + (\epsilon - b)/\beta$. Let $\delta = e^{-k}$. Let η be a random variable distributed according to $Lap(1)$. Let c be a β -smooth upper bound on DS_f for a function $f : X \rightarrow \mathbb{R}$, where X is Banach space and d_X is the distance corresponding to the norm of X . Define $g(\vec{x}) := f(\vec{x}) + \frac{c(\vec{x})}{b} \cdot \eta$. Then

- for any $\vec{x}_1, \vec{x}_2 \in X$, $(\epsilon \cdot L, 2\delta) \in d_{DP}(g(\vec{x}_1), g(\vec{x}_2))$, where $L = d_X(\vec{x}_1, \vec{x}_2)$;
- (in particular,) g is $(\epsilon, 2e^\epsilon \delta)$ -differentially private.

If, additionally for any two points $\vec{x}_1, \vec{x}_2 \in X$ there exists a shortest path h in X , such that c is monotonic along that path, then the factor “2” in previous statements can be removed.

The proofs of these theorems is analogous to the proofs of Theorem 3.53 and Theorem 3.56, so we do not repeat them here. The only difference is that we are using Definition 3.56 for smoothness, Definition 3.54 for differentiability, and the mean value theorem for Banach spaces and Fréchet derivative (Theorem 3.61).

The derivative sensitivity of Definition 3.55 depends on the actual input \vec{x} and hence corresponds to A-sensitivity w.r.t. component (Definition 3.48). It can be generalized to B-sensitivity (Definition 3.49) in a straightforward way by taking $\sup_{z \in X, [\vec{z}]_\rho = [\vec{x}]_\rho} (DS_f(\vec{z}))$. The partitioning ρ defines the norm $\|\cdot\|$ of the underlying Banach space. While we cannot construct a Banach space for *any* possible ρ , we can do it for certain classes of ρ .

From components to Banach spaces. Formally, a component is an equivalence relation ρ on a set X . This relation can be viewed as a partitioning of X to disjoint subsets. For these subsets, we can define a distance induced by the distance $d_X(\cdot, \cdot)$ on X . In our practical applications, the Hausdorff distance seems to be the most reasonable choice, since it gives an upper bound on the distance between the closest elements of two subsets, and in the context of databases it would mean that we are considering some change in one particular column, letting all other columns be as similar as possible (ideally, they should remain the same, but it is not always possible).

Let ρ_1, \dots, ρ_n be such that $X/(\rho_1 \sqcap \dots \sqcap \rho_n) \sim X$. In order to apply derivative sensitivity, we need that the partitionings X/ρ_i would form some Banach spaces, not necessarily over real numbers. The norm in this Banach space should correspond to the distance $\tilde{d}_X(\cdot, \cdot)$ defined on elements of X/ρ_i .

One possibility is to embed a metric space to a Banach space using *Kuratowski embedding*. For any metric space (M, d) , denoting by $C_b(M)$ the Banach space of all bounded continuous real-valued functions on M with the supremum norm, we can take the map $\Phi : M \rightarrow C_b(M)$ defined by $\Phi(x)(y) := d(x, y) - d(x_0, y)$ for some fixed $x_0 \in M$. This map transforms elements of M to elements of the Banach space $C_b(M)$, and it is an isometry. In our case $M := X/\rho$, and the distance is $\tilde{d}_X(\cdot, \cdot)$. The functions $f : X/\rho \rightarrow \mathcal{P}(\mathbb{R})$ should now be transformed to analogous functions $\hat{f} : C_b(X/\rho) \rightarrow \mathcal{P}(\mathbb{R})$ as $\hat{f}(X) := f(\Phi^{-1}(X))$. We may take a more elaborated isometry Φ if it is possible.

Example 3.7. Let X be a set of vectors of length n , equipped with ℓ_p -norm. Let ρ_i be defined in a way $x \rho_i x' \iff x_i = x'_i$. For all $X_1, X_2 \in X/\rho_i$, we have $\tilde{d}_X(X_1, X_2) = \max_{\vec{x} \in X_1, \vec{x}' \in X_2} (d_X(\vec{x}, \vec{x}')) = |x_i - x'_i|$, where x_i and x'_i do not depend on the choice of $\vec{x} \in X_1$ and $\vec{x}' \in X_2$. Fixing $\vec{x}_0 := [\vec{0}]_{\rho_i}$, we get $\Phi(X_1)(X_2) := \tilde{d}_X(X_1, X_2) - \tilde{d}_X([\vec{0}]_{\rho_i}, X_2) = |x_i - x'_i| - |0 - x'_i|$. Here each function $\Phi(X_1)$ is in one-to-one correspondence with a real number x_i , which is the same for all $\vec{x} \in X_1$. A simpler way to define an isometry would be $\Psi : X/\rho_i \rightarrow \mathbb{R}$ as $\Psi([x]_{\rho_i}) = x_i$, which is defined correctly since x_i is the same for all elements of $[x]_{\rho_i}$. This function is an isometry w.r.t. the distance $|\cdot|$ of \mathbb{R} since $\tilde{d}_X([x]_{\rho_i}, [x']_{\rho_i}) = |x_i - x'_i|$ for all $\vec{x}, \vec{x}' \in X$.

Coming up with an isometry Φ for a given component ρ and operating with smooth derivatives in resulting Banach space can be difficult in practice. Let us provide a simpler way of constructing Banach spaces, which is well suitable for databases. The following three lemmas give us the basic combinators for statements about derivative sensitivity, reducing the task of finding the derivative sensitivity of a particular function with respect to a particular norm on its domain, to a series of tasks from basic calculus.

Lemma 3.64. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$, and let \mathbb{R}^n be equipped with the norm ℓ_p . Then $\|df_{\vec{x}}\|$ is the ℓ_q -norm of the gradient vector $\nabla f(\vec{x})$, where $q = \frac{p}{p-1}$ (if $p = 1$ then $q = \infty$ and vice versa).

Proof. Let $\nabla f(\vec{x}) = (a_i)_{i=1}^n$. Assuming $a_i \neq 0$ for all i (otherwise remove the indices i for which $a_i = 0$ from the summations containing a_i):

$$\begin{aligned} |df_{\vec{x}}(\vec{y})| &= |\nabla f(\vec{x}) \cdot \vec{y}| \leq \sum_{i=1}^n |a_i| |y_i| \\ &= \sum |a_i|^{\frac{p}{p-1}} \cdot \frac{|y_i|}{|a_i|^{\frac{1}{p-1}}} \\ &\leq \left(\sum |a_i|^{\frac{p}{p-1}} \right) \left(\frac{\sum |y_i|^p}{\sum |a_i|^{\frac{p}{p-1}}} \right)^{\frac{1}{p}} \\ &= \left(\sum |a_i|^{\frac{p}{p-1}} \right)^{\frac{p-1}{p}} \left(\sum |y_i|^p \right)^{\frac{1}{p}} \\ &= \|\nabla f(\vec{x})\|_q \cdot \|\vec{y}\|_p \end{aligned}$$

for all $\vec{y} \in X$. The second inequality used here is the weighted power means inequality with exponents 1 and p . Equality is achievable (and not only for $\vec{y} = 0$): for example, by taking $y_i = |a_i|^{\frac{1}{p-1}}$. Thus $\|\nabla f(\vec{x})\|_q$ is the smallest value of c such that for all \vec{y} , $|df_{\vec{x}}(\vec{y})| \leq c \cdot \|\vec{y}\|_p$, i.e. it is the operator norm $\|df_{\vec{x}}\|$.

The cases $p = 1$ and $p = \infty$ can be achieved as limits of the general case. □

The ℓ_q -norm is the *dual norm* of the ℓ_p -norm; we denote q by $\text{dual}(p)$.

Lemma 3.65. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ have the derivative sensitivity g with respect to the norm N . Let $a \cdot N$ denote the scaling of the output of the norm N by $a \in \mathbb{R}_+$. Then f has the derivative sensitivity g/a with respect to $a \cdot N$.

Proof. The derivative sensitivity of f at \vec{x} is the operator norm of a particular linear operator $df_{\vec{x}}$. It is equal to the minimal possible c , such that for all vectors \vec{y} , the absolute value of $df_{\vec{x}}(\vec{y}) \in \mathbb{R}$ is at most c times larger than the norm $\|\vec{y}\|_N$. If we replace N with $a \cdot N$, then the norm $\|\vec{y}\|_{a \cdot N}$ is increased by a times. Hence we may now reduce c by a times and still have the inequality. \square

Lemma 3.66. (a) Let $(V_1, \|\cdot\|_{V_1})$ and $(V_2, \|\cdot\|_{V_2})$ be Banach spaces. Let $V = V_1 \times V_2$. Let for all $(v_1, v_2) \in V$,

$$\|(v_1, v_2)\|_V = \|(\|v_1\|_{V_1}, \|v_2\|_{V_2})\|_p$$

Then $(V, \|\cdot\|_V)$ is a Banach space.

(b) Suppose furthermore that a function $f : V \rightarrow \mathbb{R}$ is differentiable at each point of V . Fix a point $v = (v_1, v_2) \in V$. Let $g : V_1 \rightarrow \mathbb{R}$ be such that $g(x_1) = f(x_1, v_2)$ and $h : V_2 \rightarrow \mathbb{R}$ be such that $h(x_2) = f(v_1, x_2)$. Let $c_1 = DS_g(v_1)$ and $c_2 = DS_h(v_2)$. Then $DS_f(v) = \|(c_1, c_2)\|_q$ where $\|\cdot\|_q$ is the dual norm of $\|\cdot\|_p$.

Proof. (a) We first prove that $(V, \|\cdot\|_V)$ is a normed vector space. We prove only the triangle inequality. The rest of the properties of norm are easy to check.

$$\begin{aligned} \|(v_1, v_2) + (v'_1, v'_2)\|_V &= \|(v_1 + v'_1, v_2 + v'_2)\|_V = \|(\|v_1 + v'_1\|_{V_1}, \|v_2 + v'_2\|_{V_2})\|_p \leq \\ &\leq \|(\|v_1\|_{V_1} + \|v'_1\|_{V_1}, \|v_2\|_{V_2} + \|v'_2\|_{V_2})\|_p \leq \\ &\leq \|(\|v_1\|_{V_1}, \|v_2\|_{V_2})\|_p + \|(\|v'_1\|_{V_1}, \|v'_2\|_{V_2})\|_p = \\ &= \|(v_1, v_2)\|_V + \|(v'_1, v'_2)\|_V \end{aligned}$$

The first inequality uses the triangle inequalities of $\|\cdot\|_{V_1}$ and $\|\cdot\|_{V_2}$ and the monotonicity of $\|\cdot\|_p$ in the absolute values of the coordinates of its argument vector. The second inequality uses the triangle inequality of $\|\cdot\|_p$.

Thus $(V, \|\cdot\|_V)$ is a normed vector space. We now prove that it is a Banach space. Consider a Cauchy sequence $\{x_n\}$ in V . Then

$$\forall \epsilon > 0. \exists N \in \mathbb{N}. \forall m, n > N. \|x_m - x_n\|_V < \epsilon$$

Let $x_n = (y_n, z_n)$ where $y_n \in V_1$ and $z_n \in V_2$. Note that

$$\|y_m - y_n\|_{V_1} = \|(y_m - y_n, 0)\|_V \leq \|(y_m - y_n, z_m - z_n)\|_V = \|x_m - x_n\|_V$$

Thus

$$\forall \epsilon > 0. \exists N \in \mathbb{N}. \forall m, n > N. \|y_m - y_n\|_{V_1} < \epsilon$$

i.e. $\{y_n\}$ is a Cauchy sequence in V_1 . Because V_1 is a Banach space, there exists $y \in V_1$ such that

$$\lim_{n \rightarrow \infty} \|y_n - y\|_{V_1} = 0$$

Similarly, we get that there exists $z \in V_2$ such that

$$\lim_{n \rightarrow \infty} \|z_n - z\|_{V_2} = 0$$

Let $x = (y, z)$. Note that

$$\|x_n - x\|_V = \|(y_n - y, z_n - z)\|_V = \|(\|y_n - y\|_{V_1}, \|z_n - z\|_{V_2})\|_p$$

Then, because $\|\cdot\|_p$ is continuous,

$$\lim_{n \rightarrow \infty} \|x_n - x\|_V = \|(\lim_{n \rightarrow \infty} \|y_n - y\|_{V_1}, \lim_{n \rightarrow \infty} \|z_n - z\|_{V_2})\|_p = \|(0, 0)\|_p = 0$$

Thus V is a Banach space.

(b)

$$c_1 = \|dg_{v_1}\|$$

$$c_2 = \|dh_{v_2}\|$$

Note that

$$\begin{aligned} & \lim_{x_1 \rightarrow 0_{V_1}} \frac{|g(v_1 + x_1) - g(v_1) - df_v(x_1, 0)|}{\|x_1\|_{V_1}} = \\ &= \lim_{x_1 \rightarrow 0_{V_1}} \frac{|f(v_1 + x_1, v_2) - f(v_1, v_2) - df_v(x_1, 0)|}{\|(x_1, 0)\|_V} = \\ &= \lim_{(x_1, 0) \rightarrow 0_V} \frac{|f(v + (x_1, 0)) - f(v) - df_v(x_1, 0)|}{\|(x_1, 0)\|_V} = \\ &= \lim_{x \rightarrow 0_V} \frac{|f(v + x) - f(v) - df_v(x)|}{\|x\|_V} = 0 \end{aligned}$$

The last equality holds by the definition of Fréchet derivative. The equality before that holds because the limit on the right-hand side exists. Then, again by the definition of Fréchet derivative, we get that the linear map that maps x_1 to $df_v(x_1, 0)$, is dg_{v_1} . Thus $df_v(x_1, 0) = dg_{v_1}(x_1)$. Similarly, we get $df_v(0, x_2) = dh_{v_2}(x_2)$. Now

$$\begin{aligned} |df_v(x_1, x_2)| &= |df_v(x_1, 0) + df_v(0, x_2)| = \\ &= |dg_{v_1}(x_1) + dh_{v_2}(x_2)| \leq |dg_{v_1}(x_1)| + |dh_{v_2}(x_2)| \leq \\ &\leq c_1 \|x_1\|_{V_1} + c_2 \|x_2\|_{V_2} \leq \|(c_1, c_2)\|_q \cdot \|(\|x_1\|_{V_1}, \|x_2\|_{V_2})\|_p = \\ &= \|(c_1, c_2)\|_q \cdot \|(x_1, x_2)\|_V \end{aligned}$$

The last inequality follows from the weighted power means inequality, similarly to the proof of Lemma 3.64. Equality is also achievable: because $c_1 = \|dg_{v_1}\|$ and $c_2 = \|dh_{v_2}\|$, there exist x_1 and x_2 that achieve equality in the second inequality. Then scale x_1 and x_2 by constants such that $\|x_1\|_{V_1}$ and $\|x_2\|_{V_2}$ (which scale by the same constants) achieve equality in the third inequality. To achieve equality in the first inequality, we may further need to multiply x_1 and/or x_2 by -1 . Thus $\|df_v\| = \|(c_1, c_2)\|_q$. \square

The following example demonstrates how the Lemmata 3.64, 3.65, 3.66 can be used to construct a suitable norm and compute derivative sensitivity in the corresponding Banach space.

Example 3.8. Consider the example of computing differentially privately the time that a ship takes to reach the port. This time can be expressed as

$$f(x, y, v) = \frac{\sqrt{x^2 + y^2}}{v}$$

$$f : \mathbb{R}^3 \rightarrow \mathbb{R}$$

where (x, y) are the coordinates of the ship (with the port at $(0, 0)$) and v is the speed of the ship. For differential privacy, we need to define distances on \mathbb{R}^3 and \mathbb{R} . Because we want to use Fréchet derivatives to compute the sensitivities, we instead define norms, which then also induce distances. For \mathbb{R} it is natural to use the absolute value norm but for \mathbb{R}^3 there are more choices.

First, let us consider the ℓ_1 -distance $\|\Delta x, \Delta y, \Delta v\|_1 = |\Delta x| + |\Delta y| + |\Delta v|$. Then moving the ship by geographical distance s in a direction parallel or perpendicular to its velocity, changes the whole input by ℓ_1 -distance s . But moving the ship in any other direction changes the whole input by ℓ_1 -distance more than s . This is unnatural. We would like the change not to depend on the direction.

Consider the ℓ_2 -distance $\|\Delta x, \Delta y, \Delta v\|_2 = \sqrt{|\Delta x|^2 + |\Delta y|^2 + |\Delta v|^2}$. Then moving the ship by geographical distance s always changes the whole input by ℓ_2 -distance s , regardless of direction. If we change the speed of the ship by u (either up or down) then the whole input changes by ℓ_2 -distance u . If we simultaneously move the ship by distance s and change its speed by u however, and numerically $s = u$ (ignoring the units) then the whole input changes by ℓ_2 -distance $\sqrt{2} \cdot s$ instead of the more natural $2s$.

Thus it is better to combine ℓ_1 - and ℓ_2 -distances. We first combine the change in the coordinates of the ship using ℓ_1 -norm, then combine the result with the change in speed using ℓ_2 -norm. Thus

$$\|(\Delta x, \Delta y, \Delta v)\| = \|(\|(\Delta x, \Delta y)\|_2, \Delta v)\|_1 = \sqrt{(\Delta x)^2 + (\Delta y)^2} + |\Delta v|$$

This still has a problem. Suppose that $x = y = 2000$ km, $v = 20$ km/h. Then changing v to 10 km/h changes the whole input by the same distance as changing y to 2010 km. But the former change seems much more important in most cases than the latter. Thus we scale the geographical change and the speed change by different constants before combining them:

$$\|(\Delta x, \Delta y, \Delta v)\| = a\sqrt{(\Delta x)^2 + (\Delta y)^2} + b|\Delta v| = \|(\|(a\Delta x, a\Delta y)\|_2, b\Delta v)\|_1$$

This norm should now be good enough.

Let us now compute the derivatives. We first compute ordinary partial derivatives, i.e. with respect to the absolute value norm.

$$\begin{aligned}\frac{\partial f}{\partial x} &= \frac{2x}{2v\sqrt{x^2 + y^2}} = \frac{x}{v\sqrt{x^2 + y^2}} \\ \frac{\partial f}{\partial y} &= \frac{2y}{2v\sqrt{x^2 + y^2}} = \frac{y}{v\sqrt{x^2 + y^2}} \\ \frac{\partial f}{\partial v} &= -\frac{\sqrt{x^2 + y^2}}{v^2}\end{aligned}$$

Now we compute the derivatives w.r.t. the scaled one-dimensional norms, i.e. x is considered not as an element of the Banach space $(\mathbb{R}, \|\cdot\|)$ but as an element of the Banach space $(\mathbb{R}, \|\cdot\|_x)$ where $\|\Delta x\|_x = a|\Delta x|$. The operator norm of the Fréchet derivative in $(\mathbb{R}, \|\cdot\|_x)$ (which is also the derivative sensitivity) is then

$$DS_{f^{y,v}}(x) = \|df_x^{y,v}\| = \left| \frac{\partial f}{a\partial x} \right| = \frac{|x|}{a|v|\sqrt{x^2 + y^2}}$$

where $f^{y,v}$ is the one-variable function defined by $f^{y,v}(x) = f(x, y, v)$, i.e. y and v are considered as constants. Similarly, we compute

$$DS_{f^{x,v}}(y) = \|df_y^{x,v}\| = \left| \frac{\partial f}{a\partial y} \right| = \frac{|y|}{a|v|\sqrt{x^2 + y^2}}$$

$$DS_{f^{x,y}}(v) = \|df_v^{x,y}\| = \left| \frac{\partial f}{b\partial v} \right| = \frac{\sqrt{x^2 + y^2}}{bv^2}$$

where $f^{x,v}$ and $f^{x,y}$ are the one-variable functions defined by $f^{x,v}(y) = f(x, y, v)$ and $f^{x,y}(v) = f(x, y, v)$.

Now we use Lemma 3.66 to combine the Banach spaces $(\mathbb{R}, \|\cdot\|_x)$ and $(\mathbb{R}, \|\cdot\|_y)$ into the Banach space $(\mathbb{R}^2, \|\cdot\|_{xy})$ where

$$\|(\Delta x, \Delta y)\|_{xy} = \|(\|\Delta x\|_x, \|\Delta y\|_y)\|_2 = \|(a\Delta x, a\Delta y)\|_2$$

We get

$$\|df_{(x,y)}^v\| = DS_{f^v}(x, y) = \|(DS_{f^{y,v}}(x), DS_{f^{x,v}}(y))\|_2 = \frac{\sqrt{2}}{a|v|}$$

Now we again use Lemma 3.66 to combine the Banach spaces $(\mathbb{R}^2, \|\cdot\|_{xy})$ and $(\mathbb{R}, \|\cdot\|_v)$ into the Banach space $(\mathbb{R}^2, \|\cdot\|_{xyv})$ where

$$\|(\Delta x, \Delta y, \Delta v)\|_{xyv} = a\sqrt{(\Delta x)^2 + (\Delta y)^2} + b|\Delta v| = \|(\|(a\Delta x, a\Delta y)\|_2, b\Delta v)\|_1$$

We get

$$\|df_{(x,y,v)}\| = DS_f(x, y, v) = \|(DS_{f^v}(x, y), DS_{f^{x,y}}(v))\|_\infty = \max\left\{\frac{\sqrt{2}}{a|v|}, \frac{\sqrt{x^2 + y^2}}{bv^2}\right\}$$

Thus we have found not only the derivative sensitivity of f (DS_f) but also its derivative sensitivities w.r.t. components ($DS_{f^y,v}, DS_{f^x,v}, DS_{f^x,y}, DS_{f^y}$). The sensitivities w.r.t. components depend on other components but this is not a problem because when computing $f(x, y, v)$ differentially privately, even w.r.t. a component, we need to use all of x, y, v anyway, so why not use all this information for computing the noise level too.

Database as a Banach space. Let us have a database with a number of tables. The schema for each table is fixed. We see that database as a point in some Banach space (thus the distance between databases is the norm of their difference), where each cell in each table corresponds to a dimension of that space. As dimensions of Banach spaces are fixed (indeed, they are vector spaces \mathbb{R}^n with some extra structure), the number of rows in each table is also fixed, and each row can be seen to have a public *identity* (similarly to [56]). In this work, we consider the following *composite* norms and seminorms as possible norms for databases:

Definition 3.57 (composite seminorm). Let $\|\cdot\|_N$ be a seminorm of the vector space \mathbb{R}^n . It is a *composite seminorm* if one of the following holds for all $\vec{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$:

- There exists $i \in [n]$, such that $\|\vec{x}\|_N = |x_i|$. Such seminorm *uses the variable* x_i .
- There exists a composite seminorm $\|\cdot\|_M$ and $a \in \mathbb{R}^+$, such that $\|\vec{x}\|_N = a \cdot \|\vec{x}\|_M$. The seminorm $\|\cdot\|_N$ *uses the same variables* as $\|\cdot\|_M$.
- There exist composite seminorms $\|\cdot\|_{M_1}, \dots, \|\cdot\|_{M_k}$ and $p \in [1, \infty]$, s.t. $\|\vec{x}\|_N = \|\|\vec{x}\|_{M_1}, \dots, \|\vec{x}\|_{M_k}\|_p$. The seminorm $\|\cdot\|_N$ uses the union of the variables used by all $\|\cdot\|_{M_i}$.

Let $\text{vars}(N)$ be the set of variables used by $\|\cdot\|_N$.

We normally define the norms for rows of each table using the constructions allowed in Def. 3.57. We then state that the norm of the table is some ℓ_p -norm of the vector of the norms of its rows (last item of Def. 3.57) and the norm of the database is some ℓ_p -norm of the vector of the norms of its tables (in this work, we usually assume that it is ℓ_1 -norm). Note that the notion of *seminorm* is used only for building blocks of composite ℓ_p -norms, and a seminorm constructed as in Def. 3.57 becomes a norm if $\text{vars}(N) = \{x_1, \dots, x_n\}$. Lemmata 3.64, 3.65, 3.66 cover all cases considered in Definition 3.57, thus providing a way to compute derivative sensitivity w.r.t. a composite norm.

In order to compute derivative sensitivity w.r.t. a database, we need to take into account that the query operates on a cross product of tables. The cells of this cross product table are correlated in a certain way, so computing derivative sensitivity for the database is more complicated than computing it for a single table. Let us show how this can be handled.

Suppose we have a database of n tables. Thus we can have a Banach space for each table with a certain norm. Let $T_i = R_i^{n_i}$ be the Banach space for the i^{th} table, where n_i is the number of rows in the i^{th} table and R_i is the Banach space for its row.

The input contains a tuple of n tables, which is an element of $D = T_1 \times \dots \times T_n$. We can make D a Banach space by combining the norms of T_i using the ℓ_1 -norm.

To make a query on the database we want to join those n tables. Consider an input $(t_1, \dots, t_n) \in D$. Let $t = t_1 \times \dots \times t_n$ be the cross product of tables. First, let us assume that the n joined tables are distinct, i.e. no table is used more than once. Each row of the cross product is an element of $R = R_1 \times \dots \times R_n$, thus t is an element of $T = R^N$ where $N = n_1 \dots n_n$.

The query contains an aggregating function $f : T \rightarrow \mathbb{R}$. We can consider $t \in T$ as one large table whose number of rows is the product of the numbers of rows of tables $t_i \in T_i$ and number of columns is the sum of the numbers of columns of tables $t_i \in T_i$. We know how to compute the derivative sensitivity of f w.r.t. the components of t specified by a subset of rows and a subset of columns of t .

Suppose we want to compute the derivative sensitivity of f w.r.t. a row r of t_i . We can compute the sensitivity w.r.t. the following component of t : the subset of rows is the set of rows affected by r , i.e. $t_r = t_1 \times \dots \times t_{i-1} \times \{r\} \times t_{i+1} \times \dots \times t_n$, the subset of columns is the set of columns corresponding to

table t_i . Because changing r by distance d changes each row of t_r by distance d , we must combine the rows of t_r by ℓ_∞ -distance to ensure that t_r also changes by distance d .

Suppose we now want to compute the sensitivity of f w.r.t. a subset s of rows of t_i . Then the subset of rows affected by s , is $t_s = t_1 \times \cdots \times t_{i-1} \times s \times t_{i+1} \times \cdots \times t_n$. Each row in t_s is affected by exactly one row in s . Let $s = r_1, \dots, r_k$ and let $t_s = \bigcup_{j=1}^k t_{r_j}$ where t_{r_j} is the subset of rows affected by r_j . The t_{r_j} are disjoint. Let $t_{r_j} = \{u_{j1}, \dots, u_{jm_j}\}$. Then the norm for u_{jm} is the same as the norm for r_j , with the additional columns having zero norm, i.e. not included into the norm. The norm for t_{r_j} is computed by combining the norms for u_{jm} using ℓ_∞ -norm. The norm for t_s is then computed by combining the norms for t_{r_j} using the norm that combined the norms of the rows of t_i , i.e. ℓ_{p_i} . If s is the entire set of all rows of t_i , we get derivative sensitivity w.r.t. the table t_i .

We have shown how to compute the derivative sensitivity w.r.t. each table t_1, \dots, t_n separately. If each table is used only once, since the norm of the database is defined as the ℓ_1 -norm of tables, by Lemma 3.64 we need to take the maximum of corresponding derivative sensitivities. If some table is used multiple times, these copies should be combined using ℓ_∞ -norm, so we have to sum up the derivative sensitivities of these copies.

3.3.8.3 Smoothing. To achieve differential privacy, we need to find a smooth upper bound on the derivative sensitivity. When differential privacy is required only w.r.t. a component then the sensitivity w.r.t. to the component only needs to be smoothed for changes in that component (changes in other components are allowed to change the smooth upper bound without smoothness restriction). After a smooth upper bound has been found, we can use Theorem 3.62 to compute f differentially privately.

Suppose we have a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and we want to find its β -smooth upper bound.

A differentiable function $f : \mathbb{R} \rightarrow \mathbb{R}$ is β -smooth if $\left| \frac{f'(x)}{f(x)} \right| \leq \beta$.

If DS_f exists then $f : X \rightarrow \mathbb{R}$ is β -smooth if $\frac{DS_f(x)}{|f(x)|} \leq \beta$.

We will now show how to compute β -smooth bounds on f and DS_f for certain functions f . First, let us state and prove some helpful lemmas.

Lemma 3.67 (a part of Lemma 2.3 of [52]). *Let $f : X \rightarrow \mathbb{R}$. For $\beta > 0$, a β -smooth upper bound on f is*

$$g_{f,\beta} = \max_{x' \in X} (f(x) \cdot e^{-\beta \cdot d(x,x')}) .$$

Lemma 3.68. *Let X_i for $i \in \{1, \dots, n\}$ be Banach spaces, $f_i : X_i \rightarrow \mathbb{R}$. Let $x = (x_1, \dots, x_n)$, and let $f(x_1, \dots, x_n) = \|f_1(x_1), \dots, f_n(x_n)\|_p$. Then $\frac{\partial f}{\partial x_i}(x) = \frac{\partial f_i}{\partial x_i}(x_i) \cdot \left(\frac{f_i(x_i)}{f(x)}\right)^{p-1} \leq \frac{\partial f_i}{\partial x_i}(x_i)$.*

Proof. Let $y_i = f_i(x_i)$ and $y = (y_1, \dots, y_n)$. We have

$$\frac{\partial f}{\partial x_i}(x) = \frac{\partial f}{\partial y_i}(y) \cdot \frac{\partial f_i}{\partial x_i}(x_i) = \left(\frac{y_i}{\|y\|_p}\right)^{p-1} \cdot \frac{\partial f_i}{\partial x_i}(x_i) .$$

Since $\frac{y_i}{\|y\|_p} = \frac{f_i(x_i)}{f(x)} = \frac{f_i(x_i)}{\|(f_i(x_i))_{i=1}^n\|_p}$, we have $\frac{f_i(x_i)}{f(x)} \leq 1$, and hence also $\left(\frac{f_i(x_i)}{f(x)}\right)^{p-1} \leq 1$, getting $\frac{\partial f}{\partial x_i}(x) \leq \frac{\partial f_i}{\partial x_i}(x_i)$. \square

Lemma 3.69. *Let X_i for $i \in \{1, \dots, n\}$ be Banach spaces. Let $x = (x_1, \dots, x_n)$, and let $f(x) = \|f_1(x), \dots, f_n(x)\|_p$. Then $\frac{\partial f}{\partial x_i}(x) = \sum_{j=1}^n \left(\frac{f_j(x)}{f(x)}\right)^{p-1} \cdot \frac{\partial f_j}{\partial x_i}(x)$. This can be upper bounded as:*

1. $\sum_{j=1}^n \frac{\partial f_j}{\partial x_i}(x)$;
2. $\max_{j=1}^n \frac{f_j(x)}{f_j(x)} \cdot \frac{\partial f_j}{\partial x_i}(x)$.

Proof. Let $y_j = f_j(x)$, $z = \sum_{j=1}^n y_j^p$. We have

$$\begin{aligned} \frac{\partial f}{\partial x_i}(x) &= \frac{\partial f}{\partial z}(z) \cdot \sum_{j=1}^n \left(\frac{\partial f_j(x)^p}{\partial y_j}(y_j) \cdot \frac{\partial f_j}{\partial x_i}(x) \right) \\ &= \sum_{j=1}^n \left(\frac{y_j}{\|y\|_p} \right)^{p-1} \cdot \frac{\partial f_j}{\partial x_i}(x) \\ &= \sum_{j=1}^n \left(\frac{f_j(x)}{f(x)} \right)^{p-1} \cdot \frac{\partial f_j}{\partial x_i}(x). \end{aligned}$$

As in the proof of Lemma 3.68, $\left(\frac{y_j}{\|y\|_p} \right)^{p-1} = \left(\frac{f_j(x)}{f(x)} \right)^{p-1} \leq 1$. We get $\frac{\partial f}{\partial x_i}(x) \leq \sum_{j=1}^n \frac{\partial f_j}{\partial x_i}(x)$. We can proceed with the inequality in another way.

$$\begin{aligned} \sum_{j=1}^n \left(\frac{f_j(x)}{f(x)} \right)^{p-1} \cdot \frac{\partial f_j}{\partial x_i}(x) &= \frac{\sum_{j=1}^n f_j(x)^{p-1} \cdot \frac{\partial f_j}{\partial x_i}(x)}{f(x)^{p-1}} \\ &= \frac{\sum_{j=1}^n f_j(x)^p \cdot \frac{\partial f_j}{\partial x_i}(x) \cdot \frac{1}{f_j(x)}}{f(x)^{p-1}} \\ &\leq \max_{j=1}^n \left(\frac{1}{f_j(x)} \cdot \frac{\partial f_j}{\partial x_i}(x) \right) \cdot \frac{\sum_{j=1}^n f_j(x)^p}{f(x)^{p-1}} \\ &= \max_{j=1}^n \left(\frac{1}{f_j(x)} \cdot \frac{\partial f_j}{\partial x_i}(x) \right) \cdot \frac{f(x)^p}{f(x)^{p-1}} \\ &= \max_{j=1}^n \left(\frac{f(x)}{f_j(x)} \cdot \frac{\partial f_j}{\partial x_i}(x) \right) \cdot \square \end{aligned}$$

Lemma 3.70. Let $f(x) : \mathbb{R} \rightarrow \mathbb{R}$ be β_f -smooth, and let $g(x) : \mathbb{R} \rightarrow \mathbb{R}$ be β_g -smooth.

1. If $f(x), g(x) > 0$, then $f(x) + g(x)$ is $\max(\beta_f, \beta_g)$ -smooth;
2. $f(x) \cdot g(x)$ is $\beta_f + \beta_g$ -smooth;
3. $f(x) / g(x)$ is $\beta_f + \beta_g$ -smooth.

Proof. We have:

1. $\left| \frac{(f(x)+g(x))'}{f(x)+g(x)} \right| = \frac{|f'(x)+g'(x)|}{|f(x)+g(x)|} \leq \frac{|f'(x)|+|g'(x)|}{|f(x)+g(x)|} \leq \max \left(\left| \frac{f'(x)}{f(x)} \right|, \left| \frac{g'(x)}{g(x)} \right| \right) \leq \max(\beta_f, \beta_g)$.
2. $\left| \frac{(f(x) \cdot g(x))'}{f(x) \cdot g(x)} \right| = \left| \frac{f'(x) \cdot g(x) + f(x) \cdot g'(x)}{f(x) \cdot g(x)} \right| \leq \left| \frac{f'(x)}{f(x)} \right| + \left| \frac{g'(x)}{g(x)} \right| \leq \beta_f + \beta_g$.
3. $\left| \frac{(f(x)/g(x))'}{f(x)/g(x)} \right| = \left| \frac{f'(x) \cdot g(x) - f(x) \cdot g'(x)}{g(x)^2} \cdot \frac{g(x)}{f(x)} \right| = \left| \frac{f'(x)}{f(x)} - \frac{g'(x)}{g(x)} \right| \leq \left| \frac{f'(x)}{f(x)} \right| + \left| \frac{g'(x)}{g(x)} \right| \leq \beta_f + \beta_g$.

□

Lemma 3.71. Let X_i for $i \in \{1, \dots, n\}$ be Banach spaces, and let $X = \prod_{i=1}^n X_i$. Let $f_i : X_i \rightarrow \mathbb{R}$ be β_i -smooth. Then, $f(x_1, \dots, x_n) = \|f_1(x_1), \dots, f_n(x_n)\|_p$ is $\|(\beta_i)_{i=1}^n\|_p$ -smooth as well as $\max_{i=1}^n(\beta_i)$ -smooth, where the norm of X is the $\ell_{\text{dual}(p)}$ -combination of the norms of all X_i .

Proof. Let $X = \prod_{i=1}^n X_i$ and $x = (x_1, \dots, x_n)$. Let $\beta = \max_i \beta_i$. By Lemma 3.68, an upper bound on $\frac{\partial f}{\partial x_i}(x)$ is $c_i(x) = f'_i(x_i)$. We have

$$|c_i(x)| = |f'_i(x_i)| \leq \text{DS}_{f_i}(x_i) = |f_i(x_i)| \cdot \frac{\text{DS}_{f_i}(x_i)}{|f_i(x_i)|} \leq |f_i(x_i)| \cdot \beta_i .$$

By Lemma 3.64 and Lemma 3.66, the derivative sensitivity of f in $(X, \ell_{\frac{p}{p-1}})$ is

$$\text{DS}_f(x) = \|(c_1(x), \dots, c_n(x))\|_p$$

Using inequality $|f_i(x_i)| \leq |f(x)|$, we get

$$\frac{\text{DS}_f(x)}{|f(x)|} \leq \frac{\|(|f_i(x_i)| \cdot \beta_i)_{i=1}^n\|_p}{|f(x)|} \leq \frac{|f(x)| \cdot \|(\beta_i)_{i=1}^n\|_p}{|f(x)|} \leq \|(\beta_i)_{i=1}^n\|_p .$$

On the other hand, using inequality $\beta_i \leq \beta$, we get

$$\frac{\text{DS}_f(x)}{|f(x)|} \leq \frac{\|(|f_i(x_i)| \cdot \beta)_{i=1}^n\|_p}{|f(x)|} \leq \frac{\beta \cdot \|(|f_i(x_i)|)_{i=1}^n\|_p}{|f(x)|} = \beta \quad \square$$

Lemma 3.72. Let X_i for $i \in \{1, \dots, n\}$ be Banach spaces, $X = \prod_{i=1}^n X_i$. Let $f_i : X \rightarrow \mathbb{R}$ be β_i^j -smooth for X_j . Let $x = (x_1, \dots, x_n)$. Then, $f(x) = \|f_1(x), \dots, f_n(x)\|_p$ is $\|(\max_j \beta_i^j)_{i=1}^n\|_p$ -smooth.

Proof. Let $X = \prod_{i=1}^n X_i$ and $x = (x_1, \dots, x_n)$. By Lemma 3.69, an upper bound on $\frac{\partial f}{\partial x_i}(x)$ is $c_i(x) = \max_j \frac{f_j(x)}{f_j(x)} \cdot \frac{\partial f_j}{\partial x_i}(x)$. We have

$$|c_i(x)| = \left| \max_j \frac{f_j(x)}{f_j(x)} \cdot \frac{\partial f_j}{\partial x_i}(x) \right| \leq |f(x)| \cdot \max_j \left| \frac{\partial f_j}{\partial x_i}(x) \cdot \frac{1}{f_j(x)} \right| \leq |f(x)| \cdot \max_j \beta_i^j .$$

By Lemma 3.64 and Lemma 3.66, the derivative sensitivity of f in $(X, \ell_{\text{dual}(p)})$ is

$$\text{DS}_f(x) = \|(c_1(x), \dots, c_n(x))\|_p$$

We get

$$\frac{\text{DS}_f(x)}{|f(x)|} \leq \frac{|f(x)| \cdot \|(\max_j \beta_i^j)_{i=1}^n\|_p}{|f(x)|} \leq \|(\max_j \beta_i^j)_{i=1}^n\|_p . \quad \square$$

We are now ready to find smooth upper bounds for certain functions.

Power function. Let $f(x) = x^r$, $r \in \mathbb{R}_+$, $x > 0$. We have

$$\frac{f'(x)}{f(x)} = \frac{rx^{r-1}}{x^r} = \frac{r}{x} ; \quad \left| \frac{r}{x} \right| \leq \beta \Leftrightarrow x \geq \frac{|r|}{\beta} .$$

For $x \leq \frac{r}{\beta}$, the function $f'(x)$ achieves its maximum at the point $\frac{r}{\beta}$. By Lemma 3.67, a β -smooth upper bound on f is

$$\text{UB}_f(x) = \begin{cases} x^r & \text{if } x \geq \frac{r}{\beta} \\ e^{\beta x - r} \left(\frac{r}{\beta}\right)^r & \text{otherwise} \end{cases}$$

If $r \geq 1$, we may also find a smooth upper bound on the derivative sensitivity DS_f of f . We have

$$\text{DS}_f(x) = |f'(x)| = |r|x^{r-1} .$$

A β -smooth upper bound on DS_f is

$$\text{UB}_{\text{DS}_f}(x) = \begin{cases} rx^{r-1} & \text{if } x \geq \frac{r-1}{\beta} \\ re^{\beta x - (r-1)} \left(\frac{r-1}{\beta}\right)^{r-1} & \text{otherwise} \end{cases}$$

Exponent. Let $f(x) = e^{rx}$, $r \in \mathbb{R}$, $x \in \mathbb{R}$. We have $\text{DS}_f(x) = |f'(x)| = |r|e^{rx}$, hence:

Approved for Public Release; Distribution Unlimited.

- $\left| \frac{f'(x)}{f(x)} \right| = \frac{re^{rx}}{e^{rx}} = r$; $\left| \frac{f''(x)}{f(x)} \right| = \frac{r^2 e^{rx}}{e^{rx}} = r$.

Thus both f and DS_f are β -smooth if $|r| \leq \beta$.

Sigmoid. Consider the (sigmoid) function $\sigma(x) = \frac{e^{\alpha x}}{e^{\alpha x} + 1}$. This function can be viewed as a continuous approximation of the indicator function $I_{\mathbb{R}_+} : \mathbb{R} \rightarrow \{0, 1\}$, which is less precise for values close to 0, and the error decreases when α increases. We have:

- $\sigma'(x) = \frac{\alpha e^{\alpha x}}{(e^{\alpha x} + 1)^2}$; $\sigma''(x) = \frac{\alpha^2 e^{\alpha x}(e^{\alpha x} - 1)}{(e^{\alpha x} + 1)^3}$;
- $\left| \frac{\sigma'(x)}{\sigma(x)} \right| = \left| \alpha \cdot \frac{1}{e^{\alpha x} + 1} \right| \leq \alpha$; $\left| \frac{\sigma''(x)}{\sigma'(x)} \right| = \left| \alpha \cdot \frac{e^{\alpha x} - 1}{e^{\alpha x} + 1} \right| \leq \alpha$.

Thus both $\sigma(x)$ and $DS_{\sigma(x)} = |\sigma'(x)|$ are α -smooth. If we want less DP noise, we should decrease α , which in turn makes the sigmoid itself less precise, so there is a tradeoff.

Tauoid. Consider the function $\tau(x) = \frac{2}{e^{-\alpha x} + e^{\alpha x}}$ (let us call it a *tauoid*). This function can be viewed as a continuous approximation of the indicator function $I_{\{0\}} : \mathbb{R} \rightarrow \{0, 1\}$, which works similarly to a sigmoid. We have:

$$\begin{aligned} \tau'(x) &= -\frac{2\alpha(e^{\alpha x} - e^{-\alpha x})}{(e^{-\alpha x} + e^{\alpha x})^2} \\ &= \frac{2\alpha(e^{-\alpha x} - e^{\alpha x})}{(e^{-\alpha x} + e^{\alpha x})^2} = \frac{2\alpha e^{\alpha x}(1 - e^{2\alpha x})}{(1 + e^{2\alpha x})^2} \\ \left| \frac{\tau'(x)}{\tau(x)} \right| &= \frac{|\alpha| \cdot |e^{-\alpha x} - e^{\alpha x}|}{e^{-\alpha x} + e^{\alpha x}} \leq |\alpha| \\ |\tau'(x)| &\leq \frac{2|\alpha|e^{\alpha x}}{1 + e^{2\alpha x}} = \frac{2|\alpha|}{e^{-\alpha x} + e^{\alpha x}} \\ &= |\alpha|\tau(x) =: \text{UB}_{DS_\tau}(x) \\ \text{UB}'_{DS_\tau}(x) &= |\alpha|\tau'(x) \\ \left| \frac{\text{UB}'_{DS_\tau}(x)}{\text{UB}_{DS_\tau}(x)} \right| &= \left| \frac{\tau'(x)}{\tau(x)} \right| \leq |\alpha|. \end{aligned}$$

Thus both τ itself and UB_{DS_τ} , an upper bound on its derivative sensitivity, are α -smooth.

An ℓ_p -norm. Consider the function $f(x) = \|x\|_p = \left(\sum x_i^p \right)^{1/p}$, $x \in \mathbb{R}^n$, $x = (x_1, \dots, x_n)$. We have

$$\frac{\partial f}{\partial x_i}(x) = \frac{px_i^{p-1}}{p \left(\sum x_i^p \right)^{(p-1)/p}} = \left(\frac{x_i^p}{\sum x_i^p} \right)^{(p-1)/p}.$$

By Lemma 3.64, the derivative sensitivity of f in (\mathbb{R}^n, ℓ_p) is

$$DS_f(x) = \left(\frac{\sum x_i^p}{\sum x_i^p} \right)^{\frac{p-1}{p}} = 1.$$

This is constant and thus β -smooth for all β . The function f itself is β -smooth if $\frac{1}{\|x\|_p} \leq \beta$, i.e. if $\|x\|_p \geq \frac{1}{\beta}$. By Lemma 3.67, a β -smooth upper bound on f is

$$\text{UB}_f(x) = \begin{cases} \|x\|_p & \text{if } \|x\|_p \geq \frac{1}{\beta} \\ \frac{e^{\beta\|x\|_p - 1}}{\beta} & \text{otherwise} \end{cases}$$

This also holds for $p = \infty$.

The ℓ_∞ -norm. Let $f(x) = \|x\|_\infty = \max_i |x_i|$. We have

$$\frac{\partial f}{\partial x_i}(x) = \begin{cases} 1 & \text{if } i = \text{argmax}_j |x_j| \\ \text{undefined} & \text{if } \text{argmax}_j |x_j| \text{ is not unique} \\ 0 & \text{otherwise} \end{cases}$$

The derivative sensitivity of f in $(\mathbb{R}^n, \ell_\infty)$ is

$$\text{DS}_f(x) = \begin{cases} 1 & \text{if } \operatorname{argmax}_j |x_j| \text{ is unique} \\ \text{undefined} & \text{if } \operatorname{argmax}_j |x_j| \text{ is not unique} \end{cases}$$

Because we are interested in upper bounds on the derivative sensitivity, we define

$$\text{DS}_f(x_0) := \limsup_{x \rightarrow x_0} \text{DS}_f(x) = 1$$

for those x_0 for which $\text{DS}_f(x_0)$ is undefined. Thus $\text{DS}_f(x) = 1$, which is constant and β -smooth for all β . The smooth upper bound on the function f itself can be found similarly to the ℓ_p -norm case.

The following constructions are considered in Fig. 41.

Product. Let $f : \prod_{i=1}^n X_i \rightarrow \mathbb{R}, f(x_1, \dots, x_n) = \prod_{i=1}^n f_i(x_i)$ where X_i are Banach spaces. Let $X = \prod_{i=1}^n X_i$ and $x = (x_1, \dots, x_n)$. First, suppose that variables x_i are independent. We have $\frac{\partial f}{\partial x_i}(x) = \prod_{i \neq j=1}^n f_j(x_j) \cdot f'_i(x_i)$, and $\left| \frac{\partial f}{\partial x_i}(x) \cdot \frac{1}{f(x)} \right| = \left| \frac{f'_i(x_i)}{f_i(x_i)} \right|$, hence:

- If $\left| \frac{f'_i(x_i)}{f_i(x_i)} \right| \leq \beta$, then f is β -smooth w.r.t. x_i .
- By Lemmas 3.64 and 3.66,

$$\|df_x\| = \left\| \left(\prod_{i \neq j=1}^n f_j(x_j) \cdot f'_i(x_i) \right)_{i=1}^n \right\|_{\frac{p}{p-1}}$$

in (X, ℓ_p) , so we have $\frac{\|df_x\|}{|f(x)|} = \frac{\|df_x\|}{\left| \prod_{i=1}^n f_i(x_i) \right|} = \left\| \left(\frac{f'_i(x_i)}{f_i(x_i)} \right)_{i=1}^n \right\|_{\frac{p}{p-1}} \leq \|(\beta_i)_{i=1}^n\|_{\frac{p}{p-1}}$, where β_i is the smoothness of f_i . Hence, if f_i is β -smooth w.r.t. x_i for all i , then f is β -smooth in (X, ℓ_1) and $n\beta$ -smooth in (X, ℓ_∞) .

The derivative sensitivity of f w.r.t. x_i is $c_i(x) = \text{DS}_{f_i}(x_i) \cdot \left| \frac{f(x)}{f_i(x_i)} \right|$. The derivative sensitivity of f in (X, ℓ_p) is, by Lemma 3.66, $\text{DS}_f(x) = \|(c_1(x), \dots, c_n(x))\|_{\frac{p}{p-1}} = \left\| \left(\frac{\text{DS}_{f_i}(x_i)}{|f_i(x_i)|} \right)_{i=1}^n \right\|_{\frac{p}{p-1}} \cdot |f(x)|$.

We have $c_i(x) = \text{DS}_{f_i}(x_i) \cdot \left| \frac{f(x)}{f_i(x_i)} \right| = \text{DS}_{f_i}(x_i) \cdot \prod_{j \neq i} |f_j(x_j)|$. Since $\prod_{j \neq i} |f_j(x_j)|$ does not depend on x_i and $\text{DS}_{f_i}(x_i) \geq 0$, by Lemma 3.70, if DS_{f_i} is β -smooth in X_i then $c_i(x)$ is also β -smooth in X_i . Similarly, if $f_j(x_j)$ is β -smooth, then $c_i(x)$ is also β -smooth in X_j . Hence, if f_i and DS_{f_i} are β -smooth for all i , by Lemma 3.72, DS_f is β -smooth in (X, ℓ_1) and $n\beta$ -smooth in (X, ℓ_∞) . If DS_{f_i} are not all β -smooth then we can use their β -smooth upper bounds when computing c_i . Then we get a β -smooth upper bound on DS_f instead of the actual DS_f .

We may also consider the case where the variables x_i are fully dependent, i.e. equal (the case where they are partially dependent is currently not considered). Consider a function $f(x) = g(x) \cdot h(x)$ where $g, h : X \rightarrow \mathbb{R}_+$ and X is a Banach space. We have

$$\text{DS}_f(x) = g(x) \cdot \text{DS}_h(x) + h(x) \cdot \text{DS}_g(x) .$$

By Lemma 3.70, if g is β_g -smooth, h is β_h -smooth, DS_g is $\beta_{g'}$ -smooth, and DS_h is $\beta_{h'}$ -smooth, then DS_f is $\max(\beta_g + \beta_{h'}, \beta_h + \beta_{g'})$ -smooth. The function f itself is $(\beta_g + \beta_h)$ -smooth.

Sum. Let $f : \prod_{i=1}^n X_i \rightarrow \mathbb{R}, f(x_1, \dots, x_n) = \sum_{i=1}^n f_i(x_i)$ where X_i are Banach spaces. Let $X = \prod_{i=1}^n X_i$ and $x = (x_1, \dots, x_n)$. First, suppose that the variables x_i are independent. The derivative sensitivity of f w.r.t. x_i is $\text{DS}_{f_i}(x_i)$. By Lemmas 3.64 and 3.66, the derivative sensitivity of f in (X, ℓ_p) is $\text{DS}_f(x) = \|\text{DS}_{f_1}(x_1), \dots, \text{DS}_{f_n}(x_n)\|_{\frac{p}{p-1}}$.

- Let $f_i \geq 0$ for all $i \in \{1, \dots, n\}$ (or $f_i \leq 0$ for all $i \in \{1, \dots, n\}$) and β_i -smooth w.r.t. X_i . Now we have $|f(x)| = \sum_{i=1}^n |f_i(x_i)| = \|\sum_{i=1}^n |f_i(x_i)|\|_1$. By Lemma 3.71, $f(x)$ is $\beta := \max_i(\beta_i)$ -sensitive in (X, ℓ_p) . We do not get a good bound in the case when f_i may have different signs, since then $f_i(x)$ may cancel each other out and make $f(x)$ arbitrarily small even if $|f_i(x)|$ are large.

- Let DS_{f_i} be β_i -smooth for $i \in \{1, \dots, n\}$. By Lemma 3.71, DS_f is $\|(\beta_i)_{i=1}^n\|_{\frac{p}{p-1}}$ -smooth in (X, ℓ_p) , and if all DS_{f_i} are β -smooth, then DS_f is also β -smooth.

Consider the case where x_i are equal: $f(x) = \sum_{i=1}^n g_i(x)$ where $g_i : X \rightarrow \mathbb{R}$ and X is a Banach space. Then

$$DS_f(x) = \sum_{i=1}^n DS_{g_i}(x)$$

By Lemma 3.70, if all DS_{g_i} are β -smooth then DS_f is β -smooth. If all g_i are non-negative and β -smooth then f is β -smooth.

Min / max. Let $f : \prod_{i=1}^n X_i \rightarrow \mathbb{R}$, $f(x_1, \dots, x_n) = \min_{i=1}^n f_i(x_i)$ where X_i are Banach spaces (the case with max instead of min is similar). Let $X = \prod_{i=1}^n X_i$ and $x = (x_1, \dots, x_n)$. Let the variables x_i be independent.

If for all i , f_i is β -smooth in X_i then f is β -smooth in (X, ℓ_p) . The same holds with max or sum (with non-negative f_i) or $\ell_{p'}$ -norm instead of min.

The derivative sensitivity of f w.r.t. x_i is $DS_{f_i}(x_i)$ if $i = \operatorname{argmin} f_i(x_i)$ and 0 otherwise. The derivative sensitivity of f in (X, ℓ_p) is $DS_f(x) = DS_{f_i}(x_i)$ where $i = \operatorname{argmin} f_i(x_i)$. In general, DS_f is discontinuous at points where $\operatorname{argmin} f_i(x_i)$ is not unique.

A possible valid β -smooth (in (X, ℓ_p)) upper bound on DS_f is $\max c_i(x_i)$ where c_i is a β -smooth upper bound on DS_{f_i} .

Norm scaling. Let $f : X \rightarrow \mathbb{R}$ in the Banach space $(X, \|\cdot\|)$. Scaling the norm by a scales the derivative $f'(x)$ by $\frac{1}{a}$ while keeping the value of $f(x)$ the same. Hence, if f is β -smooth in $(X, \|\cdot\|)$ then it is $\frac{\beta}{a}$ -smooth in $(X, a \cdot \|\cdot\|)$.

Let $c(x)$ be a β -smooth upper bound on the derivative sensitivity of f at x in $(X, \|\cdot\|)$. Then $\frac{c(x)}{a}$ is a $\frac{\beta}{a}$ -smooth upper bound on the derivative sensitivity of f at x in $(X, a \cdot \|\cdot\|)$ by Lemma 3.65.

Sensitivity w.r.t. a larger norm. Let $f : X \rightarrow \mathbb{R}$ in the Banach space $(X, \|\cdot\|_N)$. Let $\|\cdot\|_M \geq \|\cdot\|_N$.

If f is β -smooth in $(X, \|\cdot\|_N)$, then $f(x) \leq e^{\beta\|x-x'\|_N} \cdot f(x') \leq e^{\beta\|x-x'\|_M} \cdot f(x')$ for all $x, x' \in X$, so f is also β -smooth in $(X, \|\cdot\|_M)$. The same holds about any function that is β -smooth in $(X, \|\cdot\|_N)$, including a β -smooth upper bound on the derivative sensitivity of f .

Let us show that the derivative sensitivity of f w.r.t. $\|\cdot\|_N$ is a valid upper bound on the derivative sensitivity of f w.r.t. $\|\cdot\|_M$. First, note that $\|\cdot\|_{\operatorname{dual}(N)} \geq \|\cdot\|_{\operatorname{dual}(M)}$. Indeed, by definition of a dual norm, $\|T\|_{\operatorname{dual}(M)} = \sup\{T(x) \mid \|x\|_M \leq 1\}$ for an operator T from the dual space $X \rightarrow \mathbb{R}$ of X . Since $\|x\|_N \leq \|x\|_M$, we have $\forall x : \{T(x) \mid \|x\|_N \leq 1\} \supseteq \{T(x) \mid \|x\|_M \leq 1\}$. Hence, $\|T\|_{\operatorname{dual}(N)} = \sup\{T(x) \mid \|x\|_N \leq 1\} \geq \sup\{T(x) \mid \|x\|_M \leq 1\} = \|T\|_{\operatorname{dual}(M)}$.

By definition, we have $DS_f(x) = \|df_x\|_{\operatorname{dual}(N)}$, where df_x is the Fréchet derivative of f at x . Since $\|\cdot\|_{\operatorname{dual}(N)} \geq \|\cdot\|_{\operatorname{dual}(M)}$, we have $\|df_x\|_{\operatorname{dual}(N)} \geq \|df_x\|_{\operatorname{dual}(M)}$.

Composition with a real function. Let $f(x) = h(g(x))$, $x \in X$ where $g : X \rightarrow \mathbb{R}$, $h : \mathbb{R} \rightarrow \mathbb{R}$ and X is a Banach space.

$$DS_f(x) = |h'(g(x))| \cdot DS_g(x)$$

$$\frac{DS_f(x)}{|f(x)|} = \frac{|h'(g(x))|}{|h(g(x))|} \cdot DS_g(x)$$

Suppose that h is β_h -smooth and $DS_g(x) \leq B$ for all x . Then f is $\beta_h B$ -smooth. We have

$$DS_{DS_f}(x) = |h''(g(x))|(DS_g(x))^2 + |h'(g(x))| \cdot DS_{DS_g}(x) ,$$

$$\frac{DS_{DS_f}(x)}{DS_f(x)} = \frac{|h''(g(x))|}{|h'(g(x))|} \cdot DS_g(x) + \frac{DS_{DS_g}(x)}{DS_g(x)} .$$

By Lemma 3.70, if h' is $\beta_{h'}$ -smooth, DS_g is $\beta_{g'}$ -smooth, and $DS_g(x) \leq B$ for all x then DS_f is $(\beta_{h'} B + \beta_{g'})$ -smooth.

Example 3.9. Consider an example of computing differentially privately the time it takes for the next ship to reach the port. This time can be expressed as

$$f(x_1, y_1, v_1, \dots, x_n, y_n, v_n) = \min_{i=1}^n \frac{\sqrt{x_i^2 + y_i^2}}{v_i}$$

$$f : \mathbb{R}^{3n} \rightarrow \mathbb{R}$$

where (x_i, y_i) are the coordinates of the i^{th} ship (with the port at $(0, 0)$) and v_i is its speed.

Note that v_i is in the power -1 and we do not know how to find the smooth derivative sensitivity of the function $f_{v,i}(v_i) = v_i^{-1}$ (we only know how to do it for power functions with exponent ≥ 1). Let us define $w_i = \zeta \ln v_i$. The coefficient ζ is used to control the distance by which the whole input vector changes if $\ln v_i$ is changed by 1. Similarly, we add a coefficient to the geographical coordinates: $s_i = \alpha x_i, t_i = \alpha y_i$. Then we consider $(s_1, t_1, w_1, \dots, s_n, t_n, w_n)$ as an element of the Banach space $(\mathbb{R}^{3n}, \|\cdot\|)$ where

$$\|(s_1, t_1, w_1, \dots, s_n, t_n, w_n)\| = \|(\|(\|s_1, t_1\|_2, w_1)\|_1, \dots, \|(\|s_n, t_n\|_2, w_n)\|_1)\|_p$$

Then $v_i = e^{w_i/\zeta}, x_i = \frac{s_i}{\alpha}, y_i = \frac{t_i}{\alpha}$ and

$$g(s_1, t_1, w_1, \dots, s_n, t_n, w_n) = \frac{1}{\alpha} \min_{i=1}^n \frac{\sqrt{s_i^2 + t_i^2}}{e^{w_i/\zeta}}$$

Now the derivative sensitivity of $g_{w,i}(w_i) = e^{-w_i/\zeta}$ is

$$\text{DS}_{g_{w,i}}(w_i) = \frac{1}{\zeta} e^{-w_i/\zeta}$$

which is $\frac{1}{\zeta}$ -smooth. The function $g_{w,i}$ itself is also $\frac{1}{\zeta}$ -smooth.

The derivative sensitivity of $g_{st,i}(s_i, t_i) = \sqrt{s_i^2 + t_i^2}$ in (\mathbb{R}^2, ℓ_2) is

$$\text{DS}_{g_{st,i}}(s_i, t_i) = 1$$

which is β -smooth for all β . The function $g_{st,i}$ is $\frac{1}{\zeta}$ -smooth if $\frac{1}{\sqrt{s_i^2 + t_i^2}} \leq \frac{1}{\zeta}$, i.e. if $\sqrt{s_i^2 + t_i^2} \geq \zeta$. A $\frac{1}{\zeta}$ -smooth upper bound on $g_{st,i}$ is

$$\hat{g}_{st,i}(s_i, t_i) = \begin{cases} \sqrt{s_i^2 + t_i^2} & \text{if } \sqrt{s_i^2 + t_i^2} \geq \zeta \\ \zeta e^{\frac{\sqrt{s_i^2 + t_i^2}}{\zeta} - 1} & \text{otherwise} \end{cases}$$

An upper bound on the derivative sensitivity of $g_i(s_i, t_i, w_i) = \frac{\sqrt{s_i^2 + t_i^2}}{e^{w_i/\zeta}}$ is

$$\begin{aligned} c_{g_i}(s_i, t_i, w_i) &= \left\| (\text{DS}_{g_{st,i}}(s_i, t_i) \cdot g_{w,i}(w_i), \text{DS}_{g_{w,i}}(w_i) \cdot \hat{g}_{st,i}(s_i, t_i)) \right\|_{\infty} = \\ &= \left\| \left(1 \cdot e^{-w_i/\zeta}, \frac{1}{\zeta} e^{-w_i/\zeta} \cdot \hat{g}_{st,i}(s_i, t_i) \right) \right\|_{\infty} = \frac{\max\left(1, \frac{\hat{g}_{st,i}(s_i, t_i)}{\zeta}\right)}{e^{w_i/\zeta}} \end{aligned}$$

and it is $\frac{1}{\zeta}$ -smooth because $\text{DS}_{g_{w,i}}(w_i)$, $\text{DS}_{g_{st,i}}(s_i, t_i)$, and $\hat{g}_{st,i}(s_i, t_i)$ are $\frac{1}{\zeta}$ -smooth.

A $\frac{1}{\zeta}$ -smooth upper bound on DS_{g_i} is

$$c(u) = \frac{1}{\alpha} \max_i c_{g_i}(s_i, t_i, w_i) = \frac{1}{\alpha} \max_i \frac{\max\left(1, \frac{\hat{g}_{st,i}(s_i, t_i)}{\zeta}\right)}{e^{w_i/\zeta}}$$

Approved for Public Release; Distribution Unlimited.

where $u = (s_1, t_1, w_1, \dots, s_n, t_n, w_n)$.

Now we can use Theorem 3.62 to compute an ϵ -differentially private version of g :

$$h(u) = g(u) + \frac{c(u)}{b} \cdot \eta$$

$$\epsilon = (\gamma + 1)(b + \frac{1}{\zeta})$$

$$\gamma > 1, b > 0, \eta \sim \text{GenCauchy}(\gamma)$$

To compute an ϵ -differentially private version of f , we first transform $(x_1, y_1, v_1, \dots, x_n, y_n, v_n)$ into u and then compute $h(u)$.

Smooth derivative sensitivity w.r.t. multiple tables. To achieve differential privacy, we need a β -smooth upper bound on derivative sensitivity. For this, we first compute an upper bound on the sensitivity w.r.t. a single table that is β -smooth w.r.t. all the necessary tables. We compute this for each necessary table and then take the maximum. Maximum preserves β -smoothness and because we are using ℓ_1 -norm, we get an upper bound on the sensitivity w.r.t. all needed tables. It remains to show how to ensure that the upper bound on derivative sensitivity is β -smooth w.r.t. all the necessary tables (and not only the table w.r.t. which the sensitivity was computed).

To compute an upper bound on the sensitivity w.r.t. a single table that is β -smooth w.r.t. all the tables, we force the sensitivity w.r.t. the other tables to zero. If the other tables are not sensitive (i.e. do not contain any attributes included in the database norm), we use the exact values of subexpressions instead of their smooth upper bounds. For the other tables that are sensitive (i.e. contain at least one attribute included in the database norm), we still compute β -smooth upper bounds on the values of subexpressions (but their derivative sensitivities are 0).

If a table is used more than once then we compute the derivative sensitivity w.r.t. each copy of the table separately and add the result together because changing a row in the table is equivalent to changing a row in each copy of the table. A problem occurs with smoothing. If the sensitivity w.r.t. one copy is β -smooth w.r.t. each copy separately then changing a row in one copy by distance d changes the sensitivity by up to $e^{\beta d}$ times. Changing it in all m copies changes the sensitivity by up to $e^{m\beta d}$ times. Thus the sensitivity w.r.t. the initial table (instead of one of its copies) is $m\beta$ -smooth instead of β -smooth.

To achieve β -smoothness w.r.t. each initial table, we require the sensitivities to be (for all i) $\frac{\beta}{m_i}$ -smooth w.r.t. each copy of table i where m_i is the number of times the table i is used. Our analysis does not directly support different values of β for different (copies of) tables, thus we need to scale the β at appropriate points during the analysis. Scaling at the end is not optimal because then we would achieve $\frac{\beta}{m}$ -smoothness w.r.t. each copy of each table, where $m = \max_i m_i$, even for those tables that are used less than m times. Thus we scale at the beginning (at leaf nodes) but this requires also using a scaled norm for global sensitivity because global sensitivity is used for computing smooth sensitivity at some intermediate nodes (compose versions of exponent e^x , sigmoid $\frac{e^{ax}}{e^{ax}+1}$, and tauoid $\frac{2}{e^{ax}+e^{-ax}}$).

In the analyses where we also need a non-scaled global sensitivity (e.g. used for smoothing the combined sensitivity described in Sec. 3.3.9), we just compute the global sensitivity twice w.r.t. two different norms.

Optimizing smoothness parameter. Theorem 3.62 tells how to use derivative sensitivity to achieve differential privacy in a Banach space. We prove a stronger version of this theorem that allows us to choose the best value of β (the one giving the smallest noise level) according to the actual input, instead of having to guess it in advance.

Theorem 3.73. *Let $\gamma, b, \beta \in \mathbb{R}_+$, $\gamma > 1$. Let $f : X \rightarrow \mathbb{R}$ where X is a Banach space. Let \mathcal{I} be a (possibly infinite) set of indices. For all $I \in \mathcal{I}$, let $\epsilon = (\gamma + 1)(b_I + \beta_I)$ and c_I be a β_I -smooth upper bound on DS_f . Let η be a random variable distributed according to $\text{GenCauchy}(\gamma)$. Then $g(\vec{x}) = f(\vec{x}) + h(\vec{x}) \cdot \eta$, where $h(\vec{x}) = \min_{I \in \mathcal{I}} \frac{c_I(\vec{x})}{b_I}$ (provided that the minimum exists), is ϵ -differentially private.*

Proof. Let $\vec{x}, \vec{x}' \in X$. Denote $L = \|\vec{x}' - \vec{x}\|$. We have to show that $d_{\text{dp}}(g(\vec{x}), g(\vec{x}')) \leq \epsilon L$.

Let $n \in \mathbb{N}$ be arbitrary. Let $\vec{v}_0 = \vec{x}$, $\vec{v}_n = \vec{x}'$, and $\vec{v}_i = \frac{n-i}{n} \cdot \vec{x} + \frac{i}{n} \cdot \vec{x}'$. Consider $i \in \{1, \dots, n\}$. There exist $I, J \in \mathcal{I}$ such that $h(\vec{v}_{i-1}) = \frac{c_I(\vec{v}_{i-1})}{b_I}$ and $h(\vec{v}_i) = \frac{c_J(\vec{v}_i)}{b_J}$. Note that $\frac{c_I(\vec{v}_{i-1})}{b_I} \leq \frac{c_I(\vec{v}_{i-1})}{b_J}$ and $\frac{c_I(\vec{v}_i)}{b_I} \geq \frac{c_J(\vec{v}_i)}{b_J}$. Because c_I and c_J are continuous, there exists a point \vec{u} on the (closed) segment connecting \vec{v}_{i-1} to \vec{v}_i where $\frac{c_I(\vec{u})}{b_I} = \frac{c_J(\vec{u})}{b_J}$. Let $K, M \in \{I, J\}$ be such that $\beta_K = \max(\beta_I, \beta_J)$, $\beta_M = \min(\beta_I, \beta_J)$, and $\{K, M\} = \{I, J\}$. Then

$$\left| \ln \frac{h(\vec{v}_i)}{h(\vec{v}_{i-1})} \right| \leq \left| \ln \frac{h(\vec{v}_i)}{h(\vec{u})} \right| + \left| \ln \frac{h(\vec{u})}{h(\vec{v}_{i-1})} \right| \leq \beta_J \|\vec{v}_i - \vec{u}\| + \beta_I \|\vec{u} - \vec{v}_{i-1}\| \leq \beta_K L/n .$$

The mean value theorem for Banach spaces and Fréchet derivative gives

$$|f(\vec{v}_i) - f(\vec{v}_{i-1})| \leq \|df_{\vec{t}_i}\| \cdot \|\vec{v}_i - \vec{v}_{i-1}\| \leq c_K(\vec{t}_i) \cdot \frac{L}{n} \leq e^{\beta_K L/n} \cdot c_K(\vec{v}_{i-1}) \cdot \frac{L}{n}$$

for some point \vec{t}_i on the segment connecting \vec{v}_{i-1} to \vec{v}_i . Here we also used the fact that c_K is an upper bound on DS_f .

If $h(\vec{v}_{i-1}) = \frac{c_K(\vec{v}_{i-1})}{b_K}$ then $c_K(\vec{v}_{i-1}) = b_K \cdot h(\vec{v}_{i-1})$. Otherwise

$$h(\vec{v}_{i-1}) = \frac{c_M(\vec{v}_{i-1})}{b_M} \geq \frac{e^{-\beta_M \|\vec{u} - \vec{v}_{i-1}\|} c_M(\vec{u})}{b_M} = \frac{e^{-\beta_M \|\vec{u} - \vec{v}_{i-1}\|} c_K(\vec{u})}{b_K} \geq \frac{e^{-\beta_M \|\vec{u} - \vec{v}_{i-1}\|} e^{-\beta_K \|\vec{v}_{i-1} - \vec{u}\|} c_K(\vec{v}_{i-1})}{b_K} = \frac{e^{-(\beta_I + \beta_J)L/n} c_K(\vec{v}_{i-1})}{b_K} .$$

In both cases, $c_K(\vec{v}_{i-1}) \leq e^{(\beta_I + \beta_J)L/n} \cdot b_K \cdot h(\vec{v}_{i-1})$. Thus

$$|f(\vec{v}_i) - f(\vec{v}_{i-1})| \leq e^{(\beta_I + \beta_J + \beta_K)L/n} \cdot b_K \cdot h(\vec{v}_{i-1}) \cdot \frac{L}{n}$$

and

$$d_{\text{dp}}(g(\vec{v}_{i-1}), g(\vec{v}_i)) = d_{\text{dp}}(f(\vec{v}_{i-1}) + h(\vec{v}_{i-1}) \cdot \eta, f(\vec{v}_i) + h(\vec{v}_i) \cdot \eta) \leq (\gamma + 1) \left(\frac{|f(\vec{v}_i) - f(\vec{v}_{i-1})|}{|h(\vec{v}_{i-1})|} + \frac{\beta_K L}{n} \right) \leq (\gamma + 1) \left(b_K \cdot e^{(\beta_I + \beta_J + \beta_K)L/n} \cdot \frac{L}{n} + \frac{\beta_K L}{n} \right) = (\gamma + 1)(b_K e^{(\beta_I + \beta_J + \beta_K)L/n} + \beta_K)L/n \leq (\gamma + 1)(b_K e^{3\beta L/n} + \beta_K)L/n$$

where $\beta = \sup_{I \in \mathcal{I}} \beta_I$. This supremum exists because for all $I \in \mathcal{I}$, $\beta_I = \frac{\epsilon}{\gamma+1} - b_I \leq \frac{\epsilon}{\gamma+1}$. Now

$$\begin{aligned} d_{\text{dp}}(g(\vec{x}), g(\vec{x}')) &\leq \sum_{i=1}^n d_{\text{dp}}(g(\vec{v}_{i-1}), g(\vec{v}_i)) \leq L/n \cdot (\gamma + 1) \sum_{i=1}^n \max_{K \in \mathcal{I}} b_K e^{3\beta L/n} + \beta_K \leq \\ &L/n \cdot (\gamma + 1) \sum_{i=1}^n \max_{K \in \mathcal{I}} e^{3\beta L/n} (b_K + \beta_K) = L/n \cdot e^{3\beta L/n} \sum_{i=1}^n \max_{K \in \mathcal{I}} (\gamma + 1)(b_K + \beta_K) = \\ &L/n \cdot e^{3\beta L/n} \sum_{i=1}^n \max_{K \in \mathcal{I}} \epsilon = e^{3\beta L/n} \epsilon L . \end{aligned}$$

This inequality holds for any $n \in \mathbb{N}$. If $n \rightarrow \infty$ then $e^{3\beta L/n} \rightarrow 1$ and we obtain the inequality that we had to show. \square

3.3.8.4 Derivative Sensitivity of SQL Queries. We now describe how the theory of derivative sensitivity for Banach spaces can be applied to SQL queries, computing the smooth upper bounds of their derivative sensitivities, such that an appropriate amount of noise may be added to turn them differentially private. Our theory deals with functions that return a numeric value, so the query should return a single output. We do not support DISTINCT queries, as we do not know how to continuously approximate

Continuous function	Approximated condition	name
$\frac{e^{\alpha x}}{e^{\alpha x} + 1}$	$x \geq 0$	sigmoid
$\frac{2}{e^{-\alpha x} + e^{\alpha x}}$	$x = 0$	tauoid

Table 9: Continuous Approximations of Step Functions

efficiently a function that removes repeating elements from a list of arguments. We consider queries of the form

```
SELECT aggr expr FROM t1 AS s1, ..., tn AS sn
WHERE condition ,
```

where:

- *expr* is an expression over table columns, computed as a continuous function.
- *condition* is a boolean expression over predicates $P(x) \in \{x < 0, x = 0\}$, where x is an expression of the same form as *expr*. Since all functions have to be continuous, these predicates are computed using continuous approximations to the step functions, listed in Table. 9.
- *aggr* is one of the operations SUM, COUNT, MIN, MAX.

GROUP BY queries can be simulated by generating for each group a separate query, with a filter selecting that particular group. Hence, we can group either by a public or a discrete attribute to get a finite number of groups.

Continuous approximation of an SQL query. We assume that the expression under SELECT statement is an application of a continuous function to the database. Hence, for a query without a filter, we can directly apply results of Section 3.3.8.3 to find a smooth upper bound on derivative sensitivity,

A filter that *does not depend* on sensitive data can be applied directly to the cross product of the input tables, and we may then proceed with the query without a filter.

A filter that *does depend* on sensitive data is treated as a part of the query. We treat this filter as a continuous function, applied in such a way that the discarded rows would be ignored by the aggregating function. We combine sigmoids and tauoids to obtain the approximated value of the indicator $\sigma(x_i) \in \{0, 1\}$, denoting whether the row x_i satisfies the filter.

Hence, if the filter depends on sensitive data, we have the following set-up:

- There is a set of rows $\{x_1, \dots, x_m\}$.
- There is a function f_i applied to the row x_i , returning a real number. For different rows, this function may be different, e.g. it may be determined by the public cells of the row.
- There is a filtering function σ_i applied to the row x_i . It returns a real number. It approximates a boolean condition, i.e. its values are mostly near 0 and 1.
- There is an aggregation function applied to a subset of the values $f_1(x_1), \dots, f_m(x_m)$. Only such $i \in \{1, \dots, m\}$ are selected, where the condition holds.

To convert the SQL query into a continuous function, the functions f_i and σ_i are combined as follows, depending on the aggregation function:

- SUM. The values 0 do not affect the sum, hence we compute the result as $\sum_{i=1}^m f_i(x_i) \cdot \sigma_i(x_i)$.
- COUNT: The values of f_i do not affect the result. We compute the result as $\sum_{i=1}^m \sigma_i(x_i)$, counting all entries for which $\sigma_i(x_i) = 1$. The sensitivity of such query only depends on the sensitivity of σ .

Table 10: Upper Bounds for Uni- and Multivariate functions

$f(x)$	cond.s	g , s.t. $f \leq g \sim_{ \cdot } \beta$	h , s.t. $DS_f \leq'_{ \cdot } h \sim \beta$
x^r	$r \geq 1$ $x > 0$	$\begin{cases} x^r & \text{if } x \geq \frac{r}{\beta} \\ \text{pw}_{\beta}^r(x) & \text{oth.} \end{cases}$	$\begin{cases} rx^{r-1} & \text{if } x \geq \frac{r-1}{\beta} \\ r \cdot \text{pw}_{\beta}^{r-1}(x) & \text{oth.} \end{cases}$
$f(x)$	cond.s	g , s.t. $f \leq g \sim_{ \cdot } \beta$	h , s.t. $DS_f \leq'_{ \cdot } h \sim \beta$
e^{rx}	$ r \leq \beta$	e^{rx}	$ r e^{rx}$
c	$c \in \mathbb{R}$	c	0
$\frac{e^{\alpha x}}{e^{\alpha x} + 1}$	$\beta \geq \alpha$	$\frac{e^{\alpha x}}{e^{\alpha x} + 1}$	$\frac{\alpha e^{\alpha x}}{(e^{\alpha x} + 1)^2}$
$\frac{e^{\alpha x}}{e^{\alpha x} + 1}$	$\beta < \alpha$	1	$\frac{\alpha e^{\beta x}}{(e^{\beta x} + 1)^2}$
$\frac{2e^{\alpha x}}{1 + e^{2\alpha x}}$	$\beta \geq \alpha$	$\frac{2e^{\alpha x}}{1 + e^{2\alpha x}}$	$\frac{2 \alpha e^{\alpha x}}{1 + e^{2\alpha x}}$
$f(\vec{x})$	g , s.t. $f \leq g \sim_{\ell_p} \beta$		h , s.t. $DS_f \leq'_{\ell_p} h \sim \beta$
$\ \vec{x}\ _p$	$\begin{cases} \ \vec{x}\ _p & \text{if } \ \vec{x}\ _p \geq \frac{1}{\beta} \\ \text{pw}_{\beta}^1(\ \vec{x}\ _p) & \text{oth.} \end{cases}$		1

Here $\text{pw}_{\beta}^r(x) = \left(\frac{r}{\beta}\right)^r \cdot e^{\beta x - r}$

- MIN, MAX: If $\sigma_i(x_i)$ is 0, then we need to replace the actual value $f_i(x_i)$ with some large [resp. small] value that would not affect the result of MIN [resp. MAX]. Our conversion of the SQL query proceeds by first defining $\Delta := \text{MAX}(f(x_1), \dots, f(x_n)) - \text{MIN}(f(x_1), \dots, f(x_n))$, and then computing the result by applying MIN to the values $f_i(x_i) + (1 - \sigma_i(x_i)) \cdot \Delta$. MAX is computed similarly, changing the first “+” into a “-”.

If we know that the compared values are integers and hence $d(x, x') \geq 1$ for $x \neq x'$, we can do better than using sigmoids or tauoids from Table 9, defining precise functions:

- $x > y \iff \min(1, \max(0, x - y))$.
- $x = y \iff 1 - \min(1, \max(0, |x - y|))$.

An advantage of these functions is that they do not lose precision due to addition and multiplication.

For real numbers, we may bound the precision and assume e.g. that $d(x, x') \geq 1/k$ for some $k \geq 1$, which allows to use similar functions. The sensitivity of such comparisons will be k times larger than for integers.

Inferring derivative sensitivities. Let the write-up $DS_f \leq'_N h \sim \beta$ mean that h is a β -smooth upper bound on the derivative sensitivity of f , according to the norm $\|\cdot\|_N$ on the domain of f . For compositions, we also need the upper bounds for the (absolute values of) functions f themselves; let $f \sim_N \beta$ denote that f is β -smooth, and $f \leq g \sim_N \beta$ denote that g is a β -smooth upper bound of $|f|$, again according to the norm $\|\cdot\|_N$ on the domain of f . Table 10 lists the smooth upper bounds of some simple uni- and multivariate functions and their derivative sensitivities (using absolute value as the norm on \mathbb{R}). This table is a summary of Sec. 3.3.8.3 where these upper bounds are proved. For composite functions, the rules for computing the β -smooth upper bounds are given in Fig. 41.

Query norm vs user-defined norm. The facts and rules in Table 10 and Fig. 41 are in principle sufficient to compute the smooth upper bounds of derivative sensitivity of functions resulting from SQL queries with respect to *all* composite ℓ_p -norms. Still, when we naively apply them, we end up finding the sensitivity for a particular norm that is somehow “natural” for the function. In practice, it may happen that we actually need sensitivity w.r.t. some different norm, because the data owner specified so. For example, we know how to compute the sensitivity w.r.t. the norm $\|x_1, x_2\|_1$, but are interested in differential privacy w.r.t. $\|x_1, x_2\|_2$.

$$\begin{array}{l}
\frac{\forall i : \text{DS}_{f_i} \leq'_N h_i \sim \beta}{\text{DS}_{\sum_i c_i f_i} \leq'_N \sum_i c_i h_i \sim \beta} (+D) \\
\frac{\forall i : f_i \leq g_i \sim_N \beta}{\sum_i c_i f_i \leq \sum_i c_i g_i \sim_N \beta} (+S) \\
\frac{f_i \leq g_i \sim_N \beta}{f_1 \cdot f_2 \leq g_1 \cdot g_2 \sim_N \beta_1 + \beta_2} (*S) \\
\frac{\text{DS}_f \leq'_N g \sim \beta \quad N \leq M}{\text{DS}_f \leq'_M g \sim \beta} (\leq_D) \\
\frac{f \sim_N \beta}{f \sim_{a \cdot N} \beta/a} (N_S) \\
\frac{f_2 \sim_{|\cdot|} \beta \quad \forall \vec{x} : \text{DS}_{f_1}(\vec{x}) \leq B}{f_2 \circ f_1 \sim_N \beta B} (\circ_S) \\
\frac{\text{DS}_{f_i} \leq'_N h_i \sim \beta_i \quad \forall i, j : \text{vars}(N_i) \perp \text{vars}(N_j) \quad \forall i, j : f_i \cdot f_j \geq 0}{\text{DS}_{\sum_{i=1}^k f_i} \leq'_{\ell_p(N_1, \dots, N_k)} \|h_1, \dots, h_k\|_{\text{dual}(p)} \sim \|\beta_1, \dots, \beta_k\|_{\text{dual}(p)}} (+\frac{1}{D}) \\
\frac{f_i \leq g_i \sim_N \beta \quad \text{DS}_{f_i} \leq'_N h_i \sim \beta'_i}{\text{DS}_{f_1 \cdot f_2} \leq'_N g_1 \cdot h_2 + g_2 \cdot h_1 \sim \max(\beta_1 + \beta'_2, \beta'_1 + \beta_2)} (*D) \\
\frac{f_i \leq g_i \sim_{N_i} \beta \quad \text{DS}_{f_i} \leq'_{N_i} h_i \sim \beta \quad \text{vars}(N_1) \perp \text{vars}(N_2)}{\text{DS}_{f_1 \cdot f_2} \leq'_{N_1 + N_2} g_1 \cdot h_2 + g_2 \cdot h_1 \sim \beta} (*\frac{1}{D}) \\
\frac{f_i \leq g_i \sim_{N_i} \beta \quad \text{vars}(f_1) \perp \text{vars}(f_2)}{f_1 \cdot f_2 \leq g_1 \cdot g_2 \sim_{N_1 + N_2} \beta} (*\frac{1}{S}) \\
\frac{\text{DS}_f \leq'_N g \sim \beta \quad \forall \vec{x} : \bar{g}(\vec{x}) = g(\vec{x})/a}{\text{DS}_f \leq'_{a \cdot N} \bar{g} \sim \beta/a} (N_D) \\
\frac{\text{DS}_{f_1} \leq'_N h_1 \sim \beta_1 \quad \forall \vec{x} : h_1(\vec{x}) \leq B \quad f'_2 \sim_{|\cdot|} \beta_2}{\text{DS}_{f_2 \circ f_1} \leq'_N |f'_2 \circ f_1| \cdot h_1 \sim \beta_2 B + \beta_1} (\circ_D) \\
\frac{\text{DS}_{f_i} \leq'_{N_i} h_i \sim \beta \quad \forall i, j : \text{vars}(N_i) \perp \text{vars}(N_j)}{\text{DS}_{\min\{f_1, \dots, f_k\}} \leq'_{\ell_p(N_1, \dots, N_k)} \max\{h_1, \dots, h_k\} \sim \beta} (\min^{\frac{1}{D}}) \\
\frac{f_i \leq g_i \sim_{N_i} \beta \quad \forall i, j : \text{vars}(N_i) \perp \text{vars}(N_j)}{\min\{f_1, \dots, f_k\} \leq \min\{g_1, \dots, g_k\} \sim_{\ell_p(N_1, \dots, N_k)} \beta} (\min^{\frac{1}{S}})
\end{array}$$

Figure 41: Upper Bounds for Composite Functions

Let the *query norm* (denoted N_{qr}) be the norm for which we *can* compute derivative sensitivity. Let the *user-defined norm* (denoted N_{db}) be the norm for which we *want to* compute derivative sensitivity.

Let $N \leq M$, denote that $\|\vec{x}\|_N \leq \|\vec{x}\|_M$ for all \vec{x} . If $N_{qr} \leq N_{db}$, then the rule (\leq_D) allows us to use the computed DS_f for N_{qr} also with the norm N_{db} . But if $N_{qr} \not\leq N_{db}$, then we cannot directly use the sensitivity w.r.t. N_{qr} .

According to rule (N_D) , if a the upper bound to the derivative sensitivity of the function f is β -smooth according to N_{qr} , then, its $\frac{1}{\alpha}$ -times scaled version is $\frac{\beta}{\alpha}$ -smooth in according to the norm $\alpha \cdot N_{qr}$ for any $\alpha > 0$. We compute sensitivity w.r.t. such norm $\alpha \cdot N_{qr}$, that $\alpha \cdot N_{qr} \leq N_{db}$. The sensitivity becomes $\frac{\beta}{\alpha}$ -smooth instead of β -smooth, which affects the amount of noise required to achieve differential privacy.

Lemma 3.74. *Let $\vec{x} = (x_1, \dots, x_k) \in \mathbb{R}^k$, $\vec{y} = (y_1, \dots, y_n) \in \mathbb{R}^n$, $\vec{z} = (z_1, \dots, z_m) \in \mathbb{R}^m$. If $p \geq q \geq 1$, then*

1. $\|\|\vec{x}\|_q, \|\vec{y}\|_q, z_1, \dots, z_m\|_p \leq \|\|\vec{x}\vec{y}\|_q, z_1, \dots, z_m\|_p$;
2. $\|\|\vec{x}\|_p, \|\vec{y}\|_p, z_1, \dots, z_m\|_q \geq \|\|\vec{x}\vec{y}\|_p, z_1, \dots, z_m\|_q$;

where $\vec{x}\vec{y}$ denotes concatenation. If $p = q$, then the inequalities become equalities.

Proof. Since $p, q \geq 1$, we may raise both sides of equations to the powers p or q . The main inequalities that we use in the proof are $a^n + b^n \leq (a + b)^n$ for $n \geq 1$, and $a^n + b^n \geq (a + b)^n$ for $n \leq 1$.

$$\begin{aligned}
\|\|\vec{x}\|_q, \|\vec{y}\|_q, z_1, \dots, z_m\|_p^p &= \left(\sum_{i=1}^k x_i^q \right)^{\frac{p}{q}} + \left(\sum_{i=1}^n y_i^q \right)^{\frac{p}{q}} + \sum_{i=1}^m z_i^p \\
&\leq \left(\sum_{i=1}^k x_i^q + \sum_{i=1}^n y_i^q \right)^{\frac{p}{q}} + \sum_{i=1}^m z_i^p \\
&= \|\|\vec{x}\vec{y}\|_q^p + \sum_{i=1}^m z_i^p = \|\|\vec{x}\vec{y}\|_q, z_1, \dots, z_m\|_p^p .
\end{aligned}$$

$$\begin{aligned}
\|\|\bar{x}\|_p, \|\bar{y}\|_p, z_1, \dots, z_m\|_q^q &= \left(\sum_{i=1}^k x_i^p \right)^{\frac{q}{p}} + \left(\sum_{i=1}^n y_i^p \right)^{\frac{q}{p}} + \sum_{i=1}^m z_i^q \\
&\geq \left(\sum_{i=1}^k x_i^p + \sum_{i=1}^n y_i^p \right)^{\frac{q}{p}} + \sum_{i=1}^m z_i^q \\
&= \|\|\bar{x}\bar{y}\|_p^q + \sum_{i=1}^m z_i^q = \|\|\bar{x}\bar{y}\|_p, z_1, \dots, z_m\|_q^q .
\end{aligned}$$

If $p = q$, then all inequalities in these derivations are equalities. \square

Lemma 3.75. Let N be a composite ℓ_p -norm over $\vec{x} = (x_1, \dots, x_n)$. Let composite seminorms N' and V_1, \dots, V_m be such, that $N = N'(V_1, \dots, V_m)$, and for all $i \in [n]$ let W_i be a seminorm such that $V_i \leq W_i$. Then, $N'(V_1, \dots, V_m) \leq N'(W_1, \dots, W_m)$.

Proof. Let $N = N'(V_1, \dots, V_m)$. The relation $V_i \leq W_i$ implies $\|x_1, \dots, x_n\|_{V_i} \leq \|x_1, \dots, x_n\|_{W_i}$ for all $x_1, \dots, x_n \in \mathbb{R}^n$. Define a new norm $M = N'(W_1, \dots, W_m)$. By definition of a composite ℓ_p -norm, we have the three cases for N' .

- If $N' = |x_j|$ for some $j \in [n]$, then $m = 0$, and hence $\|x_1, \dots, x_n\|_N = \|x_1, \dots, x_n\|_M = |x_j|$.
- If $N' = \alpha z$, then $m = 1$, and we have $\|x_1, \dots, x_n\|_N = \alpha \|x_1, \dots, x_n\|_{V_1}$, and $\|x_1, \dots, x_n\|_M = \alpha \|x_1, \dots, x_n\|_{W_1}$, so $N \leq M$.
- If $N' = \|z_1, \dots, z_m\|_p$, then $\|x_1, \dots, x_n\|_N = \|\|x_1, \dots, x_n\|_{V_1}, \dots, \|x_1, \dots, x_n\|_{V_m}\|_p \leq \|\|x_1, \dots, x_n\|_{W_1}, \dots, \|x_1, \dots, x_n\|_{W_m}\|_p = \|x_1, \dots, x_n\|_M$, so $N \leq M$.

In any case, we get $N \leq M$, which is equivalent to $N'(V_1, \dots, V_m) \leq N'(W_1, \dots, W_m)$. \square

Lemma 3.76. For all $x \in \mathbb{R}$, $(\alpha_1, \dots, \alpha_k) \in \mathbb{R}^k$, $(y_1, \dots, y_m) \in \mathbb{R}^m$: $\|\alpha_1 x, \dots, \alpha_k x, y_1, \dots, y_m\|_p = \|\sqrt[p]{\sum_{i=1}^k \alpha_i^p x, y_1, \dots, y_m}\|_p$.

Proof. Since an ℓ_p -norm is defined for $p \geq 1$, we may raise both sides of equation to the power p . We use the definition of ℓ_p -norm and rewrite the term.

$$\begin{aligned}
\|\alpha_1 x, \dots, \alpha_k x, y_1, \dots, y_m\|_p^p &= \sum_{i=1}^k (\alpha_i x)^p + \sum_{i=1}^m y_i^p \\
&= \left(\sum_{i=1}^k \alpha_i^p \right) x^p + \sum_{i=1}^m y_i^p \\
&= \|\sqrt[p]{\sum_{i=1}^k \alpha_i^p x, y_1, \dots, y_m}\|_p^p .
\end{aligned}$$

\square

Lemma 3.77. Let $\vec{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$. Let N be a composite ℓ_p -norm, defined over variables \vec{x} . There exist $0 \leq \alpha_i, \beta_i \in \mathbb{R}$ for $i \in [n]$, such that $\|\alpha_1 x_1, \dots, \alpha_n x_n\|_p \leq \|x_1, \dots, x_n\|_N \leq \|\beta_1 x_1, \dots, \beta_n x_n\|_q$, where:

- p is the largest ℓ_p -norm constructor in N ;
- q is the smallest ℓ_p -norm constructor in N .

A possible generic solution of estimating utility is to give an upper bound on the noise magnitude. Since widely used additive noise distributions (Laplace, Cauchy, Gaussian) are unbounded, instead of fixing a strict upper bound, we can only define a range within which the noise stays with a sufficiently high confidence.

We note that there also exist DP mechanisms with nice bounded distributions like truncated Laplace [57]. However, they only provide (ϵ, δ) -DP for $\delta > 0$. Such mechanisms can be used in the cases where the error bound should hold for sure, but the privacy bound is more relaxed, and leakage is allowed with a reasonable probability that depends on δ .

Probabilistic upper bound on noise magnitude. The noise magnitude depends linearly on the quantity $\lambda := \frac{c(t)}{b}$, where $c(t)$ is the derivative sensitivity at point t (t is the actual data) and b a parameter that depends on ϵ and β , e.g. $b = (\frac{\epsilon}{\gamma} - \beta)$ for Cauchy distribution $GenCauchy(\gamma)$. The noise that is eventually added to a numeric output is $\lambda \cdot \eta$, where η is sampled from Cauchy distribution. The value λ itself does not give user any intuition how large that noise is. Instead, we would like to get an upper bound on the added noise. Unfortunately, both Laplace and Cauchy distributions are unbounded, and such an upper bound does not exist. However, we can still report a value λ such that error stays below λ with a certain probability (confidence). Knowing that $x \sim \frac{\sqrt{2}}{\pi} \cdot \frac{1}{1+|x|^4}$, we can compute

$$\int_{-1}^1 \frac{\sqrt{2}}{\pi} \cdot \frac{1}{1+|x|^4} dx \approx 0.78 ,$$

thus fixing the probability that noise stays below $\lambda = \frac{c(t)}{b}$ to a constant $p = 0.78$. We would like to make our result more flexible and let the user choose p . Denoting PDF of noise distribution by $f_\eta(x)$, we need to find a such that

$$\int_{-a}^a f_\eta(x) dx = p .$$

For Laplace distribution, the equation $\int_{-a}^a \frac{\epsilon}{2} \cdot e^{-|x|\epsilon} dx = p$ reduces to $1 - e^{-a\epsilon} = p$, from which we get $a = \frac{-\ln(1-p)}{\epsilon}$. Unfortunately, there is no nice solution for Cauchy noise. However, since $\int_{-a}^a f_\eta(x) dx = 2 \cdot \int_0^a f_\eta(x) dx$, and $f_\eta(x)$ is monotone in $[0, \infty)$, we can use window binary search to find a for non-scaled Cauchy distribution. Now, in order to get probability p for $x = \lambda\eta$, we just need to take $a\lambda$ instead of a , as we have

$$p = \int_{-a}^a f_\eta(x) dx = \int_{-a}^a f_\eta(\lambda x) d(\lambda x) = \int_{-a\lambda}^{a\lambda} \frac{1}{\lambda} \cdot f_\eta(\lambda x) dx .$$

An upper bound on noise magnitude can be a fair estimate on its own, but it is not suitable for comparing different SQL queries. The badness of a depends on the query, its result, and its further use by the recipient of the query result. For example, the additive noise ± 5 would almost have no effect on the actual count $y = 1000$, but it would be destructive for $y = 10$. For a single real output y , we can define *relative error* as $\frac{a}{|y|}$ (which is reasonable as far as $y \neq 0$). For a vector of outputs \vec{y} , we still want to have *one* number that characterizes the error. We propose to define the error as $\frac{\|\vec{a}\|_2}{\|\vec{y}\|_2}$, where \vec{a} is the vector of errors.

Precision loss due to continuous approximation. We proposed to use sigmoids $\frac{e^{\alpha x}}{e^{\alpha x} + 1}$ as continuous approximations of indicator functions $x > 0$ to implement filtering. Using sigmoids instead of the original filters causes a precision loss, as we use the value $\frac{e^{\alpha x}}{e^{\alpha x} + 1}$ instead of 0 or 1. The difference is $\frac{1}{e^{\alpha|x|} + 1}$, which goes to 0 when $|x| \rightarrow \infty$. The larger α is, the faster the difference goes to 0, and thus the smaller the precision loss. Unfortunately, we cannot increase α indefinitely to decrease the precision loss. If we want to achieve ϵ -differential privacy then the ϵ , together with the structure of the query, determines a value α_0 such that ϵ -differential privacy can only be achieved if $\alpha < \alpha_0$.

Suppose we have a SUM query using ℓ_1 -norm to join row norms, and input rows are chosen from a certain distribution that does not change when input size (n , the number of rows) changes, which would

be a quite common scenario. Then the query result is roughly proportional in n , and the sensitivity (and thus the added noise) is roughly constant. Thus the relative error from added noise goes to 0 as $n \rightarrow \infty$. The relative error from sigmoids, however, is roughly constant as $n \rightarrow \infty$.

The larger n gets, the larger the error from sigmoids becomes relative to the error from added noise. We would like to decrease the former even if this increases the latter. We saw that we cannot increase α to achieve this. Instead, we use so-called *precise sigmoids*, keeping α constant but increasing the second parameter a as n increases.

Precise sigmoid. To get a higher precision than that of an ordinary sigmoid but still maintain α -smoothness, we use an extra parameter a in addition to α . We use the sigmoid $\sigma(x) = \frac{e^{\alpha x}}{e^{\alpha x} + 1}$ but instead of its actual sensitivity $\sigma'(x)$, we use $c(x) = \frac{ae^{\alpha x}}{(e^{\alpha x} + 1)^2}$, which is an α -smooth upper bound on $\sigma'(x)$. The smooth sensitivity is $\frac{a}{\alpha}$ times higher than that of the original sigmoid, but the difference from the precise filter value (0 or 1) is $\frac{e^{\alpha x} + 1}{e^{\alpha x}}$ times smaller. If the probability density function of x is roughly constant near $x = 0$ then, for all $y \in (0, 1)$, the probability that the difference from the precise filter value is larger than y is

$$\begin{aligned} \Pr\left(\frac{1}{e^{\alpha x} + 1} > y\right) &= \Pr\left(e^{\alpha x} < \frac{1}{y} - 1\right) = \Pr\left(x < \frac{1}{\alpha} \ln\left(\frac{1}{y} - 1\right)\right) \approx \\ &\approx \frac{\alpha}{a} \Pr\left(x < \frac{1}{\alpha} \ln\left(\frac{1}{y} - 1\right)\right) = \frac{\alpha}{a} \Pr\left(\frac{1}{e^{\alpha x} + 1} > y\right) \end{aligned}$$

i.e. $\frac{a}{\alpha}$ times smaller than in the original sigmoid.

Then, assuming the probability density function of the input is roughly constant near the pivot point of the filter, the relative error from sigmoids will be roughly $\frac{a}{\alpha}$ times smaller and the relative error from added noise will be roughly $\frac{a}{\alpha}$ times larger.

Suppose that with an ordinary sigmoid the relative error from the sigmoid is k times larger than that from added noise. Then we can take $a = \alpha \sqrt{k}$ to make the two sources of error roughly equal and get the smallest amount of total noise. For the kind of queries considered above, k is roughly proportional to n , thus a would have to be roughly proportional to \sqrt{n} . The total error would then be inversely proportional to \sqrt{n} .

If more than one sigmoid is used then the total error would still be inversely proportional to \sqrt{n} . If no sigmoids are used then the total error would be inversely proportional to n .

We can get similar results for tauoids $\frac{2}{e^{-\alpha x} + e^{\alpha x}}$, used as continuous approximation of equality $x = 0$. Here we can allow nonzero probability at the pivot point, but near the pivot point (excluding the pivot point itself) the probability density function must still be roughly constant.

If the filtered values are integers then we do not have to use sigmoids and tauoids but instead use the precise functions $x > y \iff \min(1, \max(0, x - y))$ and $x = y \iff 1 - \min(1, \max(0, |x - y|))$. Then the only error is from adding noise and it would be inversely proportional to n .

Utility loss vs privacy loss. In the previous paragraphs, we considered how precision improves asymptotically when the number of rows $n \rightarrow \infty$. Now we consider how precision depends on the required privacy level. Let us define *utility loss* as the total relative error. Let us define *privacy loss* as the differential-privacy ϵ .

First consider the case without sigmoids. The relative error from noise is $E_n = \frac{c(x)}{bf(x)}$. Here c is a β -smooth upper bound on DS_f , thus it is monotonically decreasing in β . By default, we have $b = \beta = \frac{\epsilon}{2(\gamma+1)}$. Thus $E_n = \frac{2(\gamma+1)c(x)}{\epsilon f(x)}$ and c is also monotonically decreasing in ϵ . So if we increase ϵ by a factor of $k > 1$ then E_n decreases by a factor of at least k .

The global sensitivity can be defined as $GS_f = \max_x DS_f(x)$. Note that the product of utility loss and privacy loss, $E_n \cdot \epsilon$, is monotonically decreasing in ϵ and

$$\lim_{\epsilon \rightarrow 0} E_n \cdot \epsilon = \frac{2(\gamma+1)c(x)}{f(x)} \cdot GS_f \ ,$$

$$\lim_{\epsilon \rightarrow \infty} E_n \cdot \epsilon = \frac{2(\gamma+1)c(x)}{f(x)} \cdot DS_f(x) \ .$$

Thus $E_n \cdot \epsilon$ goes from $\frac{2(\gamma+1)c(x)}{f(x)} \cdot \text{GS}_f$ to $\frac{2(\gamma+1)c(x)}{f(x)} \cdot \text{DS}_f(x)$ as ϵ goes from 0 to ∞ . Thus the larger ϵ is, the bigger the advantage of using smooth derivative sensitivity instead of global sensitivity.

Now consider the case that uses sigmoids. We optimize the tradeoff between sigmoid precision and added noise. The total relative error is $E = 2E_n$ and we get similar results for $E \cdot \epsilon$ as above for $E_n \cdot \epsilon$.

3.3.9 Combining Derivative Sensitivity for Row Multiplicities and Components. In Sec. 3.3.7, we showed how to compute smooth derivative sensitivity for relational algebra queries, where the derivatives are w.r.t. the multiplicities of rows in input tables. The distance between two databases is the number of rows we have to add to or remove from the first database to transform it into the second database. The shortcoming of this approach is that if the databases differ by only a small change in one row then the distance is the same as when they differ by a large change in that row (in both cases we need to remove one row and add another).

In Sec. 3.3.8, we showed how to compute smooth derivative sensitivity in Banach spaces, where the derivatives are Fréchet derivatives w.r.t. the values of rows in input tables. The distance between two databases quantifies the amount of changes we have to make to the values of the rows in the first database to transform it into the second database. This removes the shortcoming of the first approach but introduces another: the number of rows in the tables of the two databases must now be the same.

We would like to combine the two approaches, allowing to add and remove rows to/from the database and also make changes to individual rows so that small changes correspond to small distances.

Because in this section we will use two different kinds of derivative sensitivity (w.r.t. row multiplicities and w.r.t. components), let us call the derivative sensitivity w.r.t. components (i.e. in Banach spaces) *Banach sensitivity* or *Banach derivative sensitivity*. In ambiguous cases, we call the derivative sensitivity w.r.t. row multiplicities also *adding/removing rows sensitivity*.

3.3.9.1 Generalization of Derivative Sensitivity.

Extending the definitions. First, let us extend the notion of derivative sensitivity. We will need the following lemma.

Lemma 3.78. *Let X be a convex subset of a Banach space with norm $\| \cdot \|$, $f : X \rightarrow \mathbb{R}$ a function, $c : X \rightarrow \mathbb{R}_+$ a β -smooth upper bound on the derivative sensitivity (in the sense of Def. 3.55) of f . Then*

$$|f(x') - f(x)| \leq e^{\beta\|x' - x\|} c(x) \|x' - x\|$$

Proof. Let $x, x' \in X$. By the Mean Value Theorem for Banach spaces, there exists $\lambda \in (0, 1)$ such that

$$f(x') - f(x) = df_{(1-\lambda)x + \lambda x'}(x' - x)$$

Then

$$|f(x') - f(x)| \leq \|df_{(1-\lambda)x + \lambda x'}\| \cdot \|x' - x\| = \text{DS}_f((1-\lambda)x + \lambda x') \cdot \|x' - x\|$$

Because c is an upper bound of DS_f ,

$$\text{DS}_f((1-\lambda)x + \lambda x') \leq c((1-\lambda)x + \lambda x')$$

Because c is β -smooth and $\|((1-\lambda)x + \lambda x') - x\| = \lambda\|x' - x\|$,

$$c((1-\lambda)x + \lambda x') \leq e^{\beta\lambda\|x' - x\|} c(x) \leq e^{\beta\|x' - x\|} c(x)$$

Thus

$$|f(x') - f(x)| \leq e^{\beta\|x' - x\|} c(x) \|x' - x\|$$

□

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} P_C(x, c, \mathbf{A}) &= \lim_{\epsilon \rightarrow 0} \int_Y P_Y(x, y, \mathbf{A}) \cdot f_{YC}(y, c) dy \\
&= f_X(x | \mathbf{A}) \cdot \int_Y f_{YC}(y, c) dy \\
&= f_X(x | \mathbf{A}) .
\end{aligned}$$

$$\begin{aligned}
\lim_{\epsilon \rightarrow \infty} P_C(x, c, \mathbf{A}) &= \lim_{\epsilon \rightarrow \infty} \int_Y P_Y(x, y, \mathbf{A}) \cdot f_{YC}(y, c) dy \\
&= \int_Y f_X(x | \mathbf{C}^g = y, \mathbf{A}) f_{YC}(y, c) dy \\
&= f_X(x | \mathbf{C}^g = c, \mathbf{A}) .
\end{aligned}$$

Take the definition of $P_P(x, \mathbf{A})$ from Equation 41. We make use of the proven limits for $P_Y(x, c, \mathbf{A})$.

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} P_P(x, \mathbf{A}) &= \int_C P_C(x, c, \mathbf{A}) \cdot \frac{f_X(x|c) \cdot f_C(c)}{f_X(x)} dc \\
&= f_X(x | \mathbf{A}) \int_C \frac{f_X(x|c) \cdot f_C(c)}{f_X(x)} dc \\
&= f_X(x | \mathbf{A}) .
\end{aligned}$$

$$\begin{aligned}
\lim_{\epsilon \rightarrow \infty} P_P(x, \mathbf{A}) &= \int_C P_C(x, c, \mathbf{A}) \cdot \frac{f_X(x|c) \cdot f_C(c)}{f_X(x)} dc \\
&= \int_C f_X(x | \mathbf{C}^g = c, \mathbf{A}) \frac{f_X(x|c) \cdot f_C(c)}{f_X(x)} dc \\
&= \frac{1}{f_X(x)} \cdot \int_C f_X(x | \mathbf{C}^r = c, \mathbf{A}) f_X(x | \mathbf{C}^r = c) f_C(c) dc .
\end{aligned}$$

□

Composability. Even if estimation distribution is easy for a single output, it can be much more difficult for several outputs. We want to get composability results for the upper bounds on posterior probability obtained for a single attacker.

Since $P_C(x, c, \mathbf{A})$ and $P_P(x, \mathbf{A})$ depend on $P_Y(x, y, \mathbf{A})$, and the issue of analyzing several timepoints is integration through all possible values of z , we will only analyze composability of $P_Y(x, y, \mathbf{A})$.

Let $c = (c_1, \dots, c_n)$ be the true output. First of all, since we want to bound attacker's inference probability by $f_X(x|c)$, we need to take into account at least all elements of c on which $f_X(x|c)$ depends, and we cannot estimate it only based on $f_X(x|c_j)$ for some particular $j \in [n]$. Hence, we need to constrain ourselves to the situations $c = c_a, c_b$, where $f_X(x|c) = f_X(x|c_a)$. E.g. for a counting histogram, when the attacker wants to guess the group a particular individual belongs to, we have $f_X(x|c_1, \dots, c_n) = \frac{c_j}{m}$, where c_j is the true group to which the input belongs. Indeed, the value m does depend on the other counts as well, but we treat m as a public parameter that makes c_i correlated.

Parallel composition for independent c_a, c_b is similar to the one of DP.

Theorem 3.97. Let $C = C_a \times C_b$. Let $f_Z(z_a, z_b | \mathbf{A}) = f_{Z_a}(z_a | \mathbf{A}) \cdot f_{Z_b}(z_b | \mathbf{A})$ for all $z_a \in C_a, z_b \in C_b$. We have

$$P_Y(x, y, \mathbf{A}) = \frac{\int_{C_a} f_X(x|z, \mathbf{A}) \cdot f_{YC_a}(y, z) f_{Z_a}(z | \mathbf{A}) dz}{\int_{C_a} f_{YC_a}(y, z) f_{Z_a}(z | \mathbf{A}) dz} . \quad (43)$$

Proof. For shortness of notation, let us remove the additional knowledge \mathbf{A} from all conditional probabilities and make it implicit. This will not affect the proofs, as \mathbf{A} is found in the premises of all theorems. By definition of $P_Y(x, y, \mathbf{A})$ (Equation 38), we have

$$\begin{aligned} P_{Y(x, y, \mathbf{A})} &= \frac{\int_C f_X(x|z) f_{YC}(y, z) f_Z(z) dz}{\int_C f_{YC}(y, z) f_Z(z) dz} \\ &= \frac{\int_{C_a} \int_{C_b} f_X(x|z_a) f_Y(y_a, y_b|z_a, z_b) f_Z(z_a, z_b) dz_a dz_b}{\int_{C_a} \int_{C_b} f_Y(y_a, y_b|z_a, z_b) f_Z(z_a, z_b) dz_a dz_b} \end{aligned}$$

Assuming that the noise distribution depends only on a single output (as in the case of Laplace noise), we get $f_Y(y_a, y_b|z_a, z_b) = f_{Y_a}(y_a|z_a) \cdot f_{Y_b}(y_b|z_b)$.

For independent variables, $f_Z(z_a, z_b) = f_{Z_a}(z_a) f_{Z_b}(z_b)$. This allows to split both the numerator and the denominator into a product of two independent integrals, and we get

$$\begin{aligned} P_{Y(x, y, \mathbf{A})} &= \frac{\int_{C_a} \int_{C_b} f_X(x|z_a) f_{Y_a}(y_a|z_a) f_{Y_b}(y_b|z_b) f_{Z_a}(z_a) f_{Z_b}(z_b) dz_a dz_b}{\int_{C_a} \int_{C_b} f_{Y_a}(y_a|z_a) f_{Y_b}(y_b|z_b) f_{Z_a}(z_a) f_{Z_b}(z_b) dz_a dz_b} \\ &= \frac{\int_{C_a} f_X(x|z_a) f_{Y_a}(y_a|z_a) f_{Z_a}(z_a) dz_a \cdot \int_{C_b} f_{Y_b}(y_b|z_b) f_{Z_b}(z_b) dz_b}{\int_{C_a} f_{Y_a}(y_a|z_a) f_{Z_a}(z_a) dz_a \cdot \int_{C_b} f_{Y_b}(y_b|z_b) f_{Z_b}(z_b) dz_b} \\ &= \frac{\int_{C_a} f_X(x|z) f_{YC_a}(y, z) f_{Z_a}(z) dz}{\int_{C_a} f_{YC_a}(y, z) f_{Z_a}(z) dz} \end{aligned}$$

□

It is more complicated with sequential composition. At another extreme, c_i may be completely correlated, giving $f_{Z_b}(g(z_a)|\mathbf{C}_a^g = z_a) = 1$ for a deterministic function g . In a particular case where $g(z) = z$, we get something similar to sequential composition of DP. This basically allows to use the same z instead of z_a and z_b , so that integration over z_b can be avoided. This can be useful to analyze queries applied multiple times to the same data.

Theorem 3.98. *Let $C = C_a \times C_b$, where $C_a = C_b$. Let $f_{Z_b}(z_a|\mathbf{C}_a^g = z_a, \mathbf{A}) = 1$ for all $z_a \in C_a$. We have*

$$P_Y(x, y, \mathbf{A}) = \frac{\int_{C_a} f_X(x|z, \mathbf{A}) \cdot f_{YC_a}(y_a, y_b, z) f_{Z_a}(z | \mathbf{A}) dz}{\int_{C_a} f_{YC_a}(y_a, y_b, z) f_{Z_a}(z | \mathbf{A}) dz} \quad (44)$$

where $f_{YC_a}(y_a, y_b, z) := f_{YC_a}(y_a, z) \cdot f_{YC_a}(y_b, z)$

Proof. As $f_{Z_b}(z_a|\mathbf{C}_a^g = z_a) = 1$ for all $z_a \in C_a$, we can write $f_Z(z_a, z_b) = f_{Z_a}(z_a)$ and only integrate over C_a . We are left with

$$\begin{aligned} P_Y(x, y, \mathbf{A}) &= \frac{\int_{C_a} f_X(x|z_a) f_Y(y_a, y_b|z_a, z_a) f_Z(z_a) dz_a}{\int_{C_a} f_Y(y_a, y_b|z_a, z_a) f_Z(z_a) dz_a} \\ &= \frac{\int_{C_a} f_X(x|z, \mathbf{A}) \cdot f_{YC_a}(y_a, y_b, z) f_{Z_a}(z) dz}{\int_{C_a} f_{YC_a}(y_a, y_b, z) f_{Z_a}(z) dz} \end{aligned}$$

□

Finally, it is possible that the outputs are strongly correlated, but $g(z) \neq z$. Let us prove a small lemma first.

Lemma 3.99. Let \mathcal{M}_q be any ϵ -differentially private mechanism. Let $z, z' \in q(X)$. Let $R = \max_{x \in q^{-1}(z), x' \in q^{-1}(z')} \|x - x'\|$. For all $y, y' \in \mathcal{M}_q(X)$ we have

$$f_Y(y|z) \leq f_Y(y|z') \cdot e^{\epsilon R} .$$

Proof. We have $z = q(x)$ and $z' = q(x')$ for some $x, x' \in X$. Since $q(x)$ can be computed from x , we have $f_Y(y|q(x), x) = f_Y(y|x)$. Let $q^{-1}(z) := \{x \mid q(x) = z\}$. Denote $X_z := q^{-1}(z)$ and $X_{z'} := q^{-1}(z')$. We have

$$\begin{aligned} \frac{f_Y(y|z)}{f_Y(y|z')} &= \frac{\int_{x \in X_z} f_Y(y|x, z) f_X(x|z) dx}{\int_{x' \in X_{z'}} f_Y(y|x', z') f_X(x'|z') dx'} \\ &= \frac{\int_{x \in X_z} f_Y(y|x) f_X(x|z) dx}{\int_{x' \in X_{z'}} f_Y(y|x') f_X(x'|z') dx'} \\ &= \int_{x \in X_z} \frac{f_Y(y|x) f_X(x|z)}{\int_{x' \in X_{z'}} f_Y(y|x') f_X(x'|z') dx'} dx \\ &\leq \int_{x \in X_z} \frac{e^{\epsilon \|x - x'\|}}{\int_{x' \in X_{z'}} f_X(x'|z') dx'} f_X(x|z) dx \\ &\leq e^{\epsilon R} \frac{\int_{x \in X_z} f_X(x|z) dx}{\int_{x' \in X_{z'}} f_X(x'|z') dx'} = e^{\epsilon R} . \end{aligned}$$

□

For $g(z) \neq z$, the estimation is more difficult, as we want to get a general upper bound for all possible definitions of g . One idea is to apply the worst case bound to the Z_b part.

Theorem 3.100. Let $C = C_a \times C_b$, $f_{Z_b}(g(z_a) | \mathbf{C}_a^g = z_a, \mathbf{A}) = 1$ for a deterministic function g . We have

$$P_Y(x, y, \mathbf{A}) \leq \max_{z_a \in Z_a, z_b \in Z_b} \left(\frac{f_{Y C_b}(y_b, z_a)}{f_{Y C_b}(y_b, z_b)} \right) \cdot \frac{\int_{C_a} f_X(x|z, \mathbf{A}) \cdot f_{Y C_a}(y_a, z) f_{Z_a}(z | \mathbf{A}) dz}{\int_{C_a} f_{Y C_a}(y_a, z) f_{Z_a}(z | \mathbf{A}) dz} .$$

Proof. Let g be a deterministic function such that $f_{Z_b}(g(z_a) | \mathbf{C}_a^g = z_a, \mathbf{A}) = 1$. We have

$$P_Y(x, y, \mathbf{A}) = \frac{\int_{C_a} f_X(x|z_a) \cdot f_{Y_a}(y_a|z_a) f_{Y_b}(y_b|g(z_a)) f_{Z_a}(z_a) dz_a}{\int_{C_a} f_{Y_a}(y_a|z_a) f_{Y_b}(y_b|z_b) f_{Z_a}(z_a) dz_a} .$$

Differently from the previous case, we cannot compute the quantity more precisely unless we know g . We rewrite the expression as

$$P_Y(x, y, \mathbf{A}) = \int_{C_a} f_X(x|z) \cdot f_{Y_a}(y_a|z) \cdot f_{Z_a}(z) \cdot \frac{1}{\int_{C_a} f_{Y_a}(y_a|z_a) \frac{f_{Y_b}(y_b|z_b)}{f_{Y_b}(y_b|g(z))} f_{Z_a}(z_a) dz_a} dz .$$

A trivial upper bound on $\frac{f_{Y_b}(y_b|g(z))}{f_{Y_b}(y_b|z)}$ is $\max_{z_a \in Z_a, z_b \in Z_b} \frac{f_{Y_b}(y_b|z_a)}{f_{Y_b}(y_b|z_b)}$. This upper bound is independent of the integration variables and can be taken out of integration. □

We can instantiate Theorem 3.100 on an ϵ -DP mechanism.

Corollary 3.101. Let $C = C_1 \times \dots \times C_n$. Let an ϵ -DP mechanism w.r.t. norm $\|\cdot\|$ be applied to each C_j . Let $\max_{x, x' \in X} \|x - x'\| = m$. We have

$$P_Y(x, y, \mathbf{A}) \leq e^{\epsilon \cdot m(n-1)} \cdot \frac{\int_{C_1} f_X(x|z, \mathbf{A}) \cdot f_{Y C_1}(y_1, z) f_{Z_1}(z | \mathbf{A}) dz}{\int_{C_1} f_{Y C_1}(y_a, z) f_{Z_1}(z | \mathbf{A}) dz} .$$

```

constr ::= exact
         | total int
         | set {string}
         | range double double

         | totalUnif int
         | setUnif {string}
         | rangeUnif double double

         | setPrior {( string , float )}
         | rangePrior float {( float , float )}

         | normal double double

```

Figure 44: Possible Constraints on Table Attributes

Proof. By assumption of Corollary 3.101, we are dealing with an ϵ -DP mechanism. by Lemma 3.99, we have $\frac{f_{y_b}(y_b|g(z))}{f_{y_b}(y_b|z)} \leq e^\epsilon \cdot \max_{x,x'} \|x - x'\| \leq e^\epsilon \cdot m(n-1)$. The claim now follows directly from Theorem 3.100. \square

3.3.11 Propagation of Constraints on Attributes. In both sensitivity and guessing advantage analyses, we make use of the bounding ranges from which the values may come. While the upper bound on derivative sensitivity may be infinite for $X = \mathbb{R}$, it can be much smaller for a bounded subset $X \subset \mathbb{R}$. The user had to provide possible bounding ranges on the inputs himself. For small models, such approach is sufficient. However, in large models, one might want to compute sensitivity and guessing advantage w.r.t. some intermediate table, e.g. a receiving party may want to know how much leaks about the data that has already gone through some preprocessing. Knowing ranges of intermediate tables can also be useful for composition. While the user can set bounding ranges only on the direct input tables, we need automated constraint derivation to determine the ranges of attributes of intermediate tables as well.

The list of possible constrains on table attributes is given in Figure 44. We see that the constraints can be split to two main types: discrete and numeric.

We consider SQL queries of the form

```

SELECT a1 AS b1, ..., ak AS bk FROM t1 AS s1, ..., tn AS sn WHERE condition
                                GROUP BY a1, ..., ak-1 .

```

3.3.11.1 Numeric Attributes. The numeric constraints of Figure 44 allow to define just an interval $[a, b]$ for some $a \leq b \in \mathbb{R}$, as well as some distribution on that interval, which is currently just a concatenation of a finite number of uniformly distributed intervals. Propagation of distributions is much more complicated than propagating numbers. In our analyses, defining a range without a particular distribution gives a safe noise overestimation, so first of all we may concentrate on propagating ranges without maintaining the distribution.

First of all, we discuss how to compute output ranges for real functions $f : \mathbb{R} \rightarrow \mathbb{R}$. These would correspond to the computation under SELECT statement. In general, overestimated ranges give us higher function sensitivity and result in larger noise, so it is safe to make an overestimation of the actual range.

Interval arithmetic. An obvious solution to function output range estimation is to use interval arithmetic. We need to define operations between intervals. For example some basic arithmetic operations might be defined as

$$\begin{aligned} [a, b] + [c, d] &= [a + b, c + d] \\ -[a, b] &= [-b, -a] \\ [a, b] \times [c, d] &= [\min\{ac, ad, bc, bd\}, \max\{ac, ad, bc, bd\}] \end{aligned}$$

These operations can then be composed to find the output range for more complex expressions. Since not all such functions yield the exact output range we can divide the input interval into smaller subintervals to produce more precise results, as described e.g. in [63, p. 306].

When using interval arithmetic to evaluate functions where the same variable appears multiple times, it tends to give us very wide output intervals to the point where the result is completely useless. This happens because we assume that all the variables in the function are independent [64, p. 36]. Let us look for alternative methods that solve this problem.

Affine arithmetic. An alternative to interval arithmetic is the *affine arithmetic*. Using affine arithmetic we can write the intervals in the affine form

$$\hat{x} = x_0 + \epsilon_1 x_1 + \epsilon_2 x_2 + \dots + \epsilon_n x_n$$

where $\epsilon_i \in [-1, 1]$, $i \in [1 \dots n]$. For example, if $n = 1$, then x_0 can be viewed as the central point of the interval, and sliding $\epsilon_i \in [-1, 1]$ gives the values $[x_0 - x_1, x_0 + x_1]$. If two affine representations use the same ϵ_i , it means that these two variables are correlated.

Let us look at the function $f(x) = x - x$. Assuming $x \in [0, 3]$, we can find the output range of the function using interval arithmetic:

$$\begin{aligned} [0, 3] - [0, 3] &= [0, 3] + (-[0, 3]) \\ &= [0, 3] + [-3, 0] \\ &= [-3, 3] \end{aligned}$$

It turns out that interval arithmetic does not give us the exact answer $[0, 0]$, but instead gives us a much wider interval. Now let us find the output range using affine arithmetic:

$$\begin{aligned} (1.5 + 0.5\epsilon_1) - (1.5 + 0.5\epsilon_1) &= (1.5 - 1.5) + (0.5 - 0.5)\epsilon_1 \\ &= 0 + 0\epsilon_1 = 0 \end{aligned}$$

We get the exact result because every occurrence of variable x has the term $x_1\epsilon_1$, which cancel out after the subtraction. If we wanted to write different variables instead we would write

$$\begin{aligned} \hat{x} &= x_0 + x_1\epsilon_1 \\ \hat{y} &= y_0 + y_1\epsilon_2 \end{aligned}$$

Converting between intervals and affine forms. We can convert any affine form $\hat{x} = x_0 + x_1\epsilon_1 + \dots + x_n\epsilon_n$ to an interval as follows:

$$x = [x_0 - r, x_0 + r], \text{ where } r = \sum_{i=1}^n |x_i|.$$

It should however be kept in mind that by converting to interval form, we lose the information about the dependencies between variables. We can also get the affine form of the interval $[a, b]$

$$\hat{x} = \frac{a+b}{2} + \frac{b-a}{2}\epsilon_1,$$

as described in [64, p.47-49].

When doing this conversion, each different variable should have its own epsilon symbol. Libaffa, a C++ implementation of affine arithmetic uses a map data structure for mapping between variable names and corresponding coefficients [65]. In Haskell it can be defined like this:

```
data AAF a = AAF a (Map String a) a
```

In this definition the first field represents the 0-th term. The map binds variable names with corresponding terms. All non-affine operations such as multiplication and square root add new terms after each operation (i.e. error terms) and to keep the code simple, we collect all these terms into a single error term, which is what the last field in the AAF structure is for.

Non-affine functions. Some functions do not give us an exact affine answer. For example consider multiplying two affine forms $\hat{x} = x_0 + x_1\epsilon_1 + \dots + x_n\epsilon_n$ and $\hat{y} = y_0 + y_1\epsilon_1 + \dots + y_n\epsilon_n$. We would end up with an affine part

$$\hat{z}^a = x_0y_0 + (x_0y_1 + y_0x_1)\epsilon_1 + \dots + (x_0y_n + y_0x_n)\epsilon_n$$

and a non-affine part

$$z^* = (x_1\epsilon_1 + \dots + x_n\epsilon_n)(y_1\epsilon_1 + \dots + y_n\epsilon_n)$$

Since we want an affine approximation of the result, we have to add a new term to the affine part to approximate the non-affine part

$$\hat{z} = \hat{z}^a + (x_1 + \dots + x_n)(y_1 + \dots + y_n)\epsilon_{n+1} ,$$

as described in [64, p.70].

In practice it makes sense to collect all such terms of subsequent operations into a single term as we did in the Haskell implementation example. This makes sure that the affine forms do not grow unnecessarily large. It can even be reasonable to have more than one error term to reduce error when multiplying or raising the affine forms to some power [64, p. 81].

Min-range approximation. We can generalize non-affine functions in a more general way. Let us consider some binary non-affine function $f(\hat{x}, \hat{y})$. We can approximate this function by finding some good values for the coefficients $\alpha, \beta, \zeta \in \mathbb{R}$ in

$$f^a = \alpha\hat{x} + \beta\hat{y} + \zeta$$

Computationally one of the best methods is to use the min-range approximation. We find the coefficients in such a way that the resulting range of the function is as narrow as possible. An example of min-range approximation of a log function is given in Figure 45, where the input range $[a, b]$ gives an output range $[c, d]$, and all function values stay inside the blue skewed rectangle. Even though this method is not optimal in terms of minimizing the error, it is easier to implement and has a smaller overshoot than Chebyshev approximation, which aims to minimize the maximum error [64, p. 56,64-65].

Fixing the values of variables. Let us look at a case where we need to find the output range of a function when we know the range of one variable, and a set of exact values for a second variable. We have affine forms $\hat{x} = x_0 + x_1\epsilon_1$ and $\hat{y} = y_0 + y_2\epsilon_2$ representing the respective variables. Once we evaluate the function using these affine forms, we end up with a new affine form

$$\hat{z} = z_0 + z_1\epsilon_1 + z_2\epsilon_2 + z_3\epsilon_3 .$$

The noise symbols ϵ_1 and ϵ_2 correspond to the variables x and y respectively. If we later want to fix y to some value a we can substitute $\epsilon_2 = \frac{a}{y}$ which gives us

$$\hat{z}' = z_0 + z_1\epsilon_1 + z_2\frac{a}{y} + z_3\epsilon_3 .$$

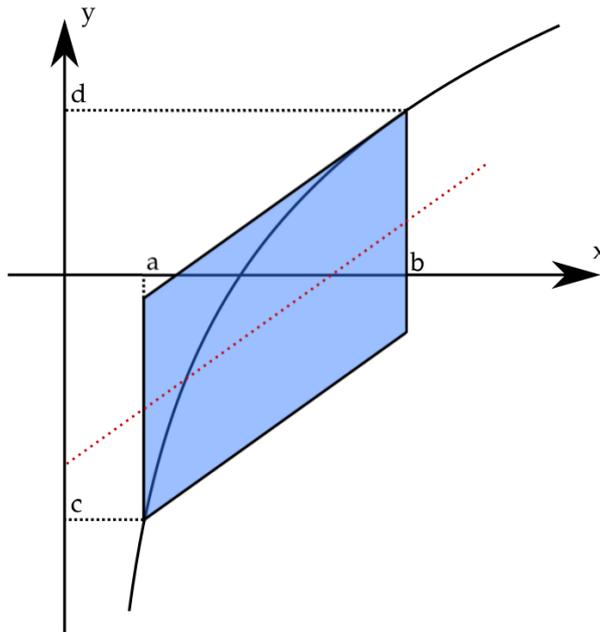


Figure 45: Min-Range Approximation for the log Function

This means we only have to find \hat{z} once and then do a relatively cheap computation to fix the values afterwards. It is also possible to fix the value to some subrange of $[y_0 - y_2, y_0 + y_2]$. This approach is very useful when we have a large number of similar computations, like applying the same function to n rows of an input table. It can also be applied in cases where a variable is represented as a discrete set of possible values rather than a range.

Subdividing input ranges. When estimating the output range of an expression given some input range, it is possible to get more accurate results by dividing the input range into smaller subdivisions. Each of these subdivisions can then be evaluated separately which results in smaller zonotopes that follow the function more closely [63, p. 306-307]. An example of dividing input range $[a, b]$ into smaller subranges for a log function is given in Figure 46. We see that the function values are now covered by the blue skewed rectangles much more tightly. When we later want to fix the input to some specific value or range, we only have to do that in the input subdivisions that contain the new value.

Distributions with unbounded domain. In practice, some data attributes may not have nice upper and lower bounds. One possible realistic distribution is the normal distribution $N(\mu, \sigma^2)$. Since the range of normally distributed inputs is unbounded, it seems that knowing that the data is distributed normally will not provide good bounds on the function output. However, the values still remain within a certain range with high probability. If the distribution has bell shape, it does not make sense to consider elements that are too far from the center. For $N(\mu, \sigma^2)$ we can take $R = \mu + 3 \cdot \sqrt{2}\sigma$, which covers $\text{erf}(3) \approx 0.9999779$ of the input space. We can then assume that our inputs are coming from the range $[\mu - R, \mu + R]$. In general, we can let η be a parameter and find ξ such that $\text{erf}(\xi) = 1 - 2^{-\eta}$. E.g. we have $\xi \approx 5.1$ for $\eta = 40$.

This approach can be generalized to any distribution with well-defined cumulative density function (CDF) $F(x)$. The probability weight of the range $[\mu - R, \mu + R]$ can be found by definition of CDF as $F(\mu + R) - F(\mu - R)$.

3.3.11.2 Filters and Groups. So far, we have shown how to compute the ranges of values computed under the SELECT statement, and how to optimize the computation of range w.r.t. the number of rows in the input table. We now want to learn how these ranges change after some grouping and/or filtering is

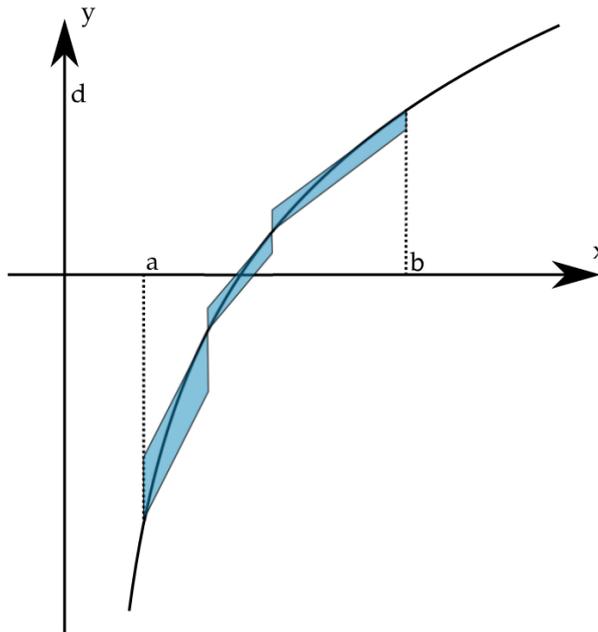


Figure 46: Subdividing the Input Range Gives More Accurate Results

applied. Similarly to the sensitivity analysis, we consider a finite number m of possible groups, so that a single GROUP BY query can be viewed as m instances of query with an additional filter that selects the group.

It remains to show how to handle the filters. It is easy to apply a public filter, as we can discard the mismatching row directly. It is more difficult with a filter that contains private values, for which we may only use ranges. Similarly to sensitivity analysis, a filter can be viewed as a function $\sigma(x) \in \{0, 1\}$, and if we know that an aggregation is applied afterwards, we can leave e.g. $\sigma(x)$ for COUNT queries and $\sigma(x) \cdot f(x)$ for SUM queries. It is more complicated for queries without aggregations, as some ranges may allow for both positive and negative filtering, so we even do not know whether the row belongs to the result or not. There are three main possible ways to go:

1. Allow leaking additional information and use exact value to determine whether to leave the row.
2. Leave only those rows that definitely satisfy the filter.
3. Remove only those rows that definitely do not satisfy the filter.

In general, in derivative sensitivity analysis we expect that the number of rows is public, so if we use the resulting table in sensitivity or guessing advantage analysis directly, then we will do it on assumption that the attacker has learned the total number of rows, which is fine, as we assume that the DP attacker is given "all the records except the one he is trying to guess" anyway.

The other two approaches would give an attacker a lower and an upper bound on the number of rows in the table. This range would have a potential application to combined sensitivity analysis, but we have not considered it so far. If the output is not used in the analysis immediately, but is applied in some subsequent query, then the third approach leaves the greatest variety of possible rows and gives a valid range of possible values, so this is what we need for computing overestimated ranges. If the following query is an aggregation, then for all these rows we will have $\sigma(x) \in \{0, 1\}$, which can be treated as an ordinary interval.

We let the analyzer decide which approach is the most suitable. This depends on the type of the analyzed query.

- **Plain SELECT-query.** The range of an output attribute is defined as the union of all intervals. Hence, a safe overapproximation would be to leave all rows for which we are not sure whether they will be filtered out or not. We apply the method (2).

- **Aggregation query.** The range of an output attribute depends on the aggregation type.
 - **COUNT.** We get the smallest possible count c_{lb} by leaving only those rows that definitely pass the filter, i.e. method (3). We get the largest possible count c_{ub} by leaving only those rows that definitely pass the filter, i.e. method (2). We define the range of a COUNT query output as $[c_{lb}, c_{ub}]$.
 - **SUM.** If we know in advance that the summed values are non-negative, then we can find the sum range similarly to the COUNT, computing the sum s_{lb} over lower bounds of ranges of the rows left by method (3), and s_{ub} over upper bounds of ranges of the rows left by method (2). If the summed values can be negative, we need to include *all* negative lower bounds into s_{lb} , and *all* positive upper bounds into s_{ub} . Finally, we define the range of a SUM query output as $[s_{lb}, s_{ub}]$.
 - **MIN/MAX.** For MIN, we compute the minimum m_{lb} over lower bounds of ranges of the rows left by (2), and the minimum m_{ub} over lower bounds of ranges of the rows left by (3). For MAX, we compute the maximum m_{lb} over upper bounds of ranges of the rows left by (3), and the maximum m_{ub} over upper bounds of ranges of the rows left by (2). We define the range of a MIN/MAX query output as $[m_{lb}, m_{ub}]$.

3.3.11.3 Discrete Attributes. So far, our analyzers are using discrete attributes only for comparison. one of the following may happen to a discrete attribute x :

- copied to the output table as it is (in plain SELECT-queries, or when grouping by x);
- used in a comparison-based filter.

In these cases, a discrete attribute is handled similarly to a numeric attribute. Instead of a *range*, we define a data structure *set*, which contains a (possibly overestimated) set of values S from which this attribute may come. We allow a *universal set*, which is analogous to the interval $[-\infty, \infty]$, and an *empty set*.

When a row is being filtered, we check whether it is possible to pass the filter given the ranges of attributes that we have. We mark whether it *always* passes the filter, *never* passes the filter, or *both outcomes are possible*. For each attribute, we constrain the output range based on the filter, possibly leaving an overapproximation if it cannot be easily extracted. E.g., if a filter involves a combination of different attributes that are not separated by AND, e.g. $x + y \leq 5$, then this dependency is not included into the resulting constraints.

If the row contains any discrete attributes, the analysis is analogous. For each $s \in S$, we check whether s passes the filter or not. If s is compared with a constant, this check is trivial. If s is compared to another attribute $s' \in S'$, then we can set output ranges of both attributes to $S \cap S'$. Similarly to numeric attributes, dependency of different attributes is not included into the resulting constraints.

3.3.12 Differential Privacy for Time Series. So far, we have studied the case where a query is executed only once. Even if we have enough budget to execute more than one query, those queries are executed on the same database, which does not change in the meantime. Now we consider the case where the database changes over time and we execute the same query at several points in time. Thus we will get a time series. We would like to release the result of the query at each of the time points in a way that all the released values together are still differentially private w.r.t. all the inputs (over all time points).

A naive approach of obtaining differential privacy for time series is to simply split the privacy budget ϵ between the n time points, releasing the query result at each time point $\frac{\epsilon}{n}$ -differentially privately. This makes the noise n times larger than when executing the query only once. In practice, the input database usually does not change very fast over time, thus the query results at different time points should also be correlated. Then we can use an approach that is used outside the differential privacy world to smoothe the noisy time series of measurements obtained from sensors, namely *Kalman filters*, which we consider

in Sec. 3.3.12.1. We also obtain a new method for generating the differential privacy noise using smooth derivative sensitivity, where we get to know the exact noise level of the released value.

Note that in the current section, we often use the phrase “local/derivative sensitivity” because most of the discussed methods apply to both local sensitivity (where the distance over inputs is integer-valued) and derivative sensitivity (where the distance over inputs is real-valued). Even when we use the phrase “derivative sensitivity”, it is often applicable to local sensitivity as well.

Similarly to Sec. 3.3.9, because in this section we will use two different kinds of derivative sensitivity (w.r.t. row multiplicities and w.r.t. components), we call the derivative sensitivity w.r.t. components (i.e. in Banach spaces), *Banach sensitivity* or *Banach derivative sensitivity*.

3.3.12.1 Kalman Filters. Differential privacy can be achieved for time series using global sensitivity. This has been done in [66] using a Kalman filter. In this section, we discuss how to use Kalman filters with local or derivative sensitivity. We also obtain a new method for generating the differential privacy noise using smooth derivative sensitivity. It can be used with a slightly smaller budget than the old method and we get to know the exact noise level of the released value. If the budget is not very small and we do not need to learn the exact noise level, the old method still gives about e times better accuracy.

Kalman filters. Kalman filters can be used to take a series of noisy measurements at different time points (e.g. GPS data) and make it less noisy. To produce an estimate at the current time point, we use not only the noisy measurement at the current time point but also those at previous time points. If the underlying value does not change too fast then the previous measurements are additional estimates of the current value, although less precise. Thus we have several (at least partially) independent estimates of the current value, with different precision. Then we take their weighted average (with more precise estimates having higher weight) to get an estimate that is more precise than any of the initial estimates.

Kalman filters work in real time, i.e. the current estimate does not depend on future measurements. Also they can be implemented in constant memory, i.e. it is not necessary to save all the previous measurements. Also, execution time per measurement does not increase with the number of measurements.

Kalman filters for differential privacy. Kalman filters seem suitable for differential privacy because differentially private queries produce noisy estimates of the query result, similarly to noisy measurements from sensors. The difference is that measurements from sensors usually have Gaussian noise but differential privacy uses Laplace noise or generalized Cauchy noise (although Gaussian noise is also possible at the expense of higher noise levels and achieving only (ϵ, δ) -DP instead of ϵ -DP).

To use a Kalman filter, we need to estimate how much the underlying value has changed from the previous time point. This is assumed to be normally distributed. We also need to estimate the noise level of the measurement. When global sensitivity is used for differential privacy, this noise level is constant and public. With local/derivative sensitivity, we need to take extra steps to determine it, which we will discuss later.

The two sources of error must be independent, and also independent of these in previous time periods. The noise levels are used to compute the weights of different measurements in computing the more accurate estimate. If the noise level estimates are incorrect or the noises are not independent then it is not guaranteed anymore that the estimate is more accurate than the measurements. Differential privacy is still guaranteed because transformations do not weaken differential privacy.

Revealing the noise level for derivative sensitivity. When global sensitivity is used for differential privacy, this noise level is constant and public. With local/derivative sensitivity, we need to take extra steps to determine it.

One way is to use a part of the privacy budget to reveal the β -smooth derivative (or local) sensitivity differentially privately. β -smoothness guarantees that the logarithm of this sensitivity has global sensitivity at most β . Thus we can use global sensitivity to reveal this logarithm differentially privately. The

problem with this approach is that we only reveal an approximation of the noise level used for revealing the query result, not the exact noise level, thus the Kalman filter will be less accurate.

Another way is to first use a part of the privacy budget to reveal a noise level differentially privately (actually we reveal the logarithm of the noise level as in the previous approach, to avoid the chance of the noise level becoming negative after adding noise) and then generate Laplace noise with exactly this noise level to reveal the query result. Then we can use the exact noise level in the Kalman filter. The problem is how to find this appropriate noise level. If we reveal the logarithm of the actual β -smooth sensitivity then the added noise may be negative and the revealed noise level thus too low to achieve differential privacy for the query result.

Instead, we add another term to the logarithm, so that after adding noise, the revealed noise level will be with probability $1 - \delta$ enough to achieve ϵ -differential privacy for the query result. Since this δ does not depend on the input distribution, we will get (ϵ, δ) -differential privacy. The revealed noise level is

$$M \cdot e^{n \cdot k + \text{Lap}(k)}$$

where

- $n \cdot k$ is the added term to make $n \cdot k + \text{Lap}(k)$ non-negative with probability $1 - \delta$;
- $M = c/b$ is the computed noise level;
- c is the β -smooth sensitivity;
- $\delta = \frac{1}{2} \cdot e^{-n}$;
- $k = \beta/a$;
- $\epsilon = a + b$;

The privacy budget ϵ is split into a and b , where a is used for revealing the noise level and b is used for revealing the query result using this noise level.

The median noise level depends on how the budget is split. Let us find the value of a for which the median noise level is minimal. The logarithm of the median noise level is

$$\frac{n\beta}{a} - \ln(\epsilon - a) + \ln c$$

The minimum is achieved when the derivative w.r.t. a is 0, i.e.

$$-\frac{n\beta}{a^2} + \frac{1}{\epsilon - a} = 0$$

$$\epsilon - a = \frac{a^2}{n\beta}$$

$$a^2 + an\beta - \epsilon n\beta = 0$$

$$(a + n\beta/2)^2 = n^2\beta^2/4 + \epsilon n\beta$$

$$a = -n\beta/2 + \sqrt{n^2\beta^2/4 + \epsilon n\beta}$$

We can try this out in Haskell:

```
let f nb eps = let a = -nb/2 + sqrt(nb^2/4 + eps*nb); b = eps-a in (exp(nb/a)/b, a)
let g nb eps = (1/(eps-nb), nb)
let h eps nb = (f nb eps, g nb eps, fst(f nb eps)/fst(g nb eps))
Prelude> h 1 0.01
((1.2276273246518064,9.512492197250393e-2),(1.0101010101010102,1.0e-2),1.2153510514052883)
Prelude> h 1 0.1
((1.9839324650212786,0.27015621187164246),(1.1111111111111112,0.1),1.7855392185191508)
Prelude> h 1 0.4
((4.418001576289371,0.46332495807108004),(1.6666666666666667,0.4),2.650800945773623)
Prelude> h 1 0.5
```

```

((5.43656365691809,0.5),(2.0,0.5),2.718281828459045)
Prelude> h 1 0.6
((6.600169279739,0.5306623862918074),(2.5,0.6),2.6400677118956)
Prelude> h 1 0.9
((11.20422267584516,0.6000000000000001),(10.000000000000002,0.9),1.120422267584516)
Prelude> h 1 0.95
((12.170092683730568,0.6092624221100721),(19.99999999999982,0.95),0.6085046341865289)
Prelude> h 1 1
((13.203179065212284,0.6180339887498949),(Infinity,1.0),0.0)
Prelude> h 1 2
((57.34058739617453,0.7320508075688772),(-1.0,2.0),-57.34058739617453)
Prelude> h 1 3
((212.31841055007064,0.7912878474779199),(-0.5,3.0),-424.6368211001413)
Prelude> h 1 4
((728.6360040835475,0.8284271247461903),(-0.3333333333333333,4.0),-2185.9080122506425)
Prelude> h 1 10
((656034.405125172,0.9160797830996161),(-0.1111111111111111,10.0),-5904309.646126548)
Prelude> h 1 20
((2.7664644276074154e10,0.9544511501033224),(-5.263157894736842e-2,20.0),-5.2562824124540894e11)

```

We see that the overhead of revealing the noise level is largest when $n\beta = \epsilon/2$, then the noise increases e times if the total budget is not increased compared to the case where noise level is not revealed. If $n\beta$ approaches ϵ then the new method even starts to use less noise than the previous method, even though it releases more information. If $n\beta \geq \epsilon$ then the previous method would require infinite noise level and cannot be used at all. The new method can be used but as $\frac{n\beta}{\epsilon}$ exceeds 1, the noise level starts to increase exponentially in $\frac{n\beta}{\epsilon}$ and soon becomes impractical. Nevertheless, we may be able to use this method practically with a few times smaller budget than the old method. This is possible because the new method uses Laplace noise whose level is determined by an exponent of another Laplace noise. This combined noise distribution has a bit heavier tails than the ordinary Laplace distribution.

3.3.12.2 Local/Derivative vs Global Sensitivity in Kalman Filters. Because Kalman filters require splitting the budget between the released values and local/derivative sensitivity cannot be practically used with very small budgets per released value, we discuss why we need local/derivative sensitivity at all. We propose a modified version of global sensitivity that often has a reasonable magnitude when the actual global sensitivity is infinite or very large. This can be used in the simpler methods of achieving DP using global sensitivity. We will see that there are still cases where derivative sensitivity is needed.

We may question why we want to use local/derivative sensitivity instead of the more standard global sensitivity. The main reason is because global sensitivity may be infinite or very large. However, we can often use something similar to global sensitivity that has a more reasonable magnitude and can still be used with the simpler methods that achieve differential privacy using global sensitivity, with a slight loss in privacy. We try to find an example where even this modified global sensitivity is not good enough and we really have to use the more complicated methods of local/derivative sensitivity.

$(1 - \delta)$ -global sensitivity. If the actual global sensitivity is infinite or very large then we may instead use a value that is with probability at least $(1 - \delta)$ an upper bound on the local (or derivative) sensitivity. Then we get ϵ -differential privacy with probability at least $(1 - \delta)$. This is similar to but not the same as (ϵ, δ) -differential privacy because it depends on the actual input distribution but (ϵ, δ) -differential privacy holds for any input distribution.

Derivative sensitivity vs $(1 - \delta)$ -global sensitivity. If the actual global sensitivity is infinite or very large then we have a choice between using either derivative sensitivity or $(1 - \delta)$ -global sensitivity. With derivative sensitivity we get classical (ϵ, δ) -differential privacy, with $(1 - \delta)$ -global sensitivity we get something similar but not the same.

To use $(1 - \delta)$ -global sensitivity, we need to know something about the input distribution. If we do not know the exact input distribution then at least we need to find some values δ and c such that with probability at least $1 - \delta$, the input is such that the derivative sensitivity at that input point is at most c .

To use derivative sensitivity, we do not need to know anything about the input distribution. Instead, we need to use a value of β and find a β -smooth upper bound on the derivative sensitivity. The noise level depends on β and on the actual input, choosing the best β can be difficult.

Also, derivative sensitivity requires a higher noise level for the same sensitivity than global or $(1 - \delta)$ -global sensitivity. It is about 2 times higher if we do not need to reveal the noise level and $2e$ times higher if we do. Thus we need the median (over an input distribution) β -smooth derivative sensitivity to be at least $2e$ times smaller than its $(1 - \delta)$ -quantile for derivative sensitivity to be advantageous over $(1 - \delta)$ -global sensitivity.

Thus the derivative sensitivity must vary over a quite large range over the possible inputs. If the query uses only one table and uses it only once and the norm used for combining row norms is ℓ_1 then there are some inputs from which the large range of sensitivity can be achieved by changing only one row. Thus β is high and derivative sensitivity cannot be practically used with small budgets.

So let us use the table twice:

```
SELECT count(*) FROM t as t1, t as t2 WHERE t1.groupId = t2.groupId
```

Let each row belong to a different individual and each individual belong to one group, denoted by the field `groupId` in the row. If an individual of row i belongs to a group with m members then the derivative sensitivity w.r.t. row i is $2m$ and is $\frac{1}{m}$ -smooth at this point. The total sensitivity of the query is determined by the largest groups thus we can make β quite small, as long as the largest groups are big enough with high probability.

With a small β , we can use a Kalman filter with derivative sensitivity (using the method from Sec. 3.3.12.1 that also reveals the noise level). The smaller the β , the larger the number of time points we can reveal before the budget is exhausted. If we need to reveal an even longer time series (e.g. an infinite time series) then the budget needs to regenerate over time. If β is smaller then it does not need to regenerate as fast as for a larger β .

3.3.12.3 An Alternative to Kalman Filter. Instead of releasing the query result at each time point, we may release the *change* in the query result from the previous time point (in addition to the query result at the first time point of the series). The changes in non-overlapping time periods are likely to be less correlated than the query results at the end of these time periods. Thus each released change does not leak too much about other released changes, which allows reusing some of the budget.

In this section, we assume that the changes are completely independent, thus the whole budget can be reused at each time point. There is still some overhead over the single execution because we are interested in the query results at different time points and each of these must be computed from several released change values. This overhead can be made polylogarithmic, which is much better than the $\Omega(\sqrt{n})$ overhead of the Kalman filter. The disadvantage is that this approach assumes that the changes in non-overlapping time periods are independent (i.e. the sets of input rows that they depend on are disjoint, or if we need user-level privacy instead of event-level, then the sets of users on whose inputs they depend on must be disjoint). Otherwise the privacy is not guaranteed, although accuracy is. On the other hand, with Kalman filters, privacy is guaranteed but accuracy depends on the input distribution. As privacy is usually more important than accuracy, we can use the Kalman filter in the general case but there is a class of queries (considered in the following sections) for which our approach works as well.

We use an algorithm similar to one in [67]. The difference is that they release the results of the query at each time point while we release the changes in the query result over certain time periods. The noises they add to each released value, are, however, not independent and are generated in roughly the same way as in our algorithm, which allows keeping the overhead polylogarithmic is both their and our algorithm. Their algorithm also achieves *pan-privacy*, a stronger version of differential privacy where the attacker can not only see the released outputs but can also look at the internal state of the mechanism

at one point in time that the attacker chooses. Our algorithm can also be modified to achieve pan-privacy, at the expense of increasing the noise level $\sqrt{2}$ times.

Suppose that changes in the query result during non-overlapping time periods are independent and normally distributed. This is also assumed for Kalman filter. Kalman filter reveals the whole query result at each time point. Those values are not independent and thus we have to split the privacy budget between them, leaving a very small budget for each time point.

Suppose that we instead reveal the change in the query result between every two successive time points, in addition to the query result at the first time point. Those revealed values are mutually independent. Then it is possible that making a same change of same magnitude in the change of query result at many time points is not more important than making it in only one time point, i.e. we could use the ℓ_∞ norm instead of the ℓ_1 norm. Then we can reuse the same budget for each time period that does not overlap with the time periods for which budget has already been used. Note that the ℓ_∞ may not always be allowed and then we still need to use the Kalman filter.

Suppose that it is allowed and we use the simplest method of using the whole budget ϵ to reveal the change in query result between each two successive time points. Let those changes be x_1, x_2, \dots, x_n and let the query result at the first time point be r_0 . Then we can reconstruct the time series as $r_i = r_0 + \sum_{j=1}^i x_j$ where r_i is the query result at time point i . The noise level of the elements will be $O(\sqrt{n})$.

We can do better. Instead of using the whole budget ϵ to reveal each x_i , we use only ϵ/s for this (where $s = \log_2 n$), then use ϵ/s to reveal each $x_{2i-1} + x_{2i}$, the same amount to reveal each $x_{4i-3} + x_{4i-2} + x_{4i-1} + x_{4i}$, and so on. I.e. for each $k = 0, 1, \dots, s-1$, we use ϵ/s to reveal each $\sum_{j=1}^{2^k} x_{2^k(i-1)+j}$. Then each r_i can be reconstructed as a sum of r_0 and m revealed values corresponding to changes in periods whose lengths are different powers of two, and m is the number of bits that are 1 in the binary representation of i . The noise level of the r_i will be $O(\sqrt{\log_2 n})$.

If we use derivative sensitivity then using only $\epsilon/\log_2 n$ for each query may be a problem because the noise level starts to increase very fast once the budget goes below a certain value. The old method of using derivative sensitivity can even not be used at all. Thus we generalize the current method to use a -ary representation instead of binary. Then we need $\epsilon/\log_a n$ budget for each query while the noise level increases to $O(\sqrt{(a-1)\log_a n})$.

If the output is a linear function of input and we use the ordinary absolute value norm for each row (this norm is separate from the norm used to combine rows) then the derivative sensitivity is constant and has no advantage over global sensitivity. To make it have an advantage, we may instead use the square root norm (changing the square root of an input row by 1 corresponds to distance 1).

So far we assumed that changes in the query result during non-overlapping time periods are independent and normally distributed. We can actually remove this restriction if we are satisfied with event-level privacy instead of user-level privacy. Then each row corresponds to an event. In the following sections, we show how to achieve differential privacy on event level.

3.3.12.4 Sensitivity w.r.t. Changes Inside Rows. In this section, we show how DP can be achieved using derivative sensitivity w.r.t. component (Sec. 3.3.8.2), i.e. *Banach sensitivity*. Note that, while Banach sensitivity requires that the number of rows in each table is fixed, we still allow to add and remove rows from timepoint to timepoint, but we only treat the content of the rows as private, and do not attempt to conceal the fact that a row has been added or removed. We will discuss how to do the latter in Sec. 3.3.12.5.

Queries with a single table. We start from a simple case where the query applies a function on each event and sums the results of the functions from all events. The result contributed by each event must not depend on other events. Thus the class of queries that we can compute is the SQL SUM queries on a single table without self-joins:

SELECT SUM($f(a_1, \dots, a_n)$) FROM t

Approved for Public Release; Distribution Unlimited.

We can also allow filters because these can be embedded in the function f (multiplying the original f by a function that returns 1 for rows satisfying the filter and 0 for other rows):

$$\text{SELECT SUM}(f(a_1, \dots, a_n)) \text{ FROM } t \text{ WHERE } p(a_1, \dots, a_n)$$

where a_i are the attributes (columns) of table t . In this section, we will further assume that any filter is already included in the function f .

At each time point, the query is evaluated on those rows that are already in the table, i.e. those events that have already occurred. Over time, rows can only be added to the table, never removed or changed.

Note that the times when rows arrive are not private, only the content of the rows is. This is not a problem if the rows arrive at regular intervals, but if the intervals are irregular then they are leaked.

Let us first compute the sensitivity w.r.t. a row (a_1, \dots, a_s) added at time point i . We compute a β -smooth upper bound on $\text{DS}_f(a_1, \dots, a_s)$. Let it be c_i . We also compute the change at time point i , which is $x_i = f(a_1, \dots, a_s)$. Let the total number of time points be n . Then we can use $\epsilon_1 = \frac{\epsilon}{\log_2 n}$ for revealing x_i . The noise level L_i will then be computed according to the values of ϵ_1 , c_i , β , and the noise distribution that we want to use (generalized Cauchy with parameter γ achieving ϵ -DP, Laplace with $\delta_1 = \frac{\delta}{e^{\epsilon \log_2 n}}$ achieving (ϵ, δ) -DP). Then we generate noise N_i with level L_i , add it to x_i , and reveal the sum $x_i + N_i$.

In addition to revealing the noisy versions of single changes x_i , we also need to reveal some sums of changes in consecutive time periods, e.g. $x_{2i-1} + x_{2i}$ or $x_{4i-3} + x_{4i-2} + x_{4i-1} + x_{4i}$.

For all k such that i is divisible by 2^k , we need to reveal the noisy version of

$$\sum_{j=1}^{2^k} x_{i-2^k+j}$$

The noise level that we have to use, is

$$\max_{j=1}^{2^k} L_{i-2^k+j}$$

because it has to be enough to hide all the x_{i-2^k+j} . The value of x_i affects at most $\log_2 n$ revealed values. The noise level L_i potentially affects (if it happens to be the maximum) the same $\log_2 n$ revealed values. Changing row i by distance d changes each revealed value affected by x_i , by DP-distance at most $\frac{\epsilon}{\log_2 n} \cdot d$, thus the whole output time series is changed by DP-distance at most $\epsilon \cdot d$. Thus we get ϵ -DP.

Using public tables. Suppose now that in addition to the single table t , the query can use some other tables t_1, \dots, t_m but these tables are all public and constant.

$$\text{SELECT SUM}(f(a_1, \dots, a_s)) \text{ FROM } t, t_1, \dots, t_m \text{ WHERE } p(a_1, \dots, a_s)$$

Let the columns a_1, \dots, a_s be from table t and a_{s+1}, \dots, a_S from the public tables.

Let us compute the change x_i in the query result when a row (a_1, \dots, a_s) is added to table t at time point i . The row (a_1, \dots, a_s) is joined with rows from the public tables getting (0 or more) rows of the form $(a_1, \dots, a_s, \dots, a_S)$ where the a_1, \dots, a_s are the same for all obtained rows but a_{s+1}, \dots, a_S may be different. We apply the filter p so that only rows satisfying p remain. Then we apply f to each row and sum the results. The obtained sum x_i is uniquely determined by a_1, \dots, a_s , i.e. there exists a function g that computes this sum $g(a_1, \dots, a_s)$. Thus the query can still be expressed in the form

$$\text{SELECT SUM}(g(a_1, \dots, a_s)) \text{ FROM } t$$

using only the table t .

It may be inconvenient to express the function g in SQL. Thus we may still compute it as described above. Let us now compute the sensitivity w.r.t. the row i . We still take the 0 or more rows of the form $(a_1, \dots, a_s, \dots, a_S)$ where the a_1, \dots, a_s are the same for all of them. We apply the filter p , which we assume to use only public columns, private filters should be included in the function f . For each row $(a_1, \dots, a_s, \dots, a_S)$, we compute the sensitivity of $f(a_1, \dots, a_s, \dots, a_S)$ w.r.t. (a_1, \dots, a_s) . Then we sum the sensitivities and obtain c_i , the sensitivity w.r.t. row i . Once we have the c_i (and x_i) for all i , we proceed as in the single table case.

Using private tables with joins by user ID. Suppose now that in addition to the single table t , the query can use some other *private* tables t_1, \dots, t_m (these tables may be used more than once, unlike table t) but these must be joined to the table t using the user ID column which should be included for all private tables and denotes the provenance of the row. Also, the user ID of each row added to table t must be different. Rows are still added only to table t , other private tables are constant (actually we can change the rows belonging to those users who are not in table t yet, as these do not affect the query result; as soon as a row from that user is added to table t , the rows of that user cannot change anymore).

Then we can join the row added at time i to other tables and obtain 0 or more rows, all of which contain private data only from the same provenance as the added row. We can compute the sensitivity of the function f applied to each obtained row w.r.t. each used row in the original tables (note that a row in the original tables is not necessarily used in all obtained rows; when it is not used the respective sensitivity is 0). Then sum the sensitivities over obtained rows to get the total sensitivity of x_i w.r.t. each used row in the original tables (if that used row is from a table used more than once then the sensitivities w.r.t. each copy of the row must be summed). Then these sensitivities are combined according to the norm of the Banach space (the Banach space is over all of the private tables, the norms of individual tables may be combined e.g. by ℓ_1 -norm) to obtain c_i .

Because the user ID of each row added to table t is different and the rows joined to it have the same user ID, the budget does not have to be split between different added rows as each row uses data from a different provenance.

Adding more than one row from the same user. So far at most one row was added from each user into table t . Now consider the case where more than one row is added from the same user.

If those rows are added at the same time then we can compute their sensitivities w.r.t. each input row and combine them according to the norm of the Banach space. The noise is added to the total change in query result from adding all the rows, not to the change from each added row separately.

If the rows are not added at the same time, then the noise must be added to each change separately and we need to somehow split the budget between the changes. If we know that the user adds rows to table t at most m times then we can use $1/m$ part of the initial budget for each time, making noise m times higher.

If there is no upper bound on the number of times but we estimate it to be at most m with high probability then we can use $1/m$ part of the initial budget for each time and after the m^{th} time we can exclude any further rows added from that user.

Personalized differential privacy. So far, rows of different users could not be joined because then a row could affect the change in query result not only at the time when it is added but also at some later time points. Its privacy budget would be used at all these time points. To keep track of the budgets of all rows, we use *personalized differential privacy*. Then, at each time point or period we find the rows that are joined to some row added at that time, and update their budgets.

In addition to joining rows of different users, we can now also add rows to multiple tables and use those tables more than once. We can now also add rows to public tables, which also consumes some of the budget of the private rows added earlier to which the public rows are joined.

To allow a user to own more than one row (possibly in different tables) using a single budget, we add provenances to the personalized differential privacy, i.e. each row has a provenance and each provenance has a budget.

Removing and changing rows. So far rows could only be added. Personalized differential privacy also allows removing and changing them. Each row can now have addition time and removal time (the latter must be later than the former, otherwise the row will never be added). Query result will be changed and budget will be consumed at both of those times. If we want a row not to be removed at all then we may set the removal time very large (larger than the end of the time period for which we want differentially private updates of the query result).

To change a row, we will have a separate row for each version of the row, one will be removed and another added at the same time point, and both will have the same provenance. Because the addition and removal occur at the same time point, the budget will be consumed only once.

When revealing query result changes in longer time periods (of power-of-two length), we can ignore those rows (or row versions) that are added after the period starts but removed again before the period ends. Those rows will not affect the revealed value or its noise level and their budget is not consumed. Similarly, if a row is changed multiple times during the period, the budget will still be consumed only once.

3.3.12.5 Sensitivity w.r.t. Changes Inside Rows and Row Multiplicities. Banach analysis assumes that the join graph (which rows are joined with which) is public, only the content of the rows is private. This is the case both for single-query and for time-series analysis. For the time-series analysis, we also have times when rows are added/removed/changed, and these are also considered public. For the single-query analysis, we used *combined sensitivity* to make the join graph private.

To adapt combined sensitivity to the time-series analysis, we apply the combined sensitivity analysis separately for each time period (of power-of-two length) for which query change is revealed. This makes both the join graph and the times of changes private. But the personalized privacy budgets also depend on the join graph and change times, thus the budgets must become private as well. So we cannot reveal the amount of budget used anymore without adding noise to it and using more budget.

Leaks through the noise level. Unlike non-combined Banach sensitivity, with combined sensitivity it is possible that the noise level of a revealed value depends on a certain row but the value itself (before adding noise) does not. Budget must be used also in those cases, making the amount of budget needed much higher than for non-combined sensitivity. If the row r is in the database in the time period $[t_1, t_2]$ then at any time point $t \in [t_1, t_2]$, the sensitivity w.r.t. adding to another table a row r' , depends on whether r can be joined with r' , i.e. on the join key of r . Thus the budget of r is used at every time point in $[t_1, t_2]$.

The total amount of budget needed for a row r is approximately $(s + m \log_2 n)\epsilon$ where s is the total number of time points when (a version of) r is in the table, m is the number of versions of r (if the same version is removed and later added again then it is considered a new version), n is the total number of analyzed time points, ϵ is the budget used for a single query at a time point.

The noise level is $\log_2 n$ times higher than the naive method (which computes the query result separately at each time point, using $n\epsilon$ budget in total) because the query result at each time point is computed as a sum of up to $\log_2 n$ values. Thus we may actually use only $n\epsilon / \log_2 n$ budget for the naive method to get the same noise level.

Using global sensitivity. To avoid noise levels depending on input rows, we can use global sensitivity instead of local. We combine global Banach sensitivity and global adding/removing rows sensitivity. The latter requires knowing in advance how many times each row can be used in total over all time points. This may be possible for real-time analysis if there are some restrictions on the input distribution, or for batch analysis, i.e. we are releasing the time series only after all the time points have passed.

With global sensitivity, we can use the same noise level for all revealed values as for the single query of the non-time-series analysis. Then the DP-distance corresponding to a change in input is at most $\log_2 n$ times higher than for the single query. The rows of the joined table are distributed between different time points but they are the same as for the single query. The sum of the changes in the query result over the time points is equal to the result of the single query. The DP-distance depends on the sum of absolute values of query changes. Let L be the maximum possible absolute value contribution of a row (of the joined table) to the query result. Then the total sensitivity (across all time points) w.r.t. adding/removing (all uses of) a row would be $ML \log_2 n$, where M is the maximum number of times a row can be used (if this is exceeded then DP is not guaranteed).

Thus we are now using a fixed amount ϵ_1 of budget per row use, not per time point. If a row r is joined with m other rows at a certain time point t then r uses $m\epsilon_1$ budget at time point t . If the global Banach sensitivity w.r.t. a row use is K then w.r.t. all uses of a row it is $MK \log_2 n$. Combined global sensitivity is then

$$C = \max\left(\frac{ML \log_2 n}{G}, MK \log_2 n\right) = M \log_2 n \cdot \max\left(\frac{L}{G}, K\right)$$

where G is the distance corresponding to adding/removing a row. The noise level for each released value is

$$\frac{M \log_2 n \cdot \max\left(\frac{L}{G}, K\right)}{\epsilon}$$

where ϵ is the total budget. As L and K are computed in the Banach analyzer and M, G, n, ϵ are also available there, the local sensitivity analyzer is not needed for this kind of combined sensitivity analysis.

We may also compute a separate global ϵ_1 sensitivity for each table t (considering only changes inside table t):

$$C_t = M \log_2 n \cdot \max\left(\frac{L}{G_t}, K_t\right)$$

where G_t is the distance corresponding to adding/removing a row in table t and K_t is the global Banach sensitivity w.r.t. a row use in table t . The noise level must be the same as for a single global sensitivity. The amount of budget needed for each row use in table t in a time period, is then

$$\frac{C_t}{C} \cdot \frac{\epsilon}{M \log_2 n}$$

Combining with local Banach sensitivity. It seems that we can also combine the global adding/removing rows sensitivity with local Banach sensitivity. So far, in the Banach analyzer, the local Banach sensitivity was computed for all uses of a row combined. To achieve this, for each row r in an input table (or table copy for multiply used tables), the rows in the joined table that depend on r (let the number of such rows be m), have their norms combined using the ℓ_∞ norm instead of the original norm, because making a change with distance d in row r corresponds to making the change (with the same distance d) in m rows of the joined table. To make those m changes of distance d have the total distance equal to the original distance d , we must use the ℓ_∞ norm to join them.

Now we need to compute the local Banach sensitivity per row use, i.e. for a single row in the joined table, as is already done for the global Banach sensitivity K . Here those m changes of distance d are allowed to have the total distance up to md because the allocated budget is multiplied by the number of times a row is used, i.e. m . Thus we use the ℓ_1 norm to combine the m rows in the joined table that depend on r .

For smoothing the sensitivity, we must take into account that the local Banach sensitivity can jump to the global one with an addition/removal of a single row. Thus the β -smooth Banach sensitivity κ w.r.t. a single row use can vary between $e^{-\beta G} K$ and K . The local Banach sensitivity can be computed for each time period separately. Because only M uses of a row can contribute to the query result, the amount of budget allowed for each time period is $\epsilon_0 = \frac{\epsilon}{M \log_2 n}$ where ϵ is the total budget. This amount is actually available for each use of a row in a time period, instead of all uses (similarly to the L above being for a row in the joined table, which corresponds to a single use of a row).

ϵ_0 also limits the possible values of β . Let $\epsilon = \epsilon_b + \epsilon_\beta$, where ϵ_b is the part of ϵ used to hide the query result and ϵ_β is the part of ϵ used to hide the noise level. Then

$$\beta = \frac{\epsilon_\beta}{(\gamma + 1)M \log_2 n}$$

$$b = \frac{\epsilon_b}{(\gamma + 1)M \log_2 n}$$

Thus β will usually be very small, making it difficult to get an advantage over global sensitivity, unless G is very high. Adding/removing all uses of a row can cause the noise level of up to $M \log_2 n$ released values to change. Thus the β -smooth Banach sensitivity κ w.r.t. a single row use can actually only be allowed to vary between $e^{-\beta G} K$ and K .

Thus the noise level for a released value for a time period with smooth Banach sensitivity $\kappa \in [e^{-\beta G} K, K]$, is

$$\frac{\max\left(\frac{L}{G}, \kappa\right)}{b} = \frac{(\gamma + 1)M \log_2 n \cdot \max\left(\frac{L}{G}, \kappa\right)}{\epsilon_b}$$

which can be up to $(\epsilon/\epsilon_b)(\gamma + 1)$ times higher than with global Banach sensitivity. If κ may be smaller than $e^{-\beta G} K$ (it cannot be larger than K) then the noise level would be

$$\frac{(\gamma + 1)M \log_2 n \cdot \max\left\{\frac{L}{G}, \kappa, e^{-\beta G} K\right\}}{\epsilon_b}$$

Compared to the combined sensitivity without time series, the noise level is about $\log_2 n$ times higher (this is for released values; for actual time series elements, it is $(\log_2 n)^{1.5}$ times) but it can be even more because the row that has the maximum κ may be used less than M times and the non-time-series noise level will then be less than $M\kappa/\epsilon_1$. Even larger increase in the noise level can be caused by the term

$$e^{-\beta G} K = e^{-\frac{\epsilon_\beta G}{(\gamma+1)M \log_2 n}} K$$

which would be

$$e^{-\frac{\epsilon_\beta G}{\gamma+1}}$$

without time series. Thus G would have to be $M \log_2 n$ times higher than for non-time-series analysis to be able to get the same advantage over global sensitivity. It can also be difficult to guess a good value of M , and a non-optimal M can further reduce accuracy.

If a table is used more than once then the G above does not need to be divided by the number of times the table is used.

Adding or removing a row r can change the noise level of those time periods where r is used, from $e^{-\beta G} K$ to K , a factor of $e^{\beta G}$. This changes the distribution of each affected released value by DP-distance $(\gamma + 1)\beta G$. The number of released values affected by adding or removing r , is $M \log_2 n$. Thus the total DP-distance by which the output distribution changes, is

$$(\gamma + 1)\beta G \cdot M \log_2 n = \epsilon_\beta G$$

We have implemented the method for both global and local Banach sensitivity. It only works correctly if the upper bound M cannot be exceeded. If it is exceeded then DP is not guaranteed.

3.3.12.6 Budget Renewal. If we want to release more than one value and want all the released values together to be B -differentially private then our privacy budget is B . A simple budget management algorithm would start from budget B and reduce it by the amount of budget used each time a new value is released, and checking that the budget would not go below zero. If each released value does not depend on all input rows or its sensitivity w.r.t. different input rows is different then we may benefit from more sophisticated budget management techniques, two of which are discussed in this section.

We discuss how to take advantage of the independence of changes even when we release query results and not changes. This approach can be used improve the accuracy of the Kalman filter. However, as in Sec. 3.3.12.3, we lose the privacy guarantee if the changes are not actually independent. The approach takes the independence into account then tracking the privacy budget. The part of budget used for releasing the parts of the values released in the past that are independent of values released in the future, is returned to the user and can be reused in the future. This only ensures differential privacy of each value at the moment when it is released. Values released in the future may weaken the privacy. To

ensure that privacy continues into future, we must in addition to tracking the amount of *used budget*, track the amount of *remaining budget*. The part of budget that will be used for releasing the parts of the values released in the future that are independent of values released in the past, is added to the remaining budget, but not above the initial budget. Every time the budget is used, the remaining budget must not go below zero and used budget must not go above the initial budget. This ensures differential privacy over the whole time period. The privacy budget seems to renew over time and this is due to the independence in the changes, not because the data subjects agree to give up some privacy over time.

For long time series, the amount of privacy budget needed to release them differentially privately with a given accuracy (noise level) grows by the number of time points n , either polylogarithmically (as in Sec. 3.3.12.3) or faster ($\Theta(\sqrt{n})$ to $\Theta(n)$ for the Kalman filter). In this section, we try to find some motivation for why the budget could renew.

Two-element time series. Let

$$X, Y \sim N(0, 1) \text{ and independent}$$

$$Z = \frac{X + Y}{\sqrt{2}}$$

Then

$$Z \sim N(0, 1)$$

The random values X and Z can be viewed as two successive elements of a time series. The distribution of Z depends only on the value of X . Thus the time series is a Markov chain. Let

$$W = X\sqrt{2} - Z$$

Then

$$X = \frac{Z + W}{\sqrt{2}}$$

$$W = \frac{X - Y}{\sqrt{2}}$$

$$W \sim N(0, 1)$$

Lemma 3.102. W is independent of Z .

Proof.

$$f_X(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$$

$$f_Y(y) = \frac{1}{\sqrt{2\pi}} e^{-y^2/2}$$

$$f_{X+Y}(u) = \int_{-\infty}^{\infty} f_X(x)f_Y(u-x)dx = \int_{-\infty}^{\infty} \frac{1}{2\pi} e^{-x^2/2} e^{-(u-x)^2/2} dx$$

$$= \int_{-\infty}^{\infty} \frac{1}{2\pi} e^{-(x^2+u^2-2ux+x^2)/2} dx$$

$$= e^{-u^2/4} \int_{-\infty}^{\infty} \frac{1}{2\pi} e^{-(u^2/2-2ux+2x^2)/2} dx$$

$$= e^{-u^2/4} \int_{-\infty}^{\infty} \frac{1}{2\pi} e^{-(u^2/2-2ux+2x^2)/2} dx$$

$$= e^{-u^2/4} \int_{-\infty}^{\infty} \frac{1}{2\pi} e^{-(u^2/4-ux+x^2)} dx$$

$$\begin{aligned}
&= e^{-u^2/4} \int_{-\infty}^{\infty} \frac{1}{2\pi} e^{-(x-\frac{u}{2})^2} dx \\
&= e^{-u^2/4} \int_{-\infty}^{\infty} \frac{1}{2\pi} e^{-\left(\frac{x-\frac{u}{2}}{1/\sqrt{2}}\right)^2/2} dx \\
&= \frac{1}{\sqrt{2\pi}} e^{-u^2/4} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\left(\frac{x-\frac{u}{2}}{1/\sqrt{2}}\right)^2/2} dx \\
&= \frac{1}{\sqrt{2} \cdot \sqrt{2\pi}} e^{-u^2/4} \int_{-\infty}^{\infty} \frac{1}{(1/\sqrt{2})\sqrt{2\pi}} e^{-\left(\frac{x-\frac{u}{2}}{1/\sqrt{2}}\right)^2/2} dx \\
&= \frac{1}{\sqrt{2} \cdot \sqrt{2\pi}} e^{-u^2/4} dx \\
&= \frac{1}{\sqrt{2} \cdot \sqrt{2\pi}} e^{-(u/\sqrt{2})^2/2} dx
\end{aligned}$$

$$f_{X+Y, X-Y}(u, v) du dv = f_X\left(\frac{u+v}{2}\right) f_Y\left(\frac{u-v}{2}\right) dx dy$$

$$du dv = \frac{1}{2} dx dy$$

$$\begin{aligned}
f_{X+Y, X-Y}(u, v) &= \frac{1}{2} f_X\left(\frac{u+v}{2}\right) f_Y\left(\frac{u-v}{2}\right) \\
&= \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}} e^{-(\frac{u+v}{2})^2/2} \cdot \frac{1}{\sqrt{2\pi}} e^{-(\frac{u-v}{2})^2/2} \\
&= \frac{1}{4\pi} e^{-(\frac{u+v}{2})^2/2 - (\frac{u-v}{2})^2/2} \\
&= \frac{1}{4\pi} e^{-(u+v)^2/8 - (u-v)^2/8} \\
&= \frac{1}{4\pi} e^{-(2u^2+2v^2)/8} \\
&= \frac{1}{4\pi} e^{-(u^2+v^2)/4} \\
&= \frac{1}{\sqrt{2} \cdot \sqrt{2\pi}} e^{-(u/\sqrt{2})^2/2} \cdot \frac{1}{\sqrt{2} \cdot \sqrt{2\pi}} e^{-(v/\sqrt{2})^2/2} \\
&= f_{X+Y}(u) f_{X-Y}(v)
\end{aligned}$$

This proves that $X + Y$ and $X - Y$ are independent. Thus also Z and W are independent. □

Thus the reverse of the time series is also a Markov chain with the same distribution.

n -element time series. We can now describe the whole time series X_1, \dots, X_n :

$$X_1 \sim N(0, 1)$$

$$Y_i \sim N(0, 1) \text{ for all } i = 1, \dots, n$$

X_1 and all Y_i are mutually independent.

$$X_{i+1} = \frac{X_i + Y_i}{\sqrt{2}}$$

Then we have

$$X_i \sim N(0, 1)$$

From the results above we also have

$$X_n \sim N(0, 1)$$

$$W_i \sim N(0, 1) \text{ for all } i = 1, \dots, n$$

X_n and all W_i are mutually independent.

$$X_{i-1} = \frac{X_i + W_i}{\sqrt{2}}$$

Differential privacy. Suppose we add Laplace noise with level L_i to each X_i and reveal the noisy result:

$$N_i \sim Lap(L_i)$$

$$R_i = X_i + N_i$$

Then R_i is $\frac{1}{L_i}$ -differentially private w.r.t. X_i . Let $\epsilon_i = \frac{1}{L_i}$. Then we may say that the amount of budget used at time point i is ϵ_i .

However, the X_i are not independent and thus the other R_j may also leak information about X_i . Suppose that $j > i$. Then

$$\begin{aligned} R_j &= X_j + N_j = X_{j-1}(\sqrt{2})^{-1} + Y_{j-1}(\sqrt{2})^{-1} + N_j \\ &= X_{j-2}(\sqrt{2})^{-2} + Y_{j-2}(\sqrt{2})^{-2} + Y_{j-1}(\sqrt{2})^{-1} + N_j = \dots \\ &= X_{j-k}(\sqrt{2})^{-k} + Y_{j-k}(\sqrt{2})^{-k} + \dots + Y_{j-m}(\sqrt{2})^{-m} + \dots + Y_{j-1}(\sqrt{2})^{-1} + N_j = \dots \\ &= X_i(\sqrt{2})^{i-j} + N_j + \sum_{m=1}^{j-i} Y_{j-m}(\sqrt{2})^{-m} \end{aligned}$$

Note that $X_i(\sqrt{2})^{i-j}$ is $(\sqrt{2})^{i-j}$ -sensitive w.r.t. X_i , thus $X_i(\sqrt{2})^{i-j} + N_j$ is $\frac{(\sqrt{2})^{i-j}}{L_j}$ -differentially private w.r.t. X_i . Because $\sum_{m=1}^{j-i} Y_{j-m}(\sqrt{2})^{-m}$ is independent of X_i and N_j , adding it does not weaken differential privacy, thus R_j is also $\frac{(\sqrt{2})^{i-j}}{L_j}$ -differentially private w.r.t. X_i .

If $j < i$ then we get

$$R_j = X_i(\sqrt{2})^{j-i} + N_j + \sum_{m=1}^{i-j} W_{j+m}(\sqrt{2})^{-m}$$

and R_j is $\frac{(\sqrt{2})^{j-i}}{L_j}$ -differentially private w.r.t. X_i .

Thus for all j , R_j is $\frac{(\sqrt{2})^{-|i-j|}}{L_j}$ -differentially private w.r.t. X_i . We may also express it as $(\sqrt{2})^{-|i-j|}\epsilon_j$ -differentially private w.r.t. X_i . The whole noisy time series (R_1, \dots, R_n) is $\sum_{j=1}^n (\sqrt{2})^{-|i-j|}\epsilon_j$ -differentially private w.r.t. X_i .

Thus the total amount of budget used is not $\sum_{j=1}^n \epsilon_j$ but only

$$B = \sum_{j=1}^n (\sqrt{2})^{-|i-j|} \epsilon_j$$

We may say that the budget *renews* to some extent. If $j < i$ then the budget used at time point j contributes $\sqrt{2}$ times less to the total used budget than the same amount of budget used at time point $j + 1$. The total amount of budget used by time j can be defined as

$$E_j = \sum_{k=1}^j (\sqrt{2})^{k-j} \epsilon_k$$

Then $E_{j+1} = \frac{E_j}{\sqrt{2}} + \epsilon_{j+1}$. Thus the amount of budget used is divided by $\sqrt{2}$ between time points j and $j + 1$, which can be viewed as budget renewal. This renewal works until time point i , when E_i becomes equal to the sum of the first i terms of B .

Let us now consider what happens after time point i . If $j > i$ then the budget used at time point j contributes $\sqrt{2}$ times less to the total used budget than the same amount of budget used at time point $j - 1$. The total amount of budget used from time j on can be defined as

$$F_j = \sum_{k=j}^n (\sqrt{2})^{j-k} \epsilon_k$$

This amount of budget must be available at time point j before using the ϵ_j at that time point. Then $F_j = (F_{j-1} - \epsilon_{j-1}) \sqrt{2}$. Thus the amount of budget available for future use is multiplied by $\sqrt{2}$ between time points $j - 1$ and j . This is also a kind of budget renewal but slightly different from the one that occurs before time point i . At time point i , F_i is equal to the sum of the last $n - i + 1$ terms of B .

Altogether,

$$B = E_i + F_i - \epsilon_i$$

The ϵ_i is subtracted because it is included in both E_i and F_i but only once in B . Thus we can use the described kinds of budget renewal using the following budget tracking algorithm:

- Let $E_0 = 0$.
- For each time point $j = 1, \dots, i - 1$:
 - Let ϵ_j be the amount of budget used at time point j .
 - Let $E_j = \frac{E_{j-1}}{\sqrt{2}} + \epsilon_j$.
 - If $E_j > (\sqrt{2})^{i-j} B$ then **fail**.
 - Reveal R_j .
- For time point $j = i$:
 - Let ϵ_j be the amount of budget used at time point j .
 - Let $E_j = \frac{E_{j-1}}{\sqrt{2}} + \epsilon_j$.
 - If $E_j > B$ then **fail**.
 - Let $F_{j+1} = (B - E_j) \sqrt{2}$.
 - Reveal R_j .
- For each time point $j = i + 1, \dots, n$:
 - Let ϵ_j be the amount of budget used at time point j .
 - Let $F_{j+1} = (F_j - \epsilon_j) \sqrt{2}$.
 - If $F_{j+1} < 0$ then **fail**.
 - Reveal R_j .

If this algorithm does not **fail** then the revealed values together are B -differentially private w.r.t. X_i . If it does **fail** then the revealed values before the **fail** together are B -differentially private w.r.t. X_i .

Differential privacy w.r.t. all X_i . The above algorithm only guarantees differential privacy w.r.t. X_i for a single i . Suppose we want that for all i , the revealed values together are B -differentially private w.r.t. X_i . One way to achieve this is to run n instances of the above algorithm in parallel, one instance (thread i) for each i to achieve differential privacy w.r.t. X_i . The threads (instances) are synchronized at each "Reveal R_j " statement, which is executed only when none of the threads has **failed** by that point.

Let us try to combine this into a sequential algorithm. Note that for all j , ϵ_j is the same in all threads. Thus also E_j is the same in all threads where it is defined. The check $E_j > (\sqrt{2})^{i-j}B$ is performed in threads $i \geq j$. To check if at least one of the threads fails here, we may check $E_j > B$.

Now consider F_j , which is not necessarily the same in all threads. It is, however, updated by the monotonic rule

$$F_{j+1} = (F_j - \epsilon_j) \sqrt{2}$$

thus if we have the F_j that is the minimum F_j among the threads where it is defined then F_{j+1} will also be the minimum among those threads. At each time point j there is one thread (thread j) where an F value is defined for the first time. To take this into account, the minimum F_{j+1} is computed as follows:

$$F_{j+1} = \min(F_j - \epsilon_j, B - E_j) \cdot \sqrt{2}$$

Once we have the minimum F_{j+1} , we can check $F_{j+1} < 0$ to test whether at least one of the threads fails at that check. Note that this check is only executed when $E_j \leq B$ thus it can be replaced with

$$F_j - \epsilon_j < 0$$

We get the following algorithm:

- Let $E_0 = 0, F_1 = B$.
- For each time point $j = 1, \dots, n$:
 - Let ϵ_j be the amount of budget used at time point j .
 - Let $E_j = \frac{E_{j-1}}{\sqrt{2}} + \epsilon_j$.
 - If $E_j > B$ then **fail**.
 - If $F_j - \epsilon_j < 0$ then **fail**.
 - Reveal R_j .
 - Let $F_{j+1} = \min(F_j - \epsilon_j, B - E_j) \cdot \sqrt{2}$.

Let us write it in the imperative style, where the values of variables can change:

- Let $E := 0, F := B$.
- For each time point $j = 1, \dots, n$:
 - Let $E := \frac{E}{\sqrt{2}}$.
 - Let ϵ be the amount of budget used at time point j .
 - Let $E := E + \epsilon$.
 - Let $F := F - \epsilon$.
 - If $E > B$ then **fail**.
 - If $F < 0$ then **fail**.
 - Reveal R_j .
 - Let $F := \min(F, B - E)$.
 - Let $F := F \sqrt{2}$.

This shows most clearly how the budget renewal works. Between any two successive time points, the used budget E is divided by $\sqrt{2}$ and the available budget F is multiplied by $\sqrt{2}$.

If, before the renewal, $E + F > B$, the F is reduced until $E + F = B$. This ensures that the total budget (already used + available for future use) is never larger than B at any time.

The used budget by itself must not become larger than B and the available budget must not become negative. This sets the restrictions on how much budget can be used at any given time point. Note that this does not show the optimal use of the budget to get the best accuracy. It only ensures differential privacy.

Optimizing the accuracy. Let us try to get the best accuracy within the given limits. Before each renewal, we must have $E + F \leq B$, so let us assume that $E + F = B$ to have the largest renewal potential. The $E + F$ renews to $\frac{E}{\sqrt{2}} + F\sqrt{2}$. If we want to get back to the previous E and F , we can use $\epsilon = (1 - \frac{1}{\sqrt{2}})E$ and $(\sqrt{2} - 1)F \geq \epsilon$. We get

$$3\epsilon \leq (\sqrt{2} - 1)(E + F) = (\sqrt{2} - 1)B$$

$$\epsilon \leq \frac{\sqrt{2} - 1}{3}B$$

Thus we can use up to $\frac{\sqrt{2}-1}{3}B$ or about 13.8% of the whole budget at each time point of an infinite time series. For a finite time series, we can use a bit more if we use up the budget over the finite time, so that in the end we have $E = B$ and $F = 0$.

Generalizations. We have investigated the case where the time series we want to reveal is a Markov chain (also in the opposite direction) whose elements are all from the standard normal distribution and the underlying value changes over time with a specific speed. It seems possible to generalize our results to the case where the elements of the time series are from other normal distributions and different elements may be from different normal distributions, and where the underlying value may change at different speeds.

Distance for differential privacy. Above we showed how to get differential privacy w.r.t. each X_i . We however assumed a certain distribution of (X_1, \dots, X_n) but differential privacy should be independent of the input distribution. Thus the input there was not actually (X_1, \dots, X_n) but $(W_2, \dots, W_i, X_i, Y_i, \dots, Y_{n-1})$ when considering differential privacy w.r.t. X_i .

Thus the input is different for each X_i but each of these inputs (and also (X_1, \dots, X_n)) uniquely determines all the others. We would like to use only (X_1, \dots, X_n) as the input and achieve the same results. Thus we need to define a distance on the (X_1, \dots, X_n) . Let h_i be the function that converts (X_1, \dots, X_n) to the corresponding $(W_2, \dots, W_i, X_i, Y_i, \dots, Y_{n-1})$:

$$h_i(X_1, \dots, X_n) = (W_2, \dots, W_i, X_i, Y_i, \dots, Y_{n-1})$$

where

$$Y_j = X_{j+1} \sqrt{2} - X_j$$

$$W_j = X_{j-1} \sqrt{2} - X_j$$

Let H_i be the inverse of this function:

$$H_i(W_2, \dots, W_i, X_i, Y_i, \dots, Y_{n-1}) = (X_1, \dots, X_n)$$

where

$$X_{j+1} = \frac{X_j + Y_j}{\sqrt{2}} \text{ for } j = i, i + 1, \dots, n - 1$$

$$X_{j-1} = \frac{X_j + W_j}{\sqrt{2}} \text{ for } j = i, i - 1, \dots, 2$$

Another way to express X_j here is

$$X_j = X_i(\sqrt{2})^{i-j} + \sum_{m=1}^{j-i} Y_{j-m}(\sqrt{2})^{-m} \text{ if } j > i$$

$$X_j = X_i(\sqrt{2})^{j-i} + \sum_{m=1}^{i-j} W_{j+m}(\sqrt{2})^{-m} \text{ if } j < i$$

Approved for Public Release; Distribution Unlimited.

Then the distance d on the (X_1, \dots, X_n) must satisfy

$$d(H_i(W_2, \dots, W_i, X_i, Y_i, \dots, Y_{n-1}), H_i(W_2, \dots, W_i, X_i + c, Y_i, \dots, Y_{n-1})) = |c|$$

since here we modify X_i by $|c|$ and recompute the X_i where $i \neq j$.

Let $g_{i,c}$ be the following function:

$$g_{i,c}(W_2, \dots, W_i, X_i, Y_i, \dots, Y_{n-1}) = (W_2, \dots, W_i, X_i + c, Y_i, \dots, Y_{n-1})$$

Note that

$$(H \circ g_{i,c} \circ h)(X_1, \dots, X_n) = (X'_1, \dots, X'_n)$$

where

$$X'_j = X_j + \frac{c}{(\sqrt{2})^{|i-j|}}$$

This describes a change in (X_1, \dots, X_n) corresponding to distance $|c|$. We may actually allow

$$|X_j - X'_j| \leq \frac{c}{(\sqrt{2})^{|i-j|}}$$

without increasing the distance $|c|$ because reducing the change in X_j or reversing the direction of the change does not increase the DP-distance change in the distribution of the released output R_j .

We have defined the distance d for only some pairs of values. To make it complete, we take the transitive closure. Then using the budget tracking algorithm described earlier, we can guarantee B -differential privacy w.r.t. distance d (if the algorithm does not fail).

3.3.12.7 Enforcing Budget Limits. In this section, we will see in which cases the budget limits can be enforced by excluding the rows whose budget becomes exhausted (as opposed to some known restrictions on the input distribution guaranteeing an upper bound on the amount of budget used). We will also consider the case where budgets can be increased over time.

Banach sensitivity with enforced budget limits. When using only Banach sensitivity (sensitivity w.r.t. changes inside rows), the budgets are public because changes inside rows do not change the amount of budget used. (The amount of budget used depends on the join keys but changing the join key is considered as exchanging the whole row instead of as a change inside the row.) Thus we can exclude rows whose budget is exhausted and still retain DP. The excluded rows also do not need to be taken into account when computing the sensitivity.

Adding/removing rows sensitivity with enforced budget limits. So far we could only compute (without revealing) the amount of budget used but could not guarantee that a given limit is not exceeded.

Let us use the following mechanism. Consider the query that counts the number of rows in the table (let this be the *joined table*) obtained by joining the N tables together in the way described above. To compute this query differentially privately, we use a personalized budget ϵ_i for each r_i . Let ϵ_0 be the amount of budget needed for each use of r_i . Thus r_i can be used $\lfloor \epsilon_i / \epsilon_0 \rfloor$ times.

Consider a row R in the joined table. Let the rows from which it is composed be r_{s_1}, \dots, r_{s_N} where $\tau_{s_i} = i$. Sort the list (s_1, \dots, s_N) in decreasing order to get (s'_1, \dots, s'_N) , which we call the *timestamp* of R . We use the lexicographic order on timestamps to order rows of the joined table by time. The first component of the timestamp is the time point when a row of the joined table is used, the rest of the components are used to order rows used at the same time point.

We sort the rows of the joined table in increasing order of timestamps. Then we go through the rows in this order. For each row R , we check if all the component rows of R have at least ϵ_0 budget left. If yes, we include R in the query result and decrease the budget of each component row by ϵ_0 . If no, we do not include R in the query result and do not change budgets.

However, this mechanism cannot work for the most general class of queries because we have the following counterexample:

SELECT * FROM t1, t2 WHERE t2.id = t1.id OR t2.id = t1.id + 1

Each of the id columns is a primary key of its table but due to the OR in the condition, each row of either table can be joined to up to 2 rows of the other table. Suppose that each table contains rows with $\text{id} = 1, 2, \dots, k$ and each row has budget for only one use. Then the joined rows that are included, are the ones with $\text{t2.id} = \text{t1.id}$. If we now remove from the second table the row with $\text{id} = 1$ then the joined rows that are included, will be the ones with $\text{t2.id} = \text{t1.id} + 1$, i.e. the ones previously excluded. Thus $k - 1$ rows become included and k rows become excluded, for a total sensitivity of $2k - 1$. The necessary DP noise makes the differentially private result useless.

Let us try if the mechanism can achieve DP for some special cases where the query is restricted in some way.

1 : n Joins. Let there be N tables numbered $1, \dots, N$. Let r_1, \dots be rows that are added to the tables in that order, i.e. r_i is added at time point i . Let τ_i be the number of the table into which r_i is added. The N tables are joined together so that for each i , tables i and $i + 1$ are joined with 1 : n join, each row of table $i + 1$ can be joined with only one row of table i but a row of table i can be joined with any number of rows of table $i + 1$. Let $r_{\rho(i)}$ be the unique row of table $\tau_i - 1$ to which r_i can be joined. If $\tau_i = 1$ then let $\rho(i) = 0$.

Now let us consider what is the sensitivity of the query result computed with the mechanism proposed in the previous paragraph w.r.t. removing a row r_k . Removing r_k removes all uses of r_k in the joined table that were included in the query result. Consider one such use, i.e. a row R in the joined table.

Lemma 3.103. *Removing R cannot cause more than one other row to become included (unless budgets can be increased over time).*

Proof. Let the rows from which R is composed be r_{s_1}, \dots, r_{s_N} where $\tau_{s_i} = i$. If we remove R then each r_{s_i} will have budget for one extra use. Suppose that this causes a row R' to be included in the query. Let the rows from which it is composed be $r_{s'_1}, \dots, r_{s'_N}$ where $\tau_{s'_i} = i$. Because R' was included only after R was removed, there must be some $r_{s'_m}$ which obtained the budget that was previously used by R , i.e. $s'_m = s_m$. We also have $s'_{m-1} = \rho(s'_m) = \rho(s_m) = s_{m-1}$. Similarly we get $s'_j = s_j$ for all $j = m, m-1, \dots, 1$. Also R' must have a later timestamp than R , otherwise R' would have been included previously instead of R .

Suppose that in addition to R' , there is another row R'' included when R is removed. W.l.o.g. let the timestamp of R'' be later than that of R' (they both must be later than R). Let the rows from which R'' is composed be $r_{s''_1}, \dots, r_{s''_N}$ where $\tau_{s''_i} = i$. Similarly to above we get that for some $\ell \geq 1$, we get $s''_j = s_j$ for all $j = \ell, \ell - 1, \dots, 1$.

Consider the case $\ell \leq m$. Then $s''_\ell = s_\ell = s'_\ell$. Excluding R allowed R'' to be included because it obtained the budget for $r_{s''_\ell}$ from R . But R' also consumes the budget for $r_{s''_\ell} = r_{s'_\ell}$ and because R' has an earlier timestamp than R'' , this budget released from R never reaches R'' . Thus removing R cannot cause R'' to become included.

Now consider the case $\ell > m$. Then $s''_m = s_m = s'_m$. Excluding R allowed R'' to be included because it obtained the budget for $r_{s''_\ell}$ from R . This budget is not consumed by R' . But R' consumes the budget for $r_{s'_m} = r_{s''_m}$ and because R' has an earlier timestamp than R'' , this budget released from R never reaches R'' . Thus removing R cannot cause R'' to become included, unless some extra budget is added for $r_{s''_m}$ between the timestamps of R' and R'' and it is not allowed to use this budget to include rows of the joined table that were previously excluded.

Thus removing R can cause at most one row to become included. □

If the extra budget cannot be used for previous rows then up to T rows can become included. Indeed, the budget released from each of the T component rows of R can be used by at most one of the rows that become included, as each row use needs the same amount, ϵ_0 .

Lemma 3.104. *Removing R cannot cause any other rows to become excluded (unless budgets can be increased over time).*

Proof. Suppose that there is a row R'' that becomes excluded when R is removed and R' is included. The timestamp of R'' must be later than R' because excluding R cannot use up the budgets needed by R'' , only including R' can. Let the rows from which R'' is composed be $r_{s'_1}, \dots, r_{s'_N}$ where $\tau_{s'_i} = i$. Then R' must have used up the budget of $r_{s'_\ell}$ for some ℓ . Thus $s'_\ell = s''_\ell$ and similarly to above we get that $s''_j = s'_j$ for all $j = \ell, \ell - 1, \dots, 1$.

Consider the case $\ell \leq m$. Then $s_\ell = s'_\ell = s''_\ell$. Thus the budget of r_{s_ℓ} was previously used by R and after removing R by R' , thus removing R cannot affect the budget of r_{s_ℓ} available for R'' , thus it cannot cause R'' to become excluded.

Now consider the case $\ell > m$. Then $s_m = s'_m = s''_m$. The budget of r_{s_m} was previously used by R and after removing R by R' , thus removing R cannot affect the budget of r_{s_m} available for R'' . If there is not enough budget of r_{s_m} to include R'' after removing R , then there was not enough budget also before removing R , thus removing R cannot cause R'' to become excluded. \square

Thus we can achieve DP for queries that use only $1 : n$ joins.

Increasing budgets. So far, budgets were not increased over time. The whole budget was available from the beginning. Note that budgets cannot be decreased over time because they are private and we do not know if there is enough budget left to take away. Let us now consider increasing budgets over time, i.e. between any two consecutive time points, we can add a non-negative number to each budget (the number may be different for different budgets). Also, the initial budgets (before the first time point) are non-negative.

So far we have only considered how many rows can become included or excluded but we have not counted the rows that change their inclusion time. If budgets are not changed or extra budgets cannot be used for rows excluded earlier then rows can be included only at the time they become available for inclusion (i.e. the last of its components is added to the database). Thus inclusion times cannot change in this case.

However, if extra budget can be used for previously excluded rows then removing R can change inclusion times of other rows. For example, suppose R_1, \dots, R_k are in increasing order of timestamps and all contain the row r . Initially r has budget for only one use and R uses it up. For each $i = 2, \dots, k$, we increase the budget of r by one use between right before the timestamp of R_i (but strictly after that of R_{i-1}). Thus for each $i = 1, \dots, k - 1$, the row R_i will be included at the timestamp of R_{i+1} and the row R_k will be excluded. Now, if we remove R , the released budget can be used to include R_k . However, this is not the only change, the rows R_1, \dots, R_{k-1} were previously included and remain included but their inclusion times change—now each R_i will be included at its own timestamp instead of that of R_{i+1} .

Thus we cannot allow using the extra budget for previously excluded rows in the joined table, i.e. for previously excluded uses of rows in the original tables. Note that the extra budget can still be used for new uses of the rows in the original tables whose previous uses were excluded. Specifically, the extra budget can be used for those and only those rows of the joined table for which at least one of its components (rows of the original tables) was added later than the extra budget.

Let us find the sensitivity of the query result w.r.t. removing a row use R when budgets can be increased. Lemma 3.103 becomes weaker:

Lemma 3.105. *If budgets can be increased then removing R cannot cause more than T other rows to become included.*

Proof. The budget ϵ_0 released from each of the T component rows of R can be used by at most one of the rows that become included, as the latter uses up the same amount, ϵ_0 . \square

Lemma 3.104 still holds:

Lemma 3.106. *If budgets can be increased then removing R cannot cause any other rows to become excluded.*

Proof. We use the proof of Lemma 3.104 but there are up to T possible choices for R' in the proof. At least one of them must have used up the budget of $r_{s'_\ell}$ for some ℓ . Choose one of those as R' . \square

Thus the sensitivity increases from 2 to $T + 1$ when budgets can be increased.

Combining with global Banach sensitivity. So far we got DP w.r.t. adding/removing rows for $1 : n$ joins for COUNT queries. Now let us combine it with global Banach sensitivity and use SUM queries. This is similar to using global sensitivity in Sec. 3.3.12.5.

First consider the case when budgets are not increased. Let L be the maximum possible absolute value contribution of a row (of the joined table) to the query result. Then the total sensitivity (across all time points) w.r.t. adding/removing (all uses of) a row would be $2LM \log_2 n$, where M is the maximum number of times a row can be used (if this is exceeded then some row uses are excluded). The factor 2 is sensitivity of the respective COUNT query w.r.t. adding/removing rows. If the global Banach sensitivity w.r.t. a row use is K then w.r.t. all uses of a row it is $MK \log_2 n$. Combined global sensitivity is then

$$C = \max\left(\frac{2LM \log_2 n}{G}, MK \log_2 n\right) = M \log_2 n \cdot \max\left(\frac{2L}{G}, K\right)$$

where G is the distance corresponding to adding/removing a row. The noise level for each released value is

$$\frac{M \log_2 n \cdot \max\left(\frac{2L}{G}, K\right)}{\epsilon}$$

where ϵ is the total budget.

Now consider the case where budgets are increased. Then the sensitivity of the respective COUNT query will be $T + 1$. Thus the factor 2 in $2L$ will become $T + 1$. The noise level for each released value will be

$$\frac{M \log_2 n \cdot \max\left(\frac{(T+1)L}{G}, K\right)}{\epsilon}$$

Combining with local Banach sensitivity. We can also combine with local Banach sensitivity. This is similar to combining global and Banach sensitivity in Sec. 3.3.12.5. Let $\epsilon = \epsilon_b + \epsilon_\beta$, where ϵ_b is the part of ϵ used to hide the query result and ϵ_β is the part of ϵ used to hide the noise level. Then

$$\beta = \frac{\epsilon_\beta}{(\gamma + 1)M \log_2 n}$$

$$b = \frac{\epsilon_b}{(\gamma + 1)M \log_2 n}$$

Let κ be β -smooth Banach sensitivity w.r.t. a single row use. The noise level for each released value will be

$$\frac{(\gamma + 1)M \log_2 n \cdot \max\left\{\frac{2L}{G}, \kappa, e^{-\beta G} K\right\}}{\epsilon_b}$$

if budgets are constant and

$$\frac{(\gamma + 1)M \log_2 n \cdot \max\left\{\frac{(T+1)L}{G}, \kappa, e^{-\beta G} K\right\}}{\epsilon_b}$$

if budgets can be increased.

3.3.13 Model-Checking Sensitivity of SQL Queries. Consider again SQL workflows, consisting of tables and queries, as defined in Sec. 3.2.1.1. We have proposed a model-checking based approach to determine the sensitivity of such workflows, which, differently from previous approaches can provide precise bounds in certain cases.

3.3.13.1 Non-GROUP BY Queries. The SQL workflows are made up of queries. We consider two types of queries here — those with out GROUP BY, and those containing it. The first kind of the queries is defined as follows.

A database query $Q = (J, F, P)$ (join—filter—project) over the database schema dbs consists of the components named below. Let us use the notation defined in Sec. 3.2.1.1.

- $J \in \{t_1, \dots, t_m\}^*$, where Z^* denotes the set of all sequences of elements of the set Z . Let $|J|$ be the length of J and $J[i]$ the i -th component of J . Let $D[J]$ denote the Cartesian product $D[J[1]] \times \dots \times D[J[|J|]]$. For $b \in D[J]$, let $b[i]$ denote the sequence of the elements of b corresponding to the table $J[i]$.
- F is a predicate on the set $D[J]$.
- P is a sequence of functions p_1, \dots, p_r , where $p_i : D[J] \rightarrow X_i$. The set X_i may be any set (it is the set of possible values in the i th column of the result table).

The application of a query $Q = (J, F, P)$ to a database $Y \in \mathcal{Y}$ proceeds as follows:

1. Let $A = \prod_{i=1}^{|J|} Y.J[i]$. Note that A is a multiset over $D[J]$.
2. Let $B \subseteq A$ be the multiset of all elements of A that satisfy F .
3. Output the multiset $\{(p_1(b), \dots, p_r(b)) \mid b \in B\}$.

Such kind of queries cover the SQL-queries containing JOIN-s and WHERE-clauses, where certain attributes of the joined tables are selected into the result. The set of projections P cover the computations over the attributes. Assuming that the predicate F and the functions p_1, \dots, p_r can be specified in SQL, the query Q corresponds to the following SQL query:

$$\text{SELECT } \overline{\llbracket p_1 \rrbracket}, \dots, \overline{\llbracket p_r \rrbracket} \text{ FROM } J[1], \dots, J[|J|] \text{ WHERE } \overline{\llbracket F \rrbracket}$$

where $\overline{\llbracket z \rrbracket}$ denotes the syntactic object (i.e. SQL expression) with semantics z .

The queries described here have the following distributivity properties, which are easy to check if one considers how a particular row may have ended up in the result of the query.

Lemma 3.107. *Let $dbs = (t_1 : r_1, \dots, t_m : r_m)$ be a database schema, $Q = (J, F, P)$ a query over dbs , and Y a database over dbs . Let $i \in \{1, \dots, m\}$ and R_1 and R_2 be relations over r_i . Let R'_1 be the result of Q on $Y[t_i \mapsto R_1]$ and R'_2 be the result of Q on $Y[t_i \mapsto R_2]$. Let R' be the result of Q on $Y[t_i \mapsto R_1 \cup R_2]$.*

- (a) *If J contains t_i only once, then $R' = R'_1 \cup R'_2$ and $R' \setminus R'_1 = R'_2$.*
- (b) *Suppose that $R_1 \cap R_2 = \emptyset$. Let i_1, \dots, i_k be the positions of t_i in J . Let $F_{i_j}^{R_2}$ be a predicate on $D[J]$, returning true on a row \mathbf{r} iff $\mathbf{r}[i_j] \in R_2$. Let Q' be the query $(J, F \wedge \bigvee_{j=1}^k F_{i_j}^{R_2}, P)$. Then the result of Q' on $Y[t_i \mapsto R_1 \cup R_2]$ is $R' \setminus R'_1$.*

Proof. Part (a). Let $t_i = J[j]$. Then by the description of application of queries above (and using the fact that for all $k \neq j$, $J[k] \neq t_i$), we have

$$R'_1 = \{(p_1(b), \dots, p_r(b)) \mid b \in \prod_{k=1}^{j-1} Y.J[k] \times R_1 \times \prod_{k=j+1}^{|J|} Y.J[k], F(b)\},$$

$$R'_2 = \{(p_1(b), \dots, p_r(b)) \mid b \in \prod_{k=1}^{j-1} Y.J[k] \times R_2 \times \prod_{k=j+1}^{|J|} Y.J[k], F(b)\},$$

$$R' = \{(p_1(b), \dots, p_r(b)) \mid b \in \prod_{k=1}^{j-1} Y.J[k] \times (R_1 \cup R_2) \times \prod_{k=j+1}^{|J|} Y.J[k], F(b)\}. \text{ Thus } R' = R'_1 \cup R'_2 \text{ and (because we are using multisets, not sets) } R' \setminus R'_1 = R'_2.$$

Part (b). W.l.o.g. (as reordering elements of a tuple is a set isomorphism) we can assume that $i_1 = 1, \dots, i_k = k$. Then the result of Q' on $Y[t_i \mapsto R_1 \cup R_2]$ is $R'' = \{(p_1(b), \dots, p_r(b)) \mid b \in (R_1 \cup R_2)^k \times \prod_{j=k+1}^{|J|} Y.J[j], F(b) \wedge \bigvee_{j=1}^k (b[j] \in R_2)\}$. Because for $j = 1, \dots, k$, we have $b[j] \in R_1 \cup R_2$, and $R_1 \cap R_2 = \emptyset$,

the condition $\bigvee_{j=1}^k (b[j] \in R_2) = \neg \bigwedge_{j=1}^k (b[j] \notin R_2) = \neg \bigwedge_{j=1}^k (b[j] \in R_1) = ((b[1], \dots, b[k]) \notin R_1^k)$. Thus $R'' = \{(p_1(b), \dots, p_r(b)) \mid b \in ((R_1 \cup R_2)^k \setminus R_1^k) \times \prod_{j=k+1}^{|J|} Y.J[j], F(b)\} = \{(p_1(b), \dots, p_r(b)) \mid b \in (R_1 \cup R_2)^k \times \prod_{j=k+1}^{|J|} Y.J[j], F(b)\} \setminus \{(p_1(b), \dots, p_r(b)) \mid b \in R_1^k \times \prod_{j=k+1}^{|J|} Y.J[j], F(b)\} = R' \setminus R_1'$. \square

Set semantics. When evaluating the queries in set semantics, a relation R over the schema r is a subset of $D[r]$ instead of a multiset over $D[r]$, and the multiset constructor in Step 3 of Sec. 3.3.13.1 is replaced with set constructor, i.e. duplicates are removed. We actually only require the output relation to be a set, the input relations can still be multisets.

3.3.13.2 GROUP BY Queries. A GROUP BY query $Q = (J, F, G, P_G, P_A)$ (*join—filter—group—project groups—aggregate*) consists of the following components:

- J and F are the same as for non-GROUP BY queries.
- G is similar to the P for non-GROUP BY queries. It is a sequence of functions g_1, \dots, g_r , where $g_i : D[J] \rightarrow X_i$. The set X_i is arbitrary. Typically each g_i selects an attribute from $D[J]$. Let $X = X_1 \times \dots \times X_r$.
- P_G is a sequence of functions $p_1, \dots, p_{r'}$, where $p_i : X \rightarrow Y_i$. The set Y_i is arbitrary. Hence, if each g_i selected an attribute from $D[J]$ then each p_i is a function of those attributes.
- P_A is a sequence of functions h_1, \dots, h_s , where $h_i : (D[J] \rightarrow \mathbb{N}) \rightarrow Z_i$. The set Z_i is arbitrary.

The application of a query $Q = (J, F, G, P_G, P_A)$ to a database $Y \in \mathcal{Y}$ proceeds as follows:

1. Let $A = \prod_{i=1}^{|J|} Y.J[i]$. Note that A is a multiset over $D[J]$.
2. Let $B \subseteq A$ be the multiset of all elements of A that satisfy F .
3. Let $C = \{(g_1(b), \dots, g_r(b)) \mid b \in B\}$ be the set of *group keys*.
4. Let $D = \{(c, \{b \in B \mid (g_1(b), \dots, g_r(b)) = c\}) \mid c \in C\}$ be the set of *groups*.
5. Output the multiset $\{(p_1(c), \dots, p_{r'}(c), h_1(E), \dots, h_s(E)) \mid (c, E) \in D\}$.

Assuming that the predicate F and the functions $g_1, \dots, g_r, p_1, \dots, p_{r'}, h_1, \dots, h_s$ can be specified in SQL, the query Q corresponds to the following SQL query:

```
SELECT  $\overline{[p_1]}$ , ...,  $\overline{[p_{r'}]}$ ,  $\overline{[h_1]}$ , ...,  $\overline{[h_s]}$  FROM  $J[1], \dots, J[|J|]$ 
WHERE  $\overline{[F]}$  GROUP BY  $\overline{[g_1]}$ , ...,  $\overline{[g_r]}$  .
```

Note that $\overline{[h_i]}$ are the aggregations of attributes, e.g. SUM, COUNT, MAX etc.

3.3.13.3 Queries Combined using Set Operations. If Q_1 and Q_2 are supported queries over set semantics then we can also support queries $Q_1 \cup Q_2$, $Q_1 \cap Q_2$, and $Q_1 \setminus Q_2$. If the results of queries Q_1 and Q_2 are R_1 and R_2 , respectively, then the results of the queries $Q_1 \cup Q_2$, $Q_1 \cap Q_2$, and $Q_1 \setminus Q_2$ are $R_1 \cup R_2$, $R_1 \cap R_2$, and $R_1 \setminus R_2$, respectively. Set operations can be combined any number of times but only at the top level, e.g. $Q_1 \setminus ((Q_2 \cup Q_3) \cap Q_4)$.

3.3.13.4 Sensitivity of queries without GROUP BY. Consider the queries described in Sec. 3.3.13.1. Note that we defined these queries to use multiset semantics, not set semantics.

In this analysis, we consider the “canonical” distance over databases, meaning that for datasets R, R' over the same schema we define their distance as the cardinality of their symmetric difference $(R \setminus R') \cup (R' \setminus R)$. For databases over the same schema we define their distance as the sum of the distances of the corresponding tables.

To find the sensitivity of a query $Q = (J, F, P)$ against a database schema $db_s = (t_1 : r_1, \dots, t_m : r_m)$ with respect to the table t_i , we will find the *derivative* query $Q_{t_i}[y]$ of Q with respect to t_i . The derivative query is parameterized with a possible row y of the table t_i . The result of $Q_{t_i}[y]$ on a database Y is the multiset difference of the result of Q on Y and the result of Q on $Y[t_i \mapsto Y.t_i \setminus \{y\}]$ (database Y with one copy of row y removed from the table t_i). We will then find the maximum number of rows that may be returned by $Q_{t_i}[y]$ on Y , maximized over all possible y and Y . This is the sensitivity of Q with respect to the table t_i .

Let us first consider the case where $\text{Dis}_{r_j} \neq \emptyset$ for all j . In this case the input tables cannot contain repeated rows, i.e. these multisets are actually sets. Due to the distributivity properties (Lemma 3.107) of the query $Q = (J, F, P)$, it is very simple to find its derivatives. Let i_1, \dots, i_k be the positions that contain t_i in J . Then

$$Q_{t_i}[y] = (J, F \wedge \bigvee_{j=1}^k F_{i_j}^{(y)}, P),$$

where $F_{i_j}^{(y)}$ is defined as in Lemma 3.107. Note that we used removals instead of additions to define the derivative query to ensure that the removed row y is included in the input of the derivative. Otherwise it would not be possible to express the derivative as a query in the sense of Sec. 3.3.13.1 (it would have to return a nonempty output on an empty input table).

Let $N \in \mathbb{N}$. The question of whether $Q_{t_i}[y]$ may return at least N rows can be phrased as a quantifier-free formula Φ in the manner we describe below. If this formula is satisfiable, then the answer to the query may contain at least N rows for some y . By varying N and checking for satisfiability of resulting formulas, we will find the value where it is no longer satisfiable. Satisfiability is modulo theories we commonly associate with the domains D_1, \dots, D_n , where the attributes in the tables are picked from. Hence the theories may include Booleans, integers, reals and strings. The first three are well-supported by existing SMT solvers. We can also handle queries over string data, as long as the operations performed with them are supported by SMT solvers.

Note that because we are using multiset semantics, each combination of rows of input tables that satisfies F , produces exactly one output row, repeated output rows are not removed. Thus the formula will not depend on P .

The formula $\Phi_{Q,N}^{db_s, t_i}$ has three major parts: the variables, the functions (function variables), and the constraints. The conjunction of the constraints is the actual formula. Each variable and each function is described with its type, i.e. the set of values from which it can take values. The SMT solver tries to assign to each variable and function a value of its given type, such that the constraints become satisfied. The variables are the following:

- For each $n \in \{1, \dots, N\}$, $j \in \{1, \dots, |J|\}$, and $k \in \{1, \dots, m\}$, where $r(a_1 : D_1, \dots, a_m : D_m; \text{Dis}_r)$ is the relation schema of the table $J[j]$: a variable $x[n, j, k]$ taking values in the set D_k . These variables enumerate the attributes of N rows from $D[J]$.
- For each $k \in \{1, \dots, m\}$, where $r(a_1 : D_1, \dots, a_m : D_m; \text{Dis}_r)$ is the relation schema of the table t_i : a variable $y[k]$, taking values in the set D_k .

The functions are the following:

- For all $i \in \{1, \dots, M\}$, where $db_s = (t_1 : r_1, \dots, t_M : r_M)$, all minimal $I = \{i_1, \dots, i_k\}$ such that $\{a_{i_1}, \dots, a_{i_k}\} \in \text{Dis}_r$, and all $\ell \in \{1, \dots, m\} \setminus I$, where $r(a_1 : D_1, \dots, a_m : D_m; \text{Dis}_r)$ is the relation schema of the table t_i : a function $f[i, I, \ell] : D_{i_1} \times \dots \times D_{i_k} \rightarrow D_\ell$.
- A function $g : D[J] \rightarrow \mathbb{Z}$.

The constraints of the formula, referring to the variables and functions above, are the following.

- For each $n \in \{1, \dots, N\}$: the tuple of variables $x[n, \cdot, \cdot]$ must satisfy the predicate F .
- For each $n \in \{1, \dots, N\}$, the following disjunction over all $j \in \{1, \dots, |J|\}$, where $J[j] = t_i$, must hold:
 - Each element of the disjunction is a conjunction of the statements $y[k] = x[n, j, k]$ over all $k \in \{1, \dots, m\}$, where $r(a_1 : D_1, \dots, a_m : D_m; \text{Dis}_r)$ is the relation schema of the table t_i .

(4) $\text{Dis}_{r_j} = \emptyset$ for some $j \neq i$

In cases (1) and (2), $\Phi_{Q,N}^{dbs,t_i}$ is satisfiable iff the sensitivity of Q with respect to the table t_i is at least N .

In cases (3) and (4), $\Phi_{Q,N}^{dbs,t_i}$ is satisfiable iff the sensitivity of Q with respect to the table t_i is ∞ and not satisfiable iff the sensitivity of Q with respect to the table t_i is 0.

Proof. We first consider the sensitivity of Q (w.r.t. t_i) over the database schema $dbs' = (t_1 : r'_1, \dots, t_m : r'_m)$ where $r'_j = r_j(a_1 : D_1, \dots, a_m : D_m; \text{Dis}_{r_j})$ if $\text{Dis}_{r_j} \neq \emptyset$ and $r'_j = r'_j(a_1 : D_1, \dots, a_m : D_m; \{\{a_1, \dots, a_m\}\})$ if $\text{Dis}_{r_j} = \emptyset$. The schema dbs' forbids repeated rows but is otherwise the same as dbs . If the sensitivity of Q over dbs' is 0 then the predicate F is false for all possible inputs and thus the result of the query Q (over both dbs' and dbs) is empty for all possible inputs. Thus the sensitivity of Q over dbs is also 0.

Now suppose that the sensitivity of Q over dbs' is at least 1. Then there exists a database Y over dbs' and a row $y \in X_{r_i}$ such that adding y to $Y.t_i$ adds at least 1 row to the result of Q . Let c be a one of the rows added to the result and $b \in D[J]$ the combination of rows of input tables that produces c .

If $\text{Dis}_{r_j} = \emptyset$ for some $j \neq i$ then we can let $Y.t_j$ contain $m \in \mathbb{N}$ copies of $b[\ell]$ where $J[\ell] = t_j$, which will cause adding y to $Y.t_i$ to produce at least m added copies of c in the result. Because m may be arbitrarily large, the sensitivity of Q over dbs is infinite.

If $\text{Dis}_{r_i} = \emptyset$ and t_i occurs more than once in J (i.e. $k \geq 2$, where i_1, \dots, i_k are the positions that contain t_i in J) then at least one of $b[i_1], \dots, b[i_k]$ must be equal to y . Let ℓ be such that $b[i_\ell] = y$. Let $s \in \{i_1, \dots, i_k\} \setminus \{i_\ell\}$. Then we can let $Y.t_i$ contain $m \in \mathbb{N}$ copies of $b[s]$ (which may or may not be equal to y), which will cause adding y to $Y.t_i$ to produce at least m added copies of c in the result. Because m may be arbitrarily large, the sensitivity of Q over dbs is infinite.

If $\text{Dis}_{r_j} \neq \emptyset$ for all $j \neq i$, $\text{Dis}_{r_i} = \emptyset$, and t_i occurs only once in J then we can use part (a) of Lemma 3.107 to get that the result of $Q_{t_i}[y]$ over dbs on database Y where $y \in Y.t_i$, is equal to the result of Q over dbs on database $Y' = Y[t_i \mapsto \{y\}]$. The database Y' does not contain repeated rows in any table, thus the result is the same over dbs' and is also equal to the result of $Q_{t_i}[y]$ over dbs' . Thus the sensitivity of Q (w.r.t. t_i) over dbs is equal to the sensitivity of Q over dbs' . \square

3.3.13.6 Tables Used More than Once. The description of the computation of the sensitivity of a query in Sec. 3.3.13.4 allows using an input table more than once in the join. This is not the only way this feature can be implemented in the analyzer. Instead, the analyzer can compute the sensitivities of the query as if all the tables were different. Then it can add the sensitivities w.r.t. each copy of the table together to obtain the sensitivity w.r.t. the original table, similarly to the partial derivatives of compound multi-variable functions.

This implementation choice has the advantage that because the sensitivities w.r.t. each copy of a table may be lower than w.r.t. the original table, the SMT solver is more likely to terminate within reasonable time (see Sec. 4.3.6).

The disadvantage is that this version of the analysis is not as precise. The actual sensitivity may be lower than the obtained result (it cannot be higher because the computed sensitivities are global, i.e. they do not depend on the actual value of the input database).

3.3.13.7 Sensitivity of Queries with Set Semantics. We define the sensitivity of a query Q over set semantics w.r.t. table t_i as the maximum number of rows in the derivative $Q_{t_i}\{y\}(Y)$, maximized over all possible y and Y , where $Q_{t_i}\{y\}(Y)$ is the set difference of the result of Q on Y and the result of Q on $Y[t_i \mapsto Y.t_i \setminus \{y\}]$ (database Y with the row y (if it exists) removed from the table t_i).

To find the sensitivity of a query in set semantics (SELECT DISTINCT in SQL) we modify the formula $\Phi_{Q,N}^{dbs,t_i}$ as follows. Let the modified formula be $\Psi_{Q,N}^{dbs,t_i}$. In set semantics, repeated rows in the output are removed. Thus the formula will now depend on \bar{P} . Instead of requiring the distinctness of the combinations of rows of input tables, we require the distinctness of the rows in the answer of the query. The following variables are added:

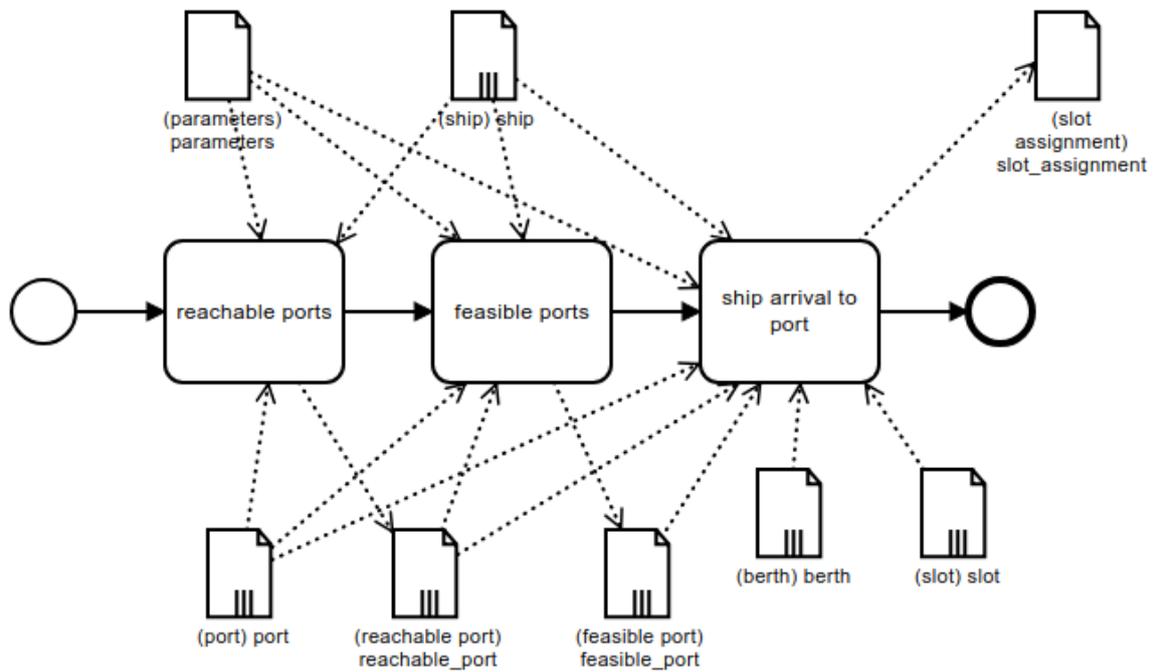


Figure 58: Workflow for the Aid Distribution Scenario

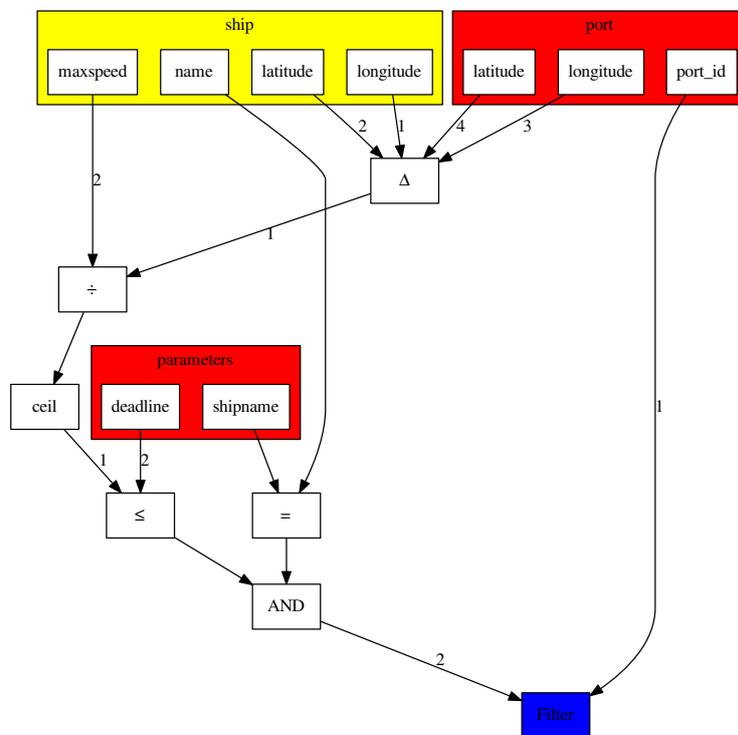


Figure 59: Leaks From the First Step of the Aid Distribution Scenario (1st Part)

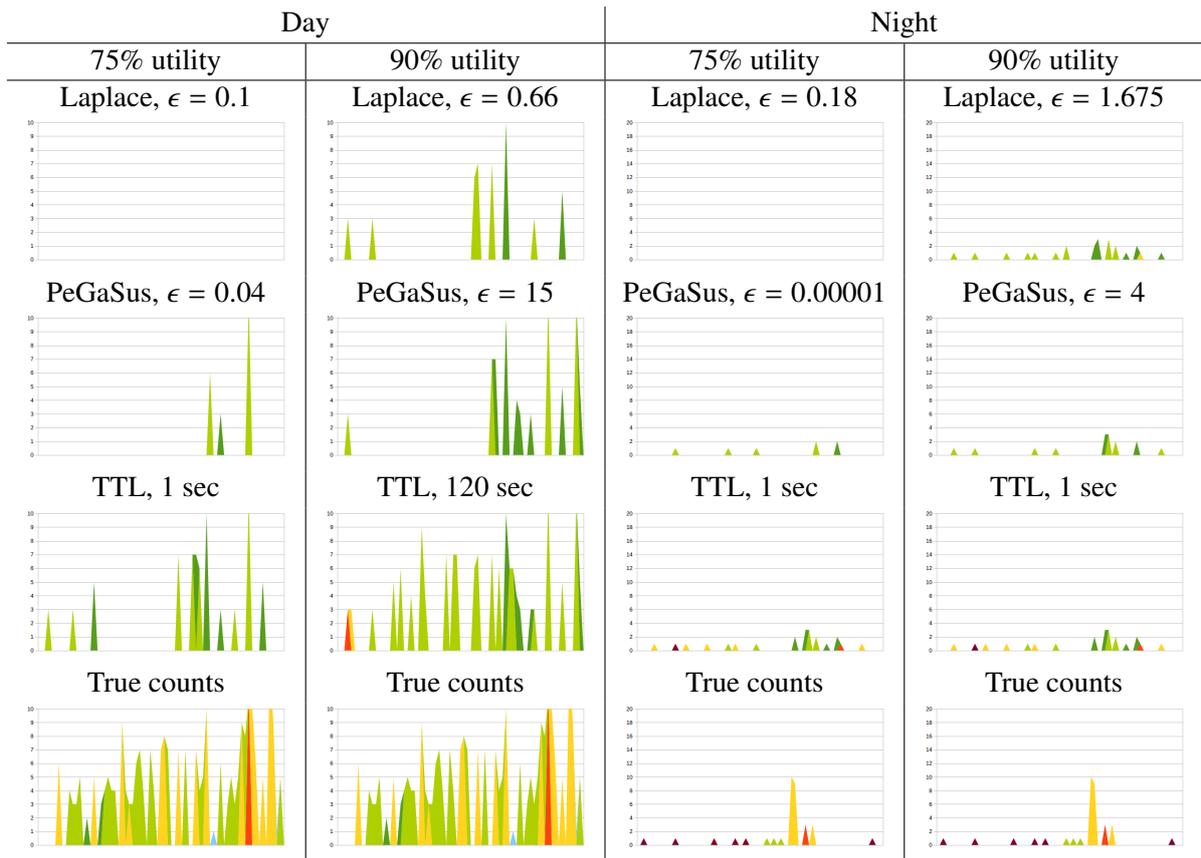


Figure 73: Comparison of Different Mechanisms for the Administrator Attacker

Table 27: Running Times (in Seconds) of Computing Posterior Probabilities for Laplace Noise

	Preprocess	$\epsilon = 0.1$	$\epsilon = 0.5$	$\epsilon = 1.0$	$\epsilon = 5.0$	$\epsilon = 10.0$	$\epsilon = 20.0$	total time
worst-case DP	4.8	0.58	0.57	0.57	0.62	0.60	0.62	8.8
$P_Y(x, y, A)$	5	5.1	5.1	5.2	5.4	5.5	5.5	38.4
$P_C(x, c, A)$	5.2	756	180	182	231	233	300	1898
$P_P(x, A)$	5.5	930	320	450	760	1020	1340	4840

different times for different ϵ values separately. We see that, while in general computation time grows with ϵ , this is also larger for $\epsilon = 0.1$ as well. The reason is that the noise of smaller ϵ has larger variance, so in the integration we need to cover a larger span of y values for which the noise is non-negligible.

Summary. Differential privacy in general consider very strong attackers that have access to almost unlimited information. In our set up with more realistic adversaries, we have seen that when the adver-

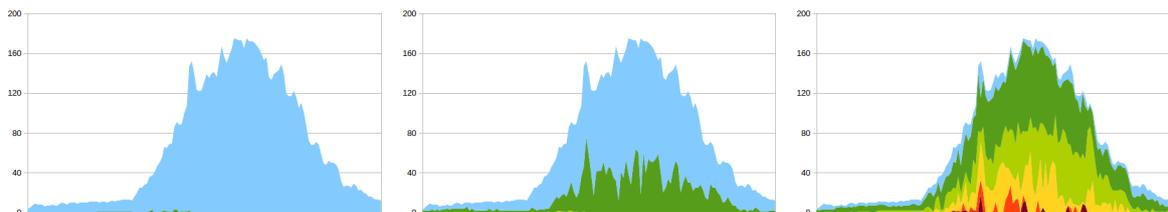


Figure 74: Prior Guessing Probabilities for $\delta = 0\%$ (Left), $\delta = 50\%$ (Middle), $\delta = 90\%$ (Right)

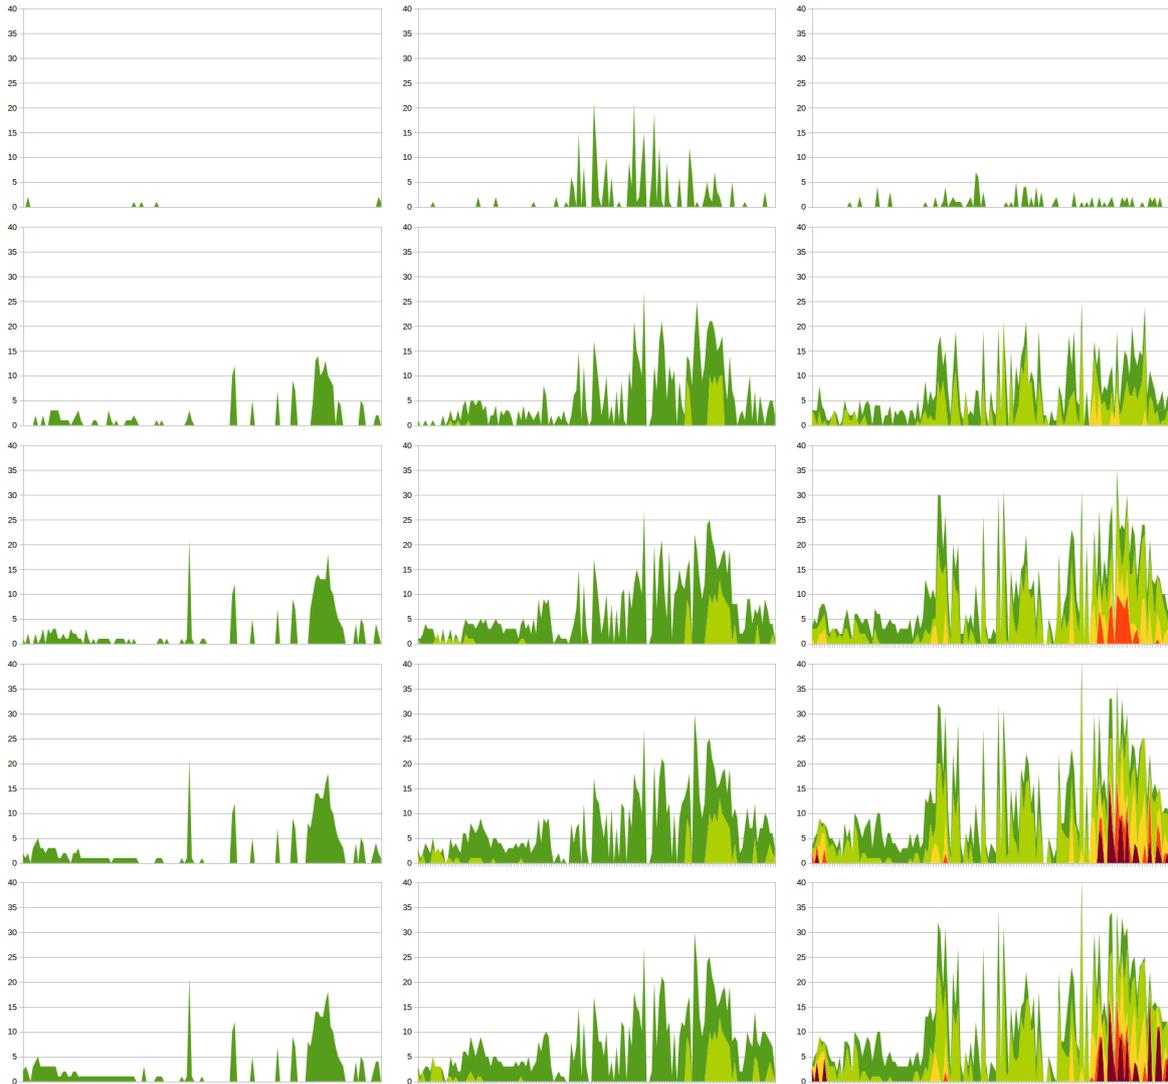


Figure 75: Posterior Guesses from Noisy Laplace Counts for $\delta = 0\%$ (Left), $\delta = 50\%$ (Middle), $\delta = 90\%$ (Right), $\epsilon \in \{0.1, 1.0, 5.0, 10.0, \infty\}$ (Top to Bottom)

sary is weaker (e.g., our external attacker). The practical privacy offered by Laplace and PeGaSus is almost the same than the one offered by TTL. Even when the practical privacy for Laplace and PeGaSus is similar, for higher utility values the formal privacy guarantee for PeGaSus is less than Laplace during the daytime and nighttime, whereas for lower utility values it is the opposite. When the adversaries become stronger (e.g., the student or administrator attacker), Laplace and PeGaSus offer more practical privacy than TTL, as it is expected. However, in some specific situations (e.g., in the afternoon when the building is less occupied) all the techniques behave similarly. Additionally, with stronger attackers the privacy loss due to the prior and adversarial knowledge at the time of the attack is high already. This means that effectively, in such situations the privacy of most of the individuals would be already compromised. Therefore, the difference between the differential privacy based techniques and TTL in terms of number of people being localized is small.

We would like to highlight that even when in terms of practical privacy the techniques behave similarly, TTL lacks of formal privacy guarantees which means that stronger attackers using more sophisticated attack methods could potentially result in higher privacy loss. Additionally, the results for Laplace and PeGaSus, which are the average over 30 counts, could potentially be worse depending on the noisy count generated in a single run at publishing time. Comparing the distributions of attacker's success for different privacy mechanisms remains out of scope of this work.

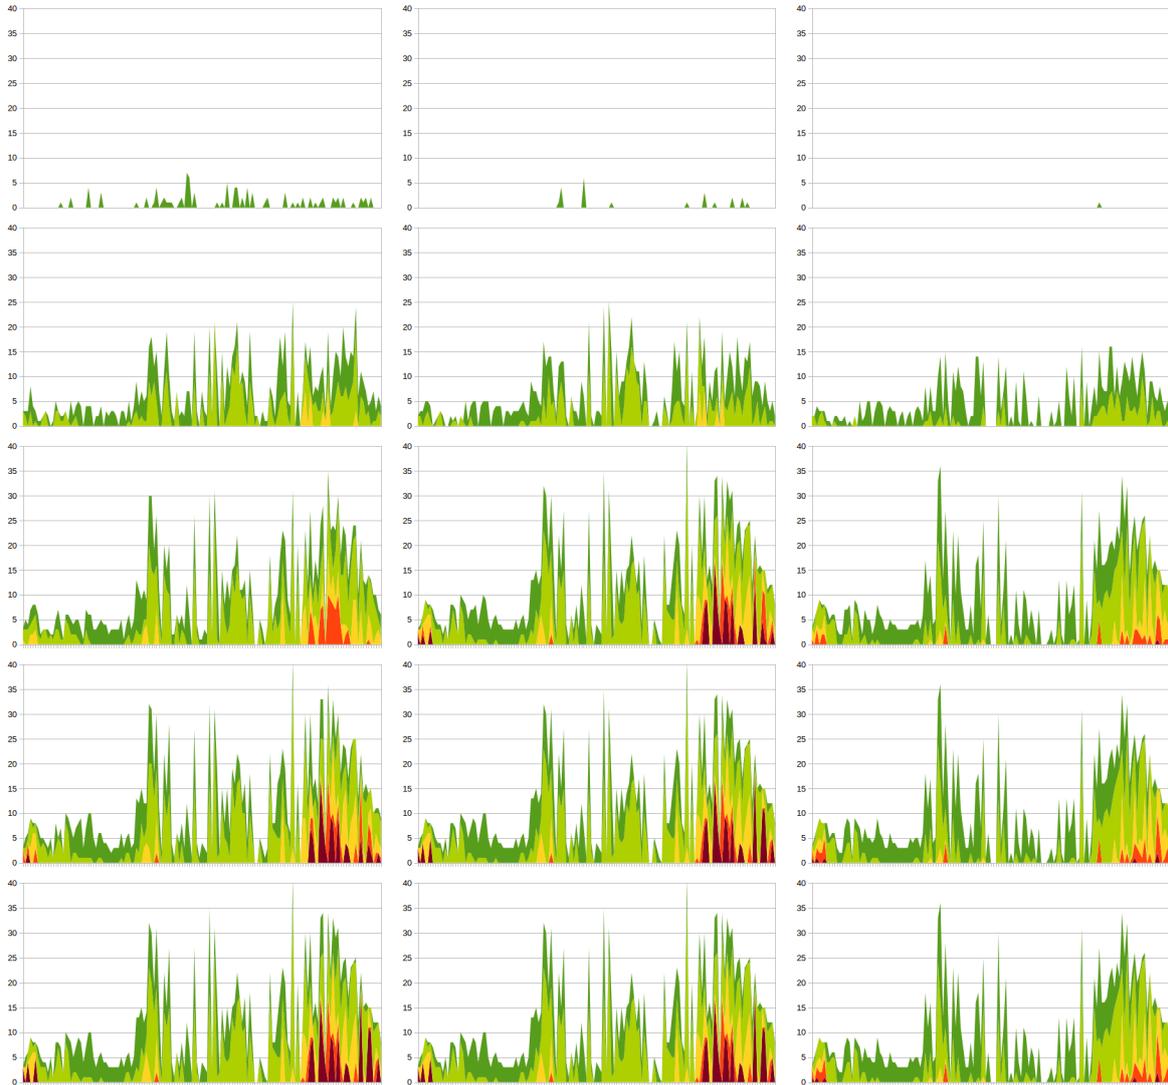


Figure 76: Posterior Guesses for $\delta = 90\%$ from Noisy Laplace Counts (Left), True Counts (Middle), and Count Distribution (Right), $\epsilon \in \{0.1, 1.0, 5.0, 10.0, \infty\}$ (Top to Bottom)

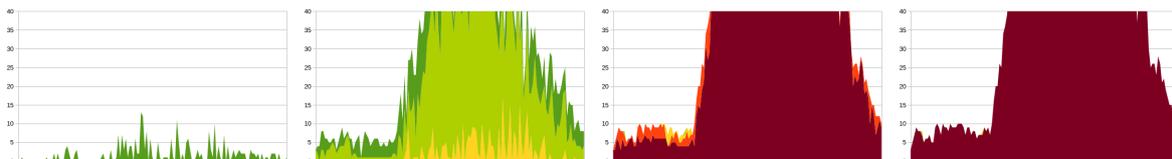


Figure 77: Posterior Guesses for Worst-Case DP, $\delta = 90\%$, $\epsilon \in \{0.1, 1.0, 5.0, 10.0\}$ (Left to Right)

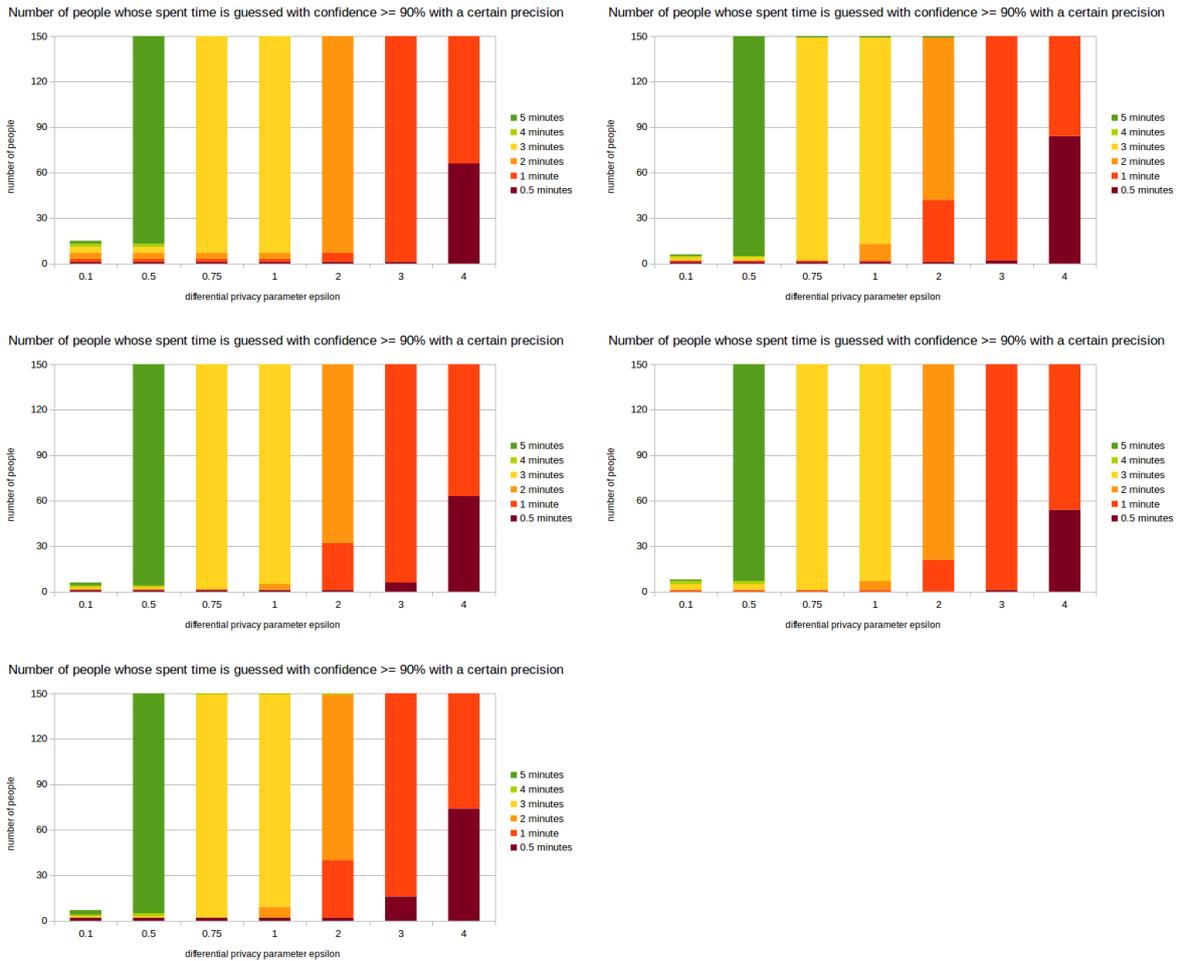


Figure 78: Precisions of Guessing a Particular Spent Time with Confidence $t = 0.9$ for Different ϵ Values and 5 Different Datasets

4.6 Secure Multiparty Computation in NAPLES

4.6.1 Deliverable Outcomes. The main deliverable outcomes of the work carried out under this sub-task are:

- Solution analysis and design document - provides an overview of any analysis carried out (including use case specification) and details the prototype implementation specifics.
- Prototype implementation, including:
 - Final easily testable product containing the respective Docker image. When run according to the usage instructions the respective Docker container should start generation of the synthetic log entries, runs the non-privacy preserving version of the algorithm, then runs the privacy-preserving version of the algorithm in SecreC using the Sharemind emulator, and then compares the results.
 - Source code of prototype implementation.
 - User and install guides for the prototype are included in code delivery as files INSTALL.md and USAGE.md.
- Final reporting (this document) to provide an overview of work carried out.

Handover of these results to the Cyber PA project will be discussed during Q1 of 2021.

4.6.2 Analysis and Prototyping Main Outcomes. An important outcome from prototyping (in addition to a working prototype and associated learnings) is the proposed event correlation approach / algorithm. In principle, the prototype event correlation algorithm attempts to construct chains of (correlated) log events by iterating over the input log entry values in timestamp order. If the value, iterated over, already matches a value of the last log entry in an existing chain, the log entry will be appended to the existing chain. Otherwise, a new chain is created. As output from the algorithm, each data owner receives log entries from detected chains which contain at least one log entry from the respective data owner. Chains with only one log entry are not considered for output.

In addition to using Pleak to model processes within the Cyber PA project, the event log correlation process was also modelled in Pleak. The model in Figure 79 visualizes the stakeholder involved in the exercise, their activities and the data elements that are used and / or generated during the processing. In addition, the scheme specifies security methods for data flows and an overview of the visibility of data elements to the stakeholders is provided. Finally, Leakage detection analysis was run on the model to give a more thorough overview of the process security guarantees).

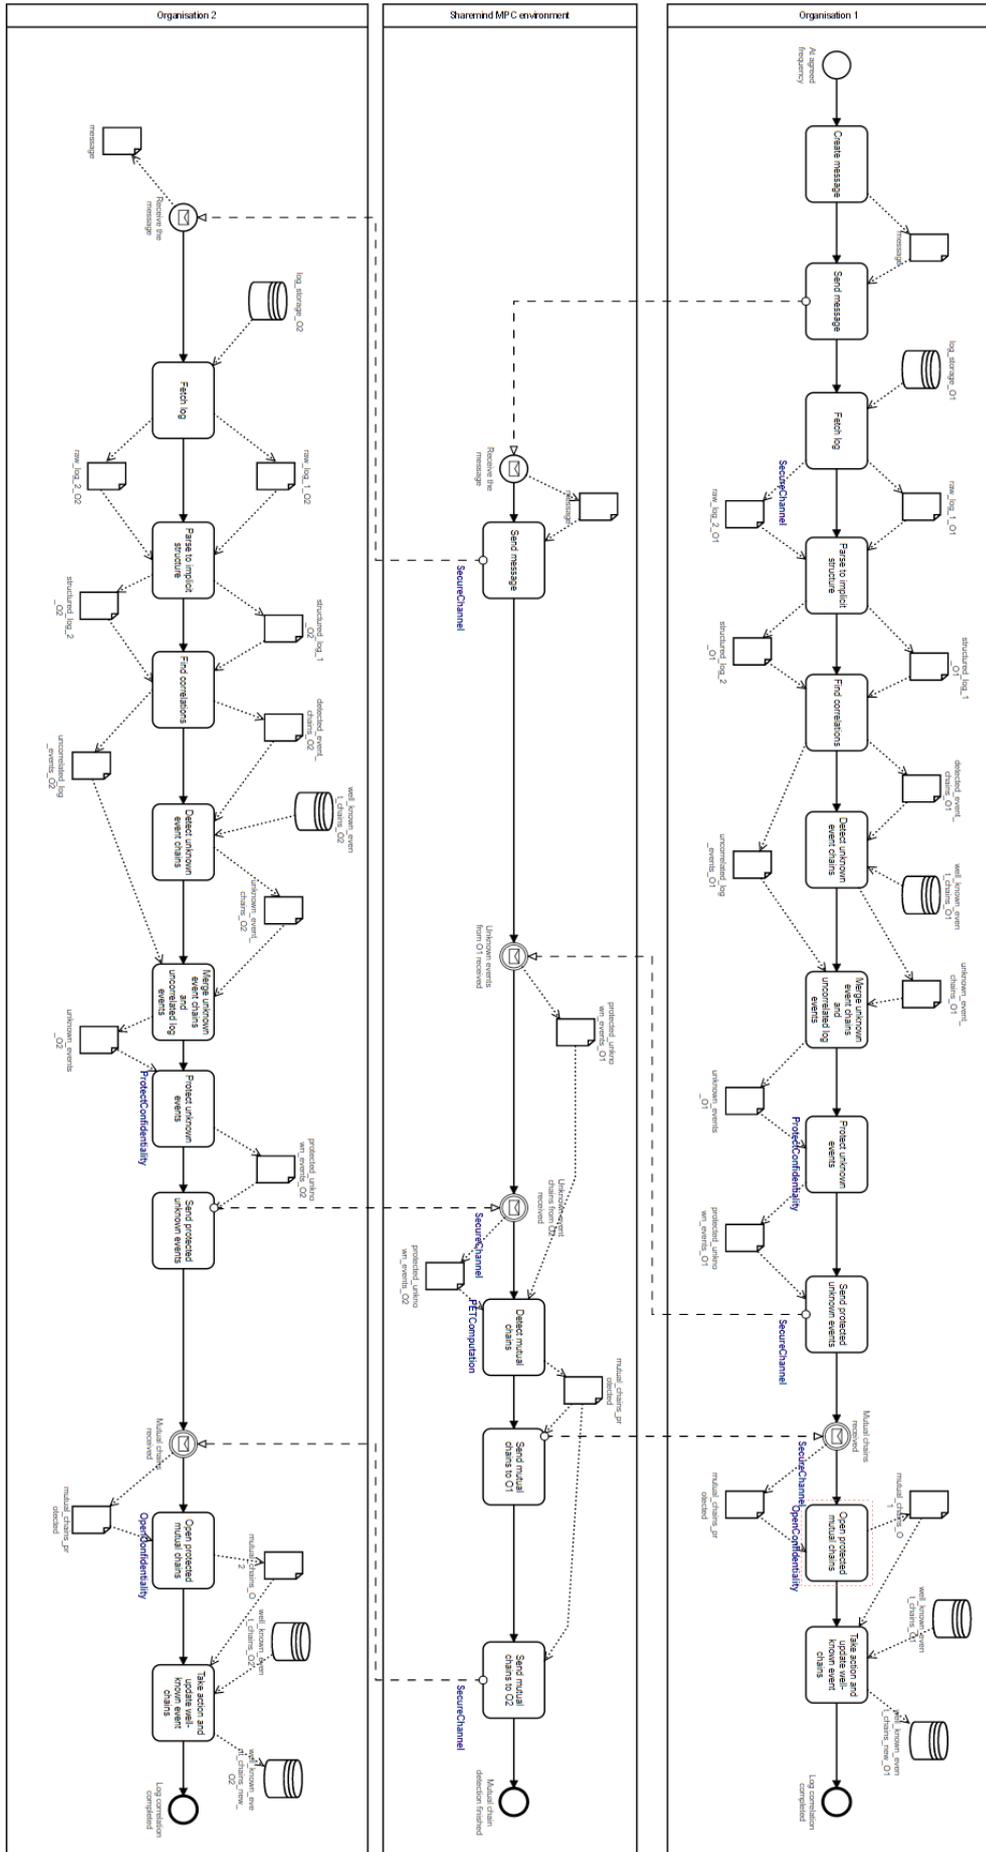


Figure 79: Event Log Correlation Process in PE-BPMN Notation

Continues from previous page

Name	Actor	Description	Trigger	Goal	Tool
Protect unknown events	ORG1/2	Unknown events are secret shared for MPC environment.	File with pre-agreed structure of unknown events created.	Protected unknown event file(s) created. NOTE: Protected here means the data file is split using secret sharing technique into 3 shares where none of the shares contain any meaningful information separately.	Sharemind CSV Importer data upload tool.
Send protected unknown events	ORG1/2	The shares of protected unknown event files are uploaded to the MPC instances.	Protected unknown event file(s) created.	Protected unknown event file uploaded to MPC environment.	Sharemind CSV Importer data upload tool.
Detect mutual chains	MPC	The unknown events of all participating organizations are gathered and similar correlation activities described in Find Correlations activity are used to detect mutual chain events.	Protected unknown event file uploaded to MPC environment.	Mutual chain events dataset created.	Sharemind MPC servers running pre-agreed SecreC code.
Send mutual chains to O1/O2	MPC	Extract from the mutual chain events dataset chains containing events contributed by ORG1/2 and make them available to ORG1/2. NOTE: In case more than 2 organizations participating in the exercise only the chain events containing events contributed by the organization are made available.	Mutual chain events dataset created.	Mutual chain events published to the ORG1/2.	Custom Sharemind MPC client application.

Continued on next page

Continues from previous page

Name	Actor	Description	Trigger	Goal	Tool
Open protected mutual chains	ORG1/2	Open the mutual chain event records for processing. NOTE: Send and Open mutual chains to O1/O2 events in real life can be considered as one as same client application is most likely used for making the query. The method is merely used to fully benefit the Pleak tool.	Mutual chain events published to the ORG1/2.	Mutual chain events available for further processing for ORG1/2.	Custom Sharemind MPC client application.
Take action and update well-known event chains	ORG1/2	Take appropriate action based on what is learned from the exercise and update the known database accordingly	Mutual chain events available for further processing for ORG1/2.	Required actions taken and well-known event chains database updated.	Tools dependent on organization internal tools.

4.6.2.3 Data Elements. Data elements are information or data artefacts created or used during the process. They are either needed as inputs for activities or output of activities. In table 30 you find Name of the element, role that created and holds the element (Owner) and short description. Same data element can be used in multiple places, data element created during one activity is very often an input for the other.

Table 30: Event Log Correlation: Overview of Data Elements

Name	Owner	Description
message	ORG1	Simple text message for informing the beginning of the exercise.
log_storage_O1/O2	ORG1/2	Log storage of the system of an organization that the organization wished to use in the exercise.
raw_log_1_O1/O2, raw_log_2_O1/O2	ORG1/2	Log file in its initial form after extracting it from systems log storage.
structured_log_1_O1/O2, structured_log_2_O1/O2	ORG1/2	Log file parsed to the implicit structure for the log correlation activity. Each organization might have different structure and can convert the file to the mutual structure later, but also it can be already using the agreed structure.
detected_event_chains_O1/O2	ORG1/2	Result of correlation detection activity where event entries are marked with common identifier that represents membership to a chain. The events are ordered sequentially by time of appearance.
uncorrelated_log_events_O1/O2	ORG1/2	List of structured log events not belonging to any chain. It is structured the same way as structured_log_1_O1/O2
well_known_event_chains_O1/O2	ORG1/2	A database of known event chains of the organization. Structured the same way as detected_event_chains_O1/O2 so that it would be easy to query whether the detected chains match to know chains. Database is updated regularly by the organization once it detects new event chains. Also, the update takes place during the exercise.

Continued on next page

Continues from previous page

Name	Owner	Description
unknown_event_chains_O1/O2	ORG1/2	List of detected chain events that did not match to any of the known events recorded in the well_known_event_chains_O1/O2 database. Structured same way as detected_event_chains_O1/O2.
unknown_events_O1/O2	ORG1/2	List of events that could not be assigned to any chain during the correlation process. Structure is pre-agreed by stakeholders and is most likely a CSV file as this is the form required by Sharemind CSV Importer.
protected_unknown_events_O1/O2	ORG1/2	A secret-shared copy of unknown_events_O1/O2. Sharemind CSV Importer tool secret shares the unknown_events file into 3 shares where none of the shares contain any meaningful information separately but is processable as normal unknown_events file in Sharemind MPC environment. NB! In case MPC instances collude, the information is likely extractable from the shares.
mutual_chains_protected	MPC	The result of Sharemind MPC mutual chain detection activity where events of all the contributing organizations are correlated into mutual chain events. It is structured in the pre-agreed format so that the organizations can clearly understand and process the information. Still the data is in the protected (i.e., secret-shared) form and each share is controlled separately by one of the non-colluding MPC parties.
mutual_chains_O1/O2	ORG1/2	Subset of mutual_chains_protected, where only chains containing events contributed by the organization are present. Also, the data is now in the open form and fully readable and processable to the organization. To emphasize the organization sees only chains that contains events it contributed. In case there are more than 2 organizations involved in the exercise the events not containing any events contributed by organization X are not included. The data is structured the same way as mutual_chains_protected.
well_known_event_chains_new_O1/O2	ORG1/2	The same data element as well_known_event_chains_O1/O2 except containing updates of findings of the exercise.

4.6.2.4 Security Measures. Security methods used in the process, highlighted in blue in the process model:

- **SecureChannel** - Data transmission marked with SecureChannel in the PE-BPMN figure means that an authenticated and encrypted security channel is in use. Technically, we expect the use of standard solutions (e.g. TLS, Transport Layer Security) with up-to-date technical solutions. The communication channel set up in this way does not leak data to the communication service provider or other parties involved in the data transmission.
- **PETComputation & ProtectConfidentiality** - Data protection mechanisms based on the capabilities of the Sharemind MPC technology.
- **OpenConfidentiality** - OpenConfidentiality is an opposite action to ProtectConfidentiality and is also a data protection mechanism based on the capabilities of the Sharemind MPC technology.

4.6.2.5 Disclosure & Leaks-When Analysis. Using Pleak tools & capabilities on the process model Fig. 79 allows to generate reports that provide an overview of data elements by stakeholders and help detect / ensure data is secured.

The Leaks-When analysis result in Figure 80 shows that the unknown events are computed by the organizations on their own and are only depending on their own input data. However, the new chains are the ones computed together and are also affected by the logs of the other party.

	log_storage_O1.log	log_storage_O2.log	well_known_event_chains_O1.chains	well_known_event_chains_O2.chains
unknown_events_O1.data	if	never	if	never
unknown_events_O2.data	never	if	never	if
well_known_event_chains_new_O1.chains	if	if	always	if
well_known_event_chains_new_O2.chains	if	if	if	always

Figure 80: Leaks-When Analysis Result for the Prototype Model

The disclosure analysis result in Table 31 provides an overview of the visibility of data elements by stakeholder. Simple disclosure report is a table which lists all stakeholders and shows, whether and how the stakeholder sees a data element. Each cell is marked either V (visible), H (hidden) or -. Marking “-“ means that the stakeholder does not see the data element in any way in the process. V means that contents of the data element are fully visible. H means that the stakeholder has the data element but this element is protected with security measures. For example, the Sharemind MPC environment has the data element containing information of mutual chains, but as it is protected, the content is not visible. In addition, cell can be marked with O, meaning that the stakeholder is the creator/owner of the data element and it is fully visible to the stakeholder. The table helps to understand what information is available to which parties.

In general, the reports confirm that neither the events nor the event chains of the contributing organizations are revealed to any other stakeholder and thus no confidential data is leaked to any unwanted participants. The Sharemind MPC hosts only see data in hidden form as shares of secret sharing.

4.6.3 Main Results. Joint situational awareness is built up from national situational awareness which in turn is built up from information gathered from local institutions and private enterprises. Local situational awareness is the result of the analytical work of the cyber security professionals. Many or even most private enterprises do not have cyber security specialists available. One way for such enterprises to participate in creation of the situational awareness is to share system logs from their computer and networking equipment with institutions with better analyzing capabilities.

As per analysis it was proposed to create an IT system log analysing engine which enables sharing potentially sensitive information in a privacy-preserving way, so that

- participating organizations would have a clean, easy, meaningful way to decide which information to share.
- participants would get strong guarantees on how the information is processed and what part in what aggregation level is published to other participating parties.

Table 31: Event Log Correlation: Simple Disclosure Report

Data element	ORG 1	ORG 2	MPC	Shared Over
detected_event_chains_O1	V	-	-	-
detected_event_chains_O2	-	V	-	-
log_storage_O1	O	-	-	-
log_storage_O2	-	O	-	-
message	V	V	V	S
mutual_chains_O1, mutual_chains_O2, mutual_chains_protected	V	V	H	S
protected_unknown_events_O1, unknown_events_O1	V	-	H	S
protected_unknown_events_O2, unknown_events_O2	-	V	H	S
raw_log_1_O1	V	-	-	-
raw_log_1_O2	-	V	-	-
raw_log_2_O1	V	-	-	-
raw_log_2_O2	-	V	-	-
structured_log_1_O1	V	-	-	-
structured_log_1_O2	-	V	-	-
structured_log_2_O1	V	-	-	-
structured_log_2_O2	-	V	-	-
uncorrelated_log_events_O1	V	-	-	-
uncorrelated_log_events_O2	-	V	-	-
unknown_event_chains_O1	V	-	-	-
unknown_event_chains_O2	-	V	-	-
well_known_event_chains_new_O1	V	-	-	-
well_known_event_chains_new_O2	-	V	-	-
well_known_event_chains_O1	O	-	-	-
well_known_event_chains_O2	-	O	-	-

V = visible, H = hidden, O = owner, S = SecureChannel

A system having the above properties was analyzed, modelled and prototyped as a part of the work task. Main findings from the prototyping activity were:

- The approaches taken give organizations ability to make clear distinction between mundane, possibly confidential every-day events and unusual, possibly harmful events which should be analyzed further. This methodology should be useful even without following information sharing and joint analysis. Therefore, it should increase the will and likelihood for enterprises to join the Cyber PA project information gathering network.
- Guarantees on how shared information is processed and revealed are provided by
 - implementing the correlation algorithm in Sharemind MPC.
 - carrying out Pleak analysis on modelled processes (Data Leakage analysis)
- Prototype benchmarks show overall asymptotic complexity $O(n^3)$ for chains detections algorithm, where n is the number of input system log records, whilst the non privacy-preserving version of chains detection has a complexity of $O(n \log n)$. This because the privacy preserving version is much more complex as it hides array access patterns, which when observed, might leak content or content distribution of the array.

For the purposes of the Cyber PA project the following prototype enhancement possibilities should be discussed, analyzed and tested:

- Improve the algorithm in a way that it does not need to look up things from arrays, thus not adding asymptotic complexity.
- Give up some insignificant amount of the privacy to significantly reduce complexity. This trade-off should be well thought-through and documented.
- Improve Sharemind MPC so required calculations do not have so much overhead.
- Consider some other privacy preserving technologies with less computational overhead, taking into account the requirements which led to preliminary implementation on Sharemind MPC.

In conclusion, results from this task are a valuable addition to the Cyber PA project. Modelling carried out in Pleak helped to further understand the designed processes in the start of prototype implementation and further helped verify privacy guarantees of the Solution. Actual prototyping of the solution provided useful feedback on the usability of the prototype in a potential Cyber PA Solution framework. All this provides an improved starting point of enabling such functionalities as a part of the Cyber PA in order to improve joint situational awareness. Additionally, the usability of both the Sharemind MPC and Pleak technologies were verified as a part of work carried out and feedback provided for both.

5 CONCLUSIONS

Our work in NAPLES has led to a wide array of modeling languages and privacy analyses for different kinds of business processes, and related executable formalisms. Our Privacy-Enhanced BPMN has been a useful modeling tool for systems with significant and complex privacy implications. Our qualitative and quantitative privacy definitions have helped to explain the nature of these implications, and our analyses have been able to precisely state, how much a system may be leaking. Regarding the definitions, a particularly important contribution has been that of *derivative sensitivity*, which clearly shows how the privacy definition is part of the privacy policy of a system, and how the latter always has to state the combinations of, and trade-offs between the leakages in various components of the data.

Almost all of our analyses have been integrated into the Pleak toolset, giving them a uniform interface for modeling the processes, specifying the privacy policies and other input parameters of the analysis, and studying the output of the analysis. The business processes can be implemented either in the PE-BPMN notation, or in plain BPMN with annotations on tasks and message flows. Pleak allows its users to inspect and analyze privacy leakages at three levels of abstraction. The first level of abstraction (the so-called “Boolean” level) allows us to see which input data sources are directly or indirectly disclosed to each party in the process. The second level (the so-called “Leaks-when” level) shows under what conditions each disclosure occurs, and which specific attributes (or functions over the input attributes) are disclosed. The third level (the “quantitative” level) allows one to measure the extent to which the disclosure reveals information about individual items (e.g. individual rows) of the input data sources. This third level also allows us to quantify to what extent the disclosure of the outputs of the process increase the probability that an attacker can guess the value of an individual item in the data source, relative to a prior probability capturing generally available knowledge about these data sources. This “attacker’s guessing advantage” model also takes into account the effect of noise added to the input data in order to achieve a given differential privacy level (cf. the ϵ parameter in differential privacy).

The Pleak toolset has been validated in the DARPA Brandeis program, modeling and analysing the systems proposed and studied by the three Collaborative Research Teams. This validation provides initial evidence that the toolset can be used to analyze realistic processes. However, the ecological validity of this validation effort is limited due to the low number of processes analyzed, and by the fact that the processes have been modeled and analyzed by the same team that developed the toolset itself (even though the processes were scoped and designed by other teams in the Brandeis project). A direction for future work is to supplement this validation with usability evaluations involving business process analysts, as well as case studies conducted by independent research teams.

In its current implementation, the Pleak toolset supports the quantitative analysis of privacy-enhanced business processes in which the computations are specified in the SQL query language. In practice though, computations may be specified in various programming languages. Another avenue for future work is to extend the set of techniques in the Pleak toolset with program analysis techniques that would allow it to handle other languages.

The Pleak toolset is also limited in terms of the range of PETs it supports. In particular, the toolset does not integrate various metrics from the field of statistical disclosure control [74], including k-anonymity and l-diversity, as well as data masking techniques such as microaggregation and data swapping. Extending Pleak in order to support a wider range of PETs is another direction for future work, which we plan to pursue as a part of technology transfer efforts.

BIBLIOGRAPHY

- [1] von Scheel, Henrik; von Rosing, Mark; Fonseca, Marianne; Hove, Maria; and Foldager, Ulrik, “Phase 1: Process Concept Evolution,” in von Rosing, Mark; von Scheel, Henrik; and Scheer, August-Wilhelm, editors, *The Complete Business Process Handbook: Body of Knowledge from Process Modeling to BPM*, Morgan Kaufmann, volume 1, 2014, pp. 1–10.
- [2] OMG, “Business Process Model and Notation (BPMN),” URL: <http://www.omg.org/spec/BPMN/2.0/>.
- [3] Pullonen, Pille; Matulevičius, Raimundas; and Bogdanov, Dan, “PE-BPMN: privacy-enhanced business process model and notation,” in *International Conference on Business Process Management*, Springer, pp. 40–56.
- [4] Pullonen, Pille; Tom, Jake; Matulevičius, Raimundas; and Toots, Aivo, “Privacy-enhanced BPMN: enabling data privacy analysis in business processes models,” *Software and Systems Modeling*, **18**(6):(2019), pp. 3235–3264, URL <https://doi.org/10.1007/s10270-019-00718-z>.
- [5] Danezis, G.; Domingo-Ferrer, J.; Hansen, M.; Hoepman, J.-H.; Metayer, D. L.; Tirtea, R.; and Schiffner, S., “Privacy and Data Protection by Design—from Policy to Engineering,” Technical report, European Union Agency for Network and Information Security, 2015.
- [6] Heurix, Johannes; Zimmermann, Peter; Neubauer, Thomas; and Fenz, Stefan, “A taxonomy for privacy enhancing technologies,” *Computers & Security*, **53**:(2015), pp. 1–17.
- [7] Clements, Tim, “GDPR data flow mapping - an approach,” , 2017, <https://www.linkedin.com/pulse/gdpr-data-flow-mapping-approach-tim>.
- [8] Archer, David W.; Bogdanov, Dan; Pinkas, Benny; and Pullonen, Pille, “Maturity and Performance of Programmable Secure Computation,” *IEEE Security and Privacy*, **14**(5):(2016), pp. 48–56.
- [9] Shamir, Adi, “How to Share a Secret,” *Communications of the ACM*, **22**(11):(1979), pp. 612–613, URL <http://doi.acm.org/10.1145/359168.359176>.
- [10] Blakley, George Robert, “Safeguarding Cryptographic Keys,” in *Proceedings of the 1979 AFIPS National Computer Conference*, AFIPS Press, pp. 313–317.
- [11] Gilboa, Niv and Ishai, Yuval, “Distributed Point Functions and Their Applications,” in *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pp. 640–658, URL http://dx.doi.org/10.1007/978-3-642-55220-5_35.
- [12] Boyle, Elette; Gilboa, Niv; and Ishai, Yuval, “Function Secret Sharing,” in *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, pp. 337–367, URL http://dx.doi.org/10.1007/978-3-662-46803-6_12.
- [13] Boyle, Elette; Gilboa, Niv; and Ishai, Yuval, “Breaking the Circuit Size Barrier for Secure Computation Under DDH,” in *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I*, pp. 509–539, URL http://dx.doi.org/10.1007/978-3-662-53018-4_19.

- [14] Boyle, Elette; Gilboa, Niv; and Ishai, Yuval, “Function Secret Sharing: Improvements and Extensions,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pp. 1292–1303, URL <http://doi.acm.org/10.1145/2976749.2978429>.
- [15] Shannon, Claude E, “Communication theory of secrecy systems,” *Bell system technical journal*, **28**(4):(1949), pp. 656–715.
- [16] Diffie, W. and Hellman, M., “New Directions in Cryptography,” *IEEE Trans. Inf. Theor.*, **22**(6):(2006), pp. 644–654, URL <http://dx.doi.org/10.1109/TIT.1976.1055638>.
- [17] Gentry, Craig, “Fully Homomorphic Encryption Using Ideal Lattices,” in *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, ACM, New York, NY, USA, STOC ’09, pp. 169–178, URL <http://doi.acm.org/10.1145/1536414.1536440>.
- [18] Paillier, Pascal, “Public-key Cryptosystems Based on Composite Degree Residuosity Classes,” in *Proceedings of the 17th International Conference on Theory and Application of Cryptographic Techniques*, Springer-Verlag, Berlin, Heidelberg, EUROCRYPT’99, pp. 223–238, URL <http://dl.acm.org/citation.cfm?id=1756123.1756146>.
- [19] Anati, Ittai; Gueron, Shay; Johnson, Simon; and Scarlata, Vincent, “Innovative technology for CPU based attestation and sealing,” Technical report, Intel Corporation, 8 2013, presented at the Second Workshop on Hardware and Architectural Support for Security and Privacy, HASP 2013.
- [20] “Intel Software Guard Extensions (Intel SGX),” , March 2017, <https://software.intel.com/en-us/sgx>.
- [21] Johnson, Simon; Scarlata, Vincent; Rozas, Carlos; Brickell, Ernie; and Mckeen, Frank, “Intel software guard extensions: EPID provisioning and attestation services,” *White Paper*, <https://software.intel.com/en-us/blogs/2016/03/09/intel-sgx-epid-provisioning-and-attestation-services>.
- [22] “Intel Software Guard Extensions Remote Attestation End-to-End Example,” , July 2016, <https://software.intel.com/en-us/articles/intel-software-guard-extensions-remote-attestation-end-to-end-example>.
- [23] Belluccini, Sara; Nicola, Rocco De; Dumas, Marlon; Pullonen, Pille; Re, Barbara; ; and Tiezzi, Francesco, “Verification of privacy-enhanced collaborations,” in *FormalISE@ICSE 2020: 8th International Conference on Formal Methods in Software Engineering, Seoul, Republic of Korea, July 13, 2020*, pp. 141–152.
- [24] Bergstra, Jan A and Klop, Jan Willem, “Process algebra for synchronous communication,” *Information and control*, **60**(1-3):(1984), pp. 109–137.
- [25] Bunte, Olav; Groote, Jan Friso; Keiren, Jeroen JA; Laveaux, Maurice; Neele, Thomas; de Vink, Erik P; Wesselink, Wieger; Wijs, Anton; and Willemse, Tim AC, “The mCRL2 Toolset for Analysing Concurrent Systems,” in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, Springer, pp. 21–39.
- [26] Groote, Jan Friso and Mousavi, Mohammad Reza, *Modeling and analysis of communicating systems*, MIT press, 2014.
- [27] Tšahhirov, Ilja and Laud, Peeter, “Application of Dependency Graphs to Security Protocol Analysis,” in Barthe, Gilles and Fournet, Cédric, editors, *Trustworthy Global Computing, Third Symposium, TGC 2007, Sophia-Antipolis, France, November 5-6, 2007, Revised Selected Papers*, Springer, volume 4912 of *Lecture Notes in Computer Science*, pp. 294–311, URL https://doi.org/10.1007/978-3-540-78663-4_20.

- [28] Laud, Peeter and Tšahhirov, Ilja, “A User Interface for a Game-Based Protocol Verification Tool,” in Degano, Pierpaolo and Guttman, Joshua D., editors, *Formal Aspects in Security and Trust, 6th International Workshop, FAST 2009, Eindhoven, The Netherlands, November 5-6, 2009, Revised Selected Papers*, Springer, volume 5983 of *Lecture Notes in Computer Science*, pp. 263–278, URL https://doi.org/10.1007/978-3-642-12459-4_19.
- [29] Dijkman, Remco M.; Dumas, Marlon; and Ouyang, Chun, “Semantics and analysis of business process models in BPMN,” *Information & Software Technology*, **50**(12):(2008), pp. 1281–1294, URL <https://doi.org/10.1016/j.infsof.2008.02.006>.
- [30] Dumas, Marlon; García-Bañuelos, Luciano; and Laud, Peeter, “Disclosure Analysis of SQL Workflows,” in Cybenko, George; Pym, David J.; and Fila, Barbara, editors, *5th International Workshop on Graphical Models for Security, held in conjunction with the Federated Logic Conference (FLoC) 2018, GramSec@FLoC 2018, Oxford, UK, July 8, 2018, Revised Selected Papers*, Springer, volume 11086 of *Lecture Notes in Computer Science*, pp. 51–70, URL https://doi.org/10.1007/978-3-030-15465-3_4.
- [31] Esparza, Javier; Römer, Stefan; and Vogler, Walter, “An Improvement of McMillan’s Unfolding Algorithm,” *Formal Methods in System Design*, **20**(3):(2002), pp. 285–310, URL <https://doi.org/10.1023/A:1014746130920>.
- [32] Armas-Cervantes, Abel; Baldan, Paolo; Dumas, Marlon; and García-Bañuelos, Luciano, “Diagnosing behavioral differences between business process models: An approach based on event structures,” *Inf. Syst.*, **56**:(2016), pp. 304–325, URL <https://doi.org/10.1016/j.is.2015.09.009>.
- [33] Milner, Robin, *Communicating and mobile systems - the Pi-calculus*, Cambridge University Press, 1999.
- [34] Fournet, Cédric and Gonthier, Georges, “The Reflexive CHAM and the Join-Calculus,” in Boehm, Hans-Juergen and Jr., Guy L. Steele, editors, *Conference Record of POPL’96: The 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the Symposium, St. Petersburg Beach, Florida, USA, January 21-24, 1996*, ACM Press, pp. 372–385, URL <http://doi.acm.org/10.1145/237721.237805>.
- [35] Laud, Peeter and Randmets, Jaak, “A Domain-Specific Language for Low-Level Secure Multiparty Computation Protocols,” in Ray, Indrajit; Li, Ninghui; and Kruegel, Christopher, editors, *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, ACM, pp. 1492–1503, URL <http://doi.acm.org/10.1145/2810103.2813664>.
- [36] Dwork, Cynthia, “Differential Privacy,” in Bugliesi, Michele; Preneel, Bart; Sassone, Vladimiro; and Wegener, Ingo, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, Springer, volume 4052 of *Lecture Notes in Computer Science*, pp. 1–12, URL http://dx.doi.org/10.1007/11787006_1.
- [37] Chatzikokolakis, Konstantinos; Andrés, Miguel E.; Bordenabe, Nicolás Emilio; and Palamidessi, Catuscia, “Broadening the Scope of Differential Privacy Using Metrics,” in Cristofaro, Emiliano De and Wright, Matthew, editors, *Privacy Enhancing Technologies - 13th International Symposium, PETS 2013, Bloomington, IN, USA, July 10-12, 2013. Proceedings*, Springer, volume 7981 of *Lecture Notes in Computer Science*, pp. 82–102, URL http://dx.doi.org/10.1007/978-3-642-39077-7_5.
- [38] Ebadi, Hamid; Sands, David; and Schneider, Gerardo, “Differential Privacy: Now it’s Getting Personal,” in Rajamani, Sriram K. and Walker, David, editors, *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai*,

India, January 15-17, 2015, ACM, pp. 69–81, URL <http://doi.acm.org/10.1145/2676726.2677005>.

- [39] Cousot, Patrick and Cousot, Radhia, “Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints,” in Graham, Robert M.; Harrison, Michael A.; and Sethi, Ravi, editors, *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*, ACM, pp. 238–252, URL <http://doi.acm.org/10.1145/512950.512973>.
- [40] Barthe, Gilles; D’Argenio, Pedro R.; and Rezk, Tamara, “Secure Information Flow by Self-Composition,” in *17th IEEE Computer Security Foundations Workshop, (CSFW-17 2004), 28-30 June 2004, Pacific Grove, CA, USA*, IEEE Computer Society, pp. 100–114, URL <http://doi.ieeecomputersociety.org/10.1109/CSFW.2004.17>.
- [41] Bogdanov, Dan; Laud, Peeter; and Randmets, Jaak, “Domain-Polymorphic Programming of Privacy-Preserving Applications,” in Russo, Alejandro and Tripp, Omer, editors, *Proceedings of the Ninth Workshop on Programming Languages and Analysis for Security, PLAS@ECOOP 2014, Uppsala, Sweden, July 29, 2014*, ACM, pp. 53–65, URL <http://doi.acm.org/10.1145/2637113.2637119>.
- [42] Bogdanov, Dan; Kamm, Liina; Kubo, Baldur; Rebane, Reimo; Sokk, Ville; and Talviste, Riivo, “Students and Taxes: a Privacy-Preserving Study Using Secure Computation,” *PoPETS*, **2016**(3):(2016), pp. 117–135, URL <http://www.degruyter.com/view/j/popets.2016.2016.issue-3/popets-2016-0019/popets-2016-0019.xml>.
- [43] Bogdanov, Dan; Niitsoo, Margus; Toft, Tomas; and Willemson, Jan, “High-performance secure multi-party computation for data mining applications,” *Int. J. Inf. Sec.*, **11**(6):(2012), pp. 403–418, URL <http://dx.doi.org/10.1007/s10207-012-0177-2>.
- [44] Karr, Michael, “Affine Relationships Among Variables of a Program,” *Acta Inf.*, **6**:(1976), pp. 133–151, URL <http://dx.doi.org/10.1007/BF00268497>.
- [45] Cousot, Patrick and Halbwachs, Nicolas, “Automatic Discovery of Linear Restraints Among Variables of a Program,” in Aho, Alfred V.; Zilles, Stephen N.; and Szymanski, Thomas G., editors, *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages, Tucson, Arizona, USA, January 1978*, ACM Press, pp. 84–96, URL <http://doi.acm.org/10.1145/512760.512770>.
- [46] Granger, Philippe, “Static Analysis of Linear Congruence Equalities among Variables of a Program,” in Abramsky, Samson and Maibaum, T. S. E., editors, *TAPSOFT’91: Proceedings of the International Joint Conference on Theory and Practice of Software Development, Brighton, UK, April 8-12, 1991, Volume 1: Colloquium on Trees in Algebra and Programming (CAAP’91)*, Springer, volume 493 of *Lecture Notes in Computer Science*, pp. 169–192, URL http://dx.doi.org/10.1007/3-540-53982-4_10.
- [47] Laud, Peeter, *Computationally Secure Information Flow*, Ph.D. thesis, University of Saarland, 2002.
- [48] Gover, Thomas M. and Thomas, Joy A., *Elements of Information Theory*, John Wiley & Sons, 2nd edition, 2006.
- [49] Cuff, Paul and Yu, Lanqing, “Differential Privacy as a Mutual Information Constraint,” in Weippl, Edgar R.; Katzenbeisser, Stefan; Kruegel, Christopher; Myers, Andrew C.; and Halevi, Shai, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, ACM, pp. 43–54, URL <https://doi.org/10.1145/2976749.2978308>.

- [50] Alvim, Mário S.; Andrés, Miguel E.; Chatzikokolakis, Konstantinos; Degano, Pierpaolo; and Palamidessi, Catuscia, “On the information leakage of differentially-private mechanisms,” *Journal of Computer Security*, **23**(4):(2015), pp. 427–469, URL <http://dx.doi.org/10.3233/JCS-150528>.
- [51] Kopperman, Ralph, “All Topologies Come From Generalized Metrics,” *The American Mathematical Monthly*, **95**(2):(1988), pp. 89–97, URL <http://www.jstor.org/stable/2323060>.
- [52] Nissim, Kobbi; Raskhodnikova, Sofya; and Smith, Adam D., “Smooth sensitivity and sampling in private data analysis,” in Johnson, David S. and Feige, Uriel, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, ACM, pp. 75–84, URL <http://doi.acm.org/10.1145/1250790.1250803>.
- [53] Guagliardo, Paolo and Libkin, Leonid, “A Formal Semantics of SQL Queries, Its Validation, and Applications,” *Proceedings of the VLDB Endowment (PVLDB)*.
- [54] McSherry, Frank, “Privacy integrated queries: an extensible platform for privacy-preserving data analysis,” in Çetintemel, Ugur; Zdonik, Stanley B.; Kossman, Donald; and Tatbul, Nesime, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2009, Providence, Rhode Island, USA, June 29 - July 2, 2009*, ACM, pp. 19–30, URL <http://doi.acm.org/10.1145/1559845.1559850>.
- [55] Baggett, L.W., *Functional Analysis: A Primer*, Chapman & Hall Pure and Applied Mathematics, Taylor & Francis, 1991, URL <https://books.google.ee/books?id=RKFiQgAACAAJ>.
- [56] He, Xi; Machanavajjhala, Ashwin; and Ding, Bolin, “Blowfish privacy: tuning privacy-utility trade-offs using policies,” in Dyreson, Curtis E.; Li, Feifei; and Özsu, M. Tamer, editors, *International Conference on Management of Data, SIGMOD 2014, Snowbird, UT, USA, June 22-27, 2014*, ACM, pp. 1447–1458, URL <http://doi.acm.org/10.1145/2588555.2588581>.
- [57] Geng, Quan; Ding, Wei; Guo, Ruiqi; and Kumar, Sanjiv, “Truncated Laplacian Mechanism for Approximate Differential Privacy,” *CoRR*, **abs/1810.00877**, URL <http://arxiv.org/abs/1810.00877>.
- [58] Lee, Jaewoo and Clifton, Chris, “How Much Is Enough? Choosing ϵ for Differential Privacy,” in Lai, Xuejia; Zhou, Jianying; and Li, Hui, editors, *Information Security, 14th International Conference, ISC 2011, Xi’an, China, October 26-29, 2011. Proceedings*, Springer, volume 7001 of *Lecture Notes in Computer Science*, pp. 325–340, URL https://doi.org/10.1007/978-3-642-24861-0_22.
- [59] Kifer, Daniel and Machanavajjhala, Ashwin, “Pufferfish: A Framework for Mathematical Privacy Definitions,” *ACM Trans. Database Syst.*, **39**(1):(2014), pp. 3:1–3:36, URL <http://doi.acm.org/10.1145/2514689>.
- [60] Oh, Sewoong and Viswanath, Pramod, “The Composition Theorem for Differential Privacy,” *CoRR*, **abs/1311.0776**, URL <http://arxiv.org/abs/1311.0776>.
- [61] Sommer, David M.; Meiser, Sebastian; and Mohammadi, Esfandiar, “Privacy Loss Classes: The Central Limit Theorem in Differential Privacy,” *PoPETs*, **2019**(2):(2019), pp. 245–269, URL <https://doi.org/10.2478/popets-2019-0029>.
- [62] Balle, Borja and Wang, Yu-Xiang, “Improving the Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising,” in Dy, Jennifer G. and Krause, Andreas, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, PMLR, volume 80 of *Proceedings of Machine Learning Research*, pp. 403–412, URL <http://proceedings.mlr.press/v80/balle18a.html>.

- [63] Stolfi, Jorge and de Figueiredo, Luiz Henrique, “An introduction to affine arithmetic,” *Trends in Applied and Computational Mathematics*, 4(3):(2003), pp. 297–312.
- [64] Stolfi, Jorge and Figueiredo, Luiz Henrique De, “Self-validated numerical methods and applications,” , 1997, set of course notes for the 21st Brazilian Mathematics Colloquium, IMPA, Rio de Janeiro.
- [65] Gay, Olivier; Coeurjolly, David; and Hurst, Nathan, “Libaffa-C++ affine arithmetic library for GNU/Linux,” , 2006.
- [66] Fan, Liyue and Xiong, Li, “Adaptively Sharing Time-Series with Differential Privacy,” *CoRR*, abs/1202.3461, URL <http://arxiv.org/abs/1202.3461>.
- [67] Dwork, Cynthia; Naor, Moni; Pitassi, Toniann; and Rothblum, Guy N., “Differential privacy under continual observation,” in Schulman, Leonard J., editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, ACM, pp. 715–724, URL <https://doi.org/10.1145/1806689.1806787>.
- [68] Toots, Aivo; Tuuling, Reedik; Yerokhin, Maksym; Dumas, Marlon; García-Bañuelos, Luciano; Laud, Peeter; Matulevičius, Raimundas; Pankova, Alisa; Pettai, Martin; Pullonen, Pille; and Tom, Jake, “Business Process Privacy Analysis in Pleak,” in Hähnle, Reiner and van der Aalst, Wil M. P., editors, *Fundamental Approaches to Software Engineering - 22nd International Conference, FASE 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings*, Springer, volume 11424 of *Lecture Notes in Computer Science*, pp. 306–312, URL https://doi.org/10.1007/978-3-030-16722-6_18.
- [69] Lepinski, Matthew; Levin, David; McCarthy, Daniel; Watro, Ronald; Lack, Michael; Hallenbeck, Daniel; and Slater, David, “Privacy-enhanced android for smart cities applications,” in *Smart City 360*, Springer, 2016, pp. 66–77.
- [70] Jeannet, Bertrand and Miné, Antoine, “Apron: A Library of Numerical Abstract Domains for Static Analysis,” in Bouajjani, Ahmed and Maler, Oded, editors, *Computer Aided Verification, 21st International Conference, CAV 2009, Grenoble, France, June 26 - July 2, 2009. Proceedings*, Springer, volume 5643 of *Lecture Notes in Computer Science*, pp. 661–667, URL https://doi.org/10.1007/978-3-642-02658-4_52.
- [71] “TPC BENCHMARK™ H, revision 2.17.3,” Transaction Processing Performance Council, 2017, http://www.tpc.org/TPC_Documents_Current_Versions/pdf/tpc-h_v2.17.3.pdf.
- [72] Chen, Yan; Machanavajjhala, Ashwin; Hay, Michael; and Miklau, Gerome, “PeGaSus: Data-Adaptive Differentially Private Stream Processing,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ACM, New York, NY, USA, CCS ’17, pp. 1375–1388, URL <http://doi.acm.org/10.1145/3133956.3134102>.
- [73] Virtanen, Pauli; Gommers, Ralf; Oliphant, Travis E.; Haberland, Matt; Reddy, Tyler; Cournapeau, David; Burovski, Evgeni; Peterson, Pearu; Weckesser, Warren; Bright, Jonathan; van der Walt, Stéfan J.; Brett, Matthew; Wilson, Joshua; Jarrod Millman, K.; Mayorov, Nikolay; Nelson, Andrew R. J.; Jones, Eric; Kern, Robert; Larson, Eric; Carey, CJ; Polat, İlhan; Feng, Yu; Moore, Eric W.; Vand erPlas, Jake; Laxalde, Denis; Perktold, Josef; Cimrman, Robert; Henriksen, Ian; Quintero, E. A.; Harris, Charles R; Archibald, Anne M.; Ribeiro, Antônio H.; Pedregosa, Fabian; van Mulbregt, Paul; and Contributors, SciPy 1. 0, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*.
- [74] Hundepool, Anco; Domingo-Ferrer, Josep; Franconi, Luisa; Giessing, Sarah; Nordholt, Eric Schulte; Spicer, Keith; and de Wolf, Peter-Paul, *Statistical Disclosure Control*, Wiley, 2012.

A LIST OF PUBLICATIONS

The following lists the publications (research papers, technical reports and blog posts) that have appeared so far as part of the NAPLES project.

- Marlon Dumas, Luciano García-Bañuelos, Peeter Laud: *Differential Privacy Analysis of Data Processing Workflows*. GraMSec@CSF 2016: 62-79
- Martin Pettai, Peeter Laud: *Combining Differential Privacy and Mutual Information for Analyzing Leakages in Workflows*. POST 2017: 298-319
- Pille Pullonen, Raimundas Matulevičius, Dan Bogdanov: *PE-BPMN: Privacy-Enhanced Business Process Model and Notation*. BPM 2017: 40-56 https://link.springer.com/chapter/10.1007/978-3-319-65000-5_3
- Marlon Dumas, Luciano García-Bañuelos, Peeter Laud: *Disclosure Analysis of SQL Workflows*. GraMSec@CSF 2018: 51-70
- Pille Pullonen, Jake Tom, Raimundas Matulevičius, Aivo Toots: *Privacy-enhanced BPMN: enabling data privacy analysis in business processes models*. Software & Systems Modeling, 2019 <https://link.springer.com/article/10.1007/s10270-019-00718-z>
- Aivo Toots, Reedik Tuuling, Maksym Yerokhin, Marlon Dumas, Luciano Garcia-Banuelos, Peeter Laud, Raimundas Matulevicius, Alisa Pankova, Martin Pettai, Pille Pullonen, Jake Tom: *Business Process Privacy Analysis in Pleak*. FASE, 2019. Full version available in <https://arxiv.org/abs/1902.05052>
- Aivo Toots, Reedik Tuuling, Maksym Yerokhin, Marlon Dumas, Luciano García-Bañuelos, Peeter Laud, Raimundas Matulevicius, Alisa Pankova, Martin Pettai, Pille Pullonen, Jake Tom: *Business Process Privacy Analysis in Pleak - (Extended Abstract)*. Inform. Spektrum 42(5): 354-355 (2019)
- Pille Pullonen: *How privacy enhancing technologies affect business processes*. Sharemind blog 2019 <https://sharemind.cyber.ee/studying-processes-with-privacy-technologies/>
- Peeter Laud, Alisa Pankova, Martin Pettai: *Achieving Differential Privacy using Methods from Calculus*, <https://arxiv.org/abs/1811.06343>
- Peeter Laud, Alisa Pankova: *Interpreting Epsilon of Differential Privacy in Terms of Advantage in Guessing or Approximating Sensitive Attributes*, <https://arxiv.org/abs/1911.12777>
- Sara Belluccini, Rocco De Nicola, Marlon Dumas, Pille Pullonen, Barbara Re, and Francesco Tiezzi: *Verification of privacy-enhanced collaborations*. FormaliSE 2020
- Pille Pullonen: *Analysing Privacy of Business Processes in Pleak*. Cybernetica's blog 2020 <https://cyber.ee/blog/2020/04-20/>

B LIST OF SYMBOLS AND ACRONYMS

$Ber(p)$	Bernoulli distribution with parameter p
DS_f	derivative sensitivity of a function f
GS_f	global sensitivity of a function f
$GenCauchy(\gamma)$	generalized Cauchy distribution with parameter γ
LS_f	local sensitivity of a function f
$Lap(\lambda)$	Laplace distribution with scaling λ
$N(\mu, \sigma^2)$	Normal distribution with mean μ and standard deviation σ
O	Big O; Big Oh; Big Omicron
$[m]$	the set of integers $\{1, \dots, m\}$
$[x]_\rho$	equivalence class w.r.t. relation ρ
Δ	difference
Ω	Big Omega
Θ	Big Theta
\boxtimes	the set of probability distributions over pairs, whose marginal distributions are the arguments of the operation
δ_{xy}	Kronecker symbol, equal to 1 iff $x = y$, and 0 otherwise
\downarrow	the downwards closure of a subset of a partially ordered set
η	monadic unit (in category theory)
\exists	existential quantifier
\forall	universal quantifier
$\frac{\partial f}{\partial x}$	partial derivative of a function f w.r.t. variable x
\int_X	(Lebesgue) integration over the set X
\mathbb{B}	set $\{0, 1\}$
\mathbb{N}	set of natural numbers
\mathbb{R}	set of real numbers
\mathbb{R}_+	set of positive real numbers
\mathbb{Z}	set of integers
\mathcal{D}	the set of all probability distributions over the given set
\mathcal{F}	the set of all upwards closed subsets of the argument
\mathcal{I}	the set of all downwards closed subsets of the argument
\mathcal{M}_f	a privacy enhancing mechanism applied to f
Carr	carrier set of the given generalized metric space
Pr	probability
c	the map between sets of distances for a mapping between generalized metric spaces
f	the map between carrier sets for a mapping between generalized metric spaces
d_{dp}	the <i>differential privacy distance</i> between probability distributions
supp	the support of a probability distribution
μ	monadic multiplication (in category theory)

∇f	gradient of a function f
\neg	negation (not)
\leq	if N and M are norms, $N \leq M$ denotes that $\ \vec{x}\ _N \leq \ \vec{x}\ _M$ for all \vec{x}
\prod	product
\sum	sum
\uparrow	the upwards closure of a subset of a partially ordered set
\vec{x}	vector
\vee	disjunction (or)
\wedge	conjunction (and)
d_X	distance in the set X
df_x	Fréchet derivative of f at the point x
f_X	probability density function of a random variable X
$\ \vec{x}\ _p$	ℓ_p -norm of a vector \vec{x}

AddSS	additive secret sharing
AES	Advanced Encryption Standard
AFRL	Air Force Research Laboratory
BDD	binary decision diagram
BP	Business Process
BPMN	Business Process Model and Notation
CDF	cumulative density function
CPU	Central Processing Unit
CRT	collaborative research team
D	direct dependency (in disclosure analysis)
DAG	directed acyclic graph
DARPA	Defense Advanced Research Projects Agency
DG	dependency graph
DP	differential privacy
DP-Workflow	data processing workflow
FHE	fully homomorphic encryption
FunSS	function secret sharing
GA	guessing advantage
GC	garbled circuit
GM	the category of generalized metric spaces
GPS	Global Positioning System
H	entropy (in quantitative analysis)
H	hidden (in disclosure analysis)
I	mutual information (in quantitative analysis)
I	indirect dependency (in disclosure analysis)
ID	identity

IoT	Internet of Things
IPSec	Internet Protocol Security
IT	Information Technology
KDE	kernel density estimation
MAC	Medium Access Control
MPC	Secure multiparty computation
NAPLES	Novel Tools for Analyzing Privacy Leakages
O	owner (in disclosure analysis)
OT	oblivious transfer
PA	project agreement
PAL	privacy abstraction layer
PDF	probability density function
PE	privacy enhanced
PE-BPMN	Privacy Enhanced Business Process Model and Notation
PET	privacy enhancing technology
PINQ	Privacy INtegrated Queries – a platform for privacy-preserving data analysis
PK	public key
PRESNA	Privacy Enhanced Social Network Analysis
PSDG	partial summary dependency graph
PULSAR	Private Updateable Lightweight Scalable Active Repository
RAM	Random Access Memory
REST	representational state transfer
SDG	summary dependency graph
SF	scale factor
SGX	Software Guard Extension
SIMD	single instruction multiple data
SIRD	disease states <i>Suspectible, Infected, Recovered, Deceased</i> in the Pandemic scenario
SK	secret key
SMT	satisfiability <i>modulo</i> theories
SoD	the set of distances of the given generalized metric space
SQL	Structured Query Language
SS	secret sharing
TA	technical area
TIPPERS	Testbed for IoT-based Privacy-Preserving PERvasive Spaces
TTL	time-to-live
UCI	University of California, Irvine
V	visible (in disclosure analysis)
VPN	Virtual Private Network