AFRL-RY-WP-TR-2015-0099

A RIGOROUS STATISTICAL FRAMEWORK FOR THE MATHEMATICS OF SENSING, EXPLOITATION, AND EXECUTION

Lawrence Carin

Duke University
Electrical & Computer Engineering Department

MAY 2015
Final Report

Approved for public release; distribution unlimited.

See additional restrictions described on inside pages

STINFO COPY

AIR FORCE RESEARCH LABORATORY
SENSORS DIRECTORATE
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7320
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE

NOTICE AND SIGNATURE PAGE

//Signature//

Michael Bakich, Program Manager
Assessment and Integration Branch
Layered Sensing Exploitation Division

//Signature//

Christina Schutte, Branch Chief
Assessment and Integration Branch
Layered Sensing Exploitation Division

//Signature//

Doug Hager, Deputy
Layered Sensing Exploitation Division
Sensors Directorate

REPORT DOCUMENTATION PAGE	*Form Approved OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
May 2015	Final	24 September 2011 – 28 February 2015

4. TITLE AND SUBTITLE
A RIGOROUS STATISTICAL FRAMEWORK FOR THE MATHEMATICS OF SENSING, EXPLOITATION, AND EXECUTION

5a. CONTRACT NUMBER FA8650-11-1-7150
5b. GRANT NUMBER
5c. PROGRAM ELEMENT NUMBER 61101E

6. AUTHOR(S)
Lawrence Carin

5d. PROJECT NUMBER 1000
5e. TASK NUMBER 11
5f. WORK UNIT NUMBER Y02D

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Duke University
Electrical & Computer Engineering Department
Durham, NC 27708-0291

BAE Systems

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory, Sensors Directorate
Wright-Patterson Air Force Base, OH 45433-7320
Air Force Materiel Command
United States Air Force

Defense Advanced
Research Projects Agency
675 North Randolph Street
Arlington, VA 22203-2114

10. SPONSORING/MONITORING AGENCY ACRONYM(S)
AFRL/RYAA

11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)
AFRL-RY-WP-TR-2015-0099

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. Report contains color.

14. ABSTRACT
A research program has been performed to develop next-generation mathematics for sensing, exploitation and execution (MSEE). An important focus of the research has been on a new class of nonparametric Bayesian architectures that constitute a rich modeling framework while still yielding parsimonious representations. Such models are attractive from multiple perspectives: (i) they flexibly adjust model complexity and sophistication to match the observed data, while (ii) explicitly defining model uncertainty manifested by missing data, and thereby (iii) linking utility of data to the objectives and associated models; additionally, (iv) these models are ideal for joint modeling of heterogeneous and possibly contradictory data, by sharing an inferred and typically low-dimensional latent space. In the MSEE construct, the utility of data is linked to the sensing objective, which in turn motivates and refines the associated models.

15. SUBJECT TERMS
sensing, mathematics, statistics

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT:**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** (Monitor)
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Michael Bakich
Unclassified	Unclassified	Unclassified	SAR	118	**19b. TELEPHONE NUMBER** *(Include Area Code)* N/A

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39-18

CONTENTS

I. Executive Summary

A research program has been executed to develop next-generation mathematics for sensing, exploitation and execution (MSEE). The methods considered account for stochasticity at their core, thereby being perfectly matched to defining the utility of data as a function of objective. An important focus of the research has been on a new class of *nonparametric* Bayesian architectures that constitute a rich modeling framework while still yielding parsimonious representations. Such models are attractive from multiple perspectives: (i) they flexibly adjust model complexity and sophistication to match the observed data, while (ii) explicitly defining model uncertainty manifested by missing data, and thereby (iii) linking utility of data to the objectives and associated models; additionally, (iv) these models are ideal for joint modeling of heterogeneous and possibly contradictory data, by sharing an inferred and typically low-dimensional latent space.

In the MSEE construct, the utility of data is linked to the sensing objective, which in turn motivates and refines the associated models. To assess the utility of data one must specify the objective, and from such the associated model(s). We will consider **unsupervised, semi-supervised and supervised models**, for such objectives as detection, classification, tracking and anomaly detection. The balance of **exploration and exploitation** is explicitly matched to the objective, available data, and previous experience ("life-long" learning) such that the models are *not* constituted anew for each sensing mission and objective (manifesting appropriate *transfer learning* from previous experiences). Since the models considered are explicitly statistical in nature, they are well suited to adaptivity, defining the utility of new data to the sensing objective (via design of experiments, and new non-myopic extensions, that exploit submodular characteristics of the mutual-information operator).

The focus in Phase II has been in extending the Phase I research into deep learning (Duke), and to perform a detailed test evaluation (BAE Systems). This report summarizes both of these areas of focus.

II. Bayesian Deep Learning

A. Introduction

Considerable research effort has been devoted to developing probabilistic models for documents. In the context of topic modeling, a popular approach is latent Dirichlet allocation (LDA) [1], a directed graphical model that aims to discover latent topics (word distributions) in collections of documents that are represented in bag-of-words form. Recent work focuses on linking observed word counts in a document to latent nonnegative matrix factorization, via a Poisson distribution, termed Poisson factor analysis (PFA) [2]. Different choices of priors on the latent nonnegative matrix factorization can lead to equivalent marginal distributions to LDA, as well as to the Focused Topic Model (FTM) of [3].

Additionally, hierarchical ("deep") tree-structured topic models have been developed by using structured Bayesian nonparametric priors, including the nested Chinese restaurant process (nCRP) [4], and the recently proposed nested hierarchical Dirichlet process (nHDP) [5]. The nCRP is limited because it requires that each document select topics from a single path in a tree, while the nHDP allows each document to access the entire tree by defining priors over a *base tree*. However, the relationship between two paths in these models is only explicitly given on shared parent nodes.

Another alternative for topic modeling is to develop undirected graphical models, such as the Replicated Softmax Model (RSM) [6], based on a generalization of the restricted Boltzmann machine (RBM) [7]. Also closely related to the RBM is the neural autoregressive density estimator (DocNADE) [8], a neural-network-based method, that has been shown to outperform the RSM.

Deep models, such as the Deep Belief Network (DBN) [9], the Deep Boltzmann Machine (DBM) [10], and layered Bayesian networks [11], [12], [13], [14] are becoming popular, as they consistently obtain state-of-the-art performances on a variety of machine learning tasks. A popular theme in this direction of work is to extend shallow topic models to deep counterparts. In such setting, documents arise from a cascade of layers of latent variables. For instance, DBNs and DBMs have been generalized to model documents by utilizing the RBM as building block [15], [16].

Combining ideas from traditional Bayesian topic modeling and deep models, we propose a new deep generative model for topic modeling, in which the Bayesian PFA is employed to interact with the data at the bottom layer, while the Sigmoid Belief Network (SBN) [17], a directed graphical model closely related to the RBM, is utilized to buildup binary hierarchies. Furthermore, our model is not necessarily restricted to SBN modules, and it is shown how an undirected model such as the RBM can be incorporated into the framework as well.

Compared with the original DBN and DBM, our proposed model: (*i*) tends to infer a more compact representation of the data, due to the "explaining away" effect described by [9]; (*ii*) allows for more direct exploration of the effect of a single deep hidden node through ancestral sampling; and (*iii*) can be easily incorporated into larger probabilistic models in a modular fashion. Compared with the nCRP and nHDP, our proposed model only infers topics at the bottom layer, but defines a flexible prior to capture high-order relationships between topics via a deep binary hierarchical structure.

Another important contribution we present is to develop two scalable Bayesian learning algorithms for our model: one of them based on the recently proposed *Bayesian conditional density filtering* (BCDF) algorithm [18], and the other based on the *stochastic gradient Nóse-Hoover thermostats* (SGNHT) algorithm [19]. We extend the SGNHT by introducing additional *thermostat variables* into the system, increasing the stability and convergence when compared to the original SGNHT algorithm.

B. Model Formulation

Our framework contains two parts, a Poisson factor analysis model and a deep structure based on the SBN (or RBM), detailed in the following.

C. Poisson Factor Analysis

Given a discrete matrix $\mathbf{X} \in \mathbb{Z}_+^{P \times N}$ containing counts from N documents and P words, Poisson factor analysis [2] assumes the entries of $\mathbf{X}$ are summations of $K < \infty$ latent counts, each produced by a latent factor (in the case of topic modeling, a hidden topic). We represent $\mathbf{X}$ using the following factor model

$$\mathbf{X} = \text{Pois}(\mathbf{\Phi}(\mathbf{\Theta} \circ \mathbf{H}^{(1)}))\,, \tag{1}$$

where $\mathbf{\Phi} \in \mathbb{R}_+^{P \times K}$ is the factor loading matrix. Each column of $\mathbf{\Phi}$, $\boldsymbol{\phi}_k$, encodes the relative importance of each word in topic k. $\mathbf{\Theta} \in \mathbb{R}_+^{K \times N}$ is the factor score matrix. Each column, $\boldsymbol{\theta}_n$, contains relative topic intensities specific to document n. $\mathbf{H}^{(1)} \in \{0,1\}^{K \times N}$ is a latent binary feature matrix. Each column, $\boldsymbol{h}_n^{(1)}$, defines a sparse set of topics associated with each document. For the single-layer PFA, the use of the superscript (1) on $\boldsymbol{h}_n^{(1)}$ is unnecessary; we introduce this notation here in preparation for the subsequent deep model, for which $\boldsymbol{h}_n^{(1)}$ will correspond to the associated first-layer latent binary units. The symbol $\circ$ represents the Hadamard, or element-wise multiplication of two matrices. The factor scores for document n are $\boldsymbol{\theta}_n \circ \boldsymbol{h}_n^{(1)}$.

A wide variety of algorithms have been developed by constructing PFAs with different prior specifications [20]. If $\mathbf{H}^{(1)}$ is an all-ones matrix, LDA is recovered from (1) by employing Dirichlet priors on $\boldsymbol{\phi}_k$ and $\boldsymbol{\theta}_n$, for $k = 1, \ldots, K$ and $n = 1, \ldots, N$, respectively. This version of LDA is referred to as Dir-PFA by [2]. For our proposed model, we construct PFAs by placing Dirichlet priors on $\boldsymbol{\phi}_k$ and gamma priors on $\boldsymbol{\theta}_n$. This is summarized as,

$$x_{pn} = \textstyle\sum_{k=1}^{K} x_{pnk}\,, \quad x_{pnk} \sim \text{Pois}(\phi_{pk}\theta_{kn}h_{kn}^{(1)})\,, \tag{2}$$

with priors specified as $\boldsymbol{\phi}_k \sim \text{Dir}(a_\phi, \ldots, a_\phi)$, $\theta_{kn} \sim \text{Gamma}(r_k, p_n/(1-p_n))$, $r_k \sim \text{Gamma}(\gamma_0, 1/c_0)$, and $\gamma_0 \sim \text{Gamma}(e_0, 1/f_0)$.

The novelty in our model comes from the prior for the binary feature matrix $\mathbf{H}^{(1)}$. Previously, [20] proposed a Beta-Bernoulli process prior on the columns $\{\boldsymbol{h}_n^{(1)}\}_{n=1}^N$ with $p_n = 0.5$. This model was called NB-FTM, tightly related with the focused topic model (FTM) [3]. In the work presented here, we construct $\mathbf{H}^{(1)}$ from a deep structure based on the SBN (or RBM) with binary latent units.

D. Structured Priors on the Latent Binary Matrix

The second part of our model consists of a deep structure for a binary hierarchy. To this end, we employ the SBN (or RBM). In the following we start by describing a single-layer model with SBN (or RBM), and then we generalize it to a deep model.

a) Modeling with the SBN: We assume the latent vector for document n, $\boldsymbol{h}_n^{(1)} \in \{0,1\}^{K_1}$, is binary. This matches most of the RBM and SBN literature, for which typically the *observed* data are binary. In our model, however, these binary variables are not observed; they are hidden and related to the data through the PFA in (2).

To construct a structured prior, we define another hidden set of units $\boldsymbol{h}_n^{(2)} \in \{0,1\}^{K_2}$ placed at a layer "above" $\boldsymbol{h}_n^{(1)}$. They layers are related through a set of weights defined by the matrix $\mathbf{W}^{(1)} = [\boldsymbol{w}_1^{(1)} \ \ldots \ \boldsymbol{w}_{K_1}^{(1)}]^\top \in \mathbb{R}^{K_1 \times K_2}$. An SBN model has the generative process,

$$p(h_{k_2 n}^{(2)} = 1) = \sigma(c_{k_2}^{(2)})\,, \tag{3}$$

$$p(h_{k_1 n}^{(1)} = 1 | \boldsymbol{h}_n^{(2)}) = \sigma\left((\boldsymbol{w}_{k_1}^{(1)})^\top \boldsymbol{h}_n^{(1)} + c_{k_1}^{(1)}\right)\,, \tag{4}$$

where $h_{k_1 n}^{(1)}$ and $h_{k_2 n}^{(2)}$ are elements of $\boldsymbol{h}_n^{(1)}$ and $\boldsymbol{h}_n^{(2)}$, respectively. The function $\sigma(x) \triangleq 1/(1+e^{-x})$ is the logistic function, and $c_{k_1}^1$ and $c_{k_2}^2$ are bias terms. The global parameters $\mathbf{W}^{(1)}$ are used to characterize the mapping from $\boldsymbol{h}_n^{(2)}$ to $\boldsymbol{h}_n^{(1)}$ for all documents.

b) Modeling with the RBM: The SBN is closely related to the RBM, which is a Markov random field with the same bipartite structure as the SBN. The RBM defines a distribution over a binary vector that is proportional to the exponential of its *energy*, which is defined (using the same notation as in SBN) as $E(\boldsymbol{h}_n^{(1)}, \boldsymbol{h}_n^{(2)}) =$

$$-(\boldsymbol{h}_n^{(1)})^\top \boldsymbol{c}^{(1)} - \ (\boldsymbol{h}_n^{(1)})^\top \mathbf{W}^{(1)} \boldsymbol{h}_n^{(2)} - (\boldsymbol{h}_n^{(2)})^\top \boldsymbol{c}^{(2)}\,. \tag{5}$$

In the experiments we consider both the deep SBN and deep RBM for representation of the latent binary units, which are connected to topic usage in a given document.

c) Discussion: An important benefit of SBNs over RBMs is that in the former sparsity or shrinkage priors on $\mathbf{W}^{(1)}$ can be readily imposed on the global parameters of the model, and fully Bayesian inference can be implemented as shown in [14]. The RBM relies on an approximation technique known as contrastive divergence [7], for which prior specification for the model parameters is limited.

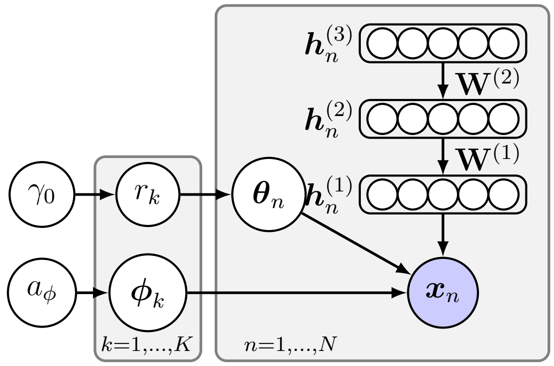

Fig. 1. Graphical model for the Deep Poisson Factor Analysis with three layers of hidden binary hierarchies. The directed binary hierarchy may be replaced by a *deep Boltzmann machine*.

E. Deep Architecture for Topic Modeling

Specifying a prior distribution on $\boldsymbol{h}_n^{(2)}$ as in (3) might be too restrictive in some cases. Alternatively, we can use another SBN prior for $\boldsymbol{h}_n^{(2)}$, in fact, we can add multiple layers as in [14] to obtain a deep architecture,

$$p(\boldsymbol{h}_n^{(1)}, \ldots, \boldsymbol{h}_n^{(L)}) = p(\boldsymbol{h}_n^{(L)}) \textstyle\prod_{\ell=2}^{L} p(\boldsymbol{h}_n^{(\ell-1)} | \boldsymbol{h}_n^{(\ell)}), \tag{6}$$

where L is the number of layers, $p(\boldsymbol{h}_n^{(L)})$ is the prior for the top layer defined as in (3), $p(\boldsymbol{h}_n^{(\ell-1)}|\boldsymbol{h}_n^{(\ell)})$ is defined in (4), and the weights $\mathbf{W}^{(\ell)} \in \mathbb{R}^{K_\ell \times K_{\ell+1}}$ and biases $\boldsymbol{c}^{(\ell)} \in \mathbb{R}^{K_\ell}$ are omitted from the conditional distributions to keep notation uncluttered. A similar deep architecture may be designed for the RBM [10].

Instead of employing the beta-Bernoulli specification for $\boldsymbol{h}_n^{(1)}$ as in the NB-FTM, which assumes independent topic usage probabilities, we propose using (6) instead as the prior for $\boldsymbol{h}_n^{(1)}$, thus

$$p(\boldsymbol{x}_n, \boldsymbol{h}_n) = p(\boldsymbol{x}_n | \boldsymbol{h}_n^{(1)}) p(\boldsymbol{h}_n^{(1)}, \ldots, \boldsymbol{h}_n^{(L)}) \,, \tag{7}$$

where $\boldsymbol{h}_n \triangleq \{\boldsymbol{h}_n^{(1)}, \ldots, \boldsymbol{h}_n^{(L)}\}$, and $p(\boldsymbol{x}_n|\boldsymbol{h}_n^{(1)})$ as in (2). The prior $p(\boldsymbol{h}_n^{(1)}|\boldsymbol{h}_n^{(2)} \ldots, \boldsymbol{h}_n^{(L)})$ can be seen as a flexible prior distribution over binary vectors that encodes high-order interactions across elements of $\boldsymbol{h}_n^{(1)}$. The graphical model for our model, Deep Poisson Factor Analysis (DPFA) is shown in Figure 1.

F. Scalable Posterior Inference

We focus on learning our model with fully Bayesian algorithms, however, emerging large-scale corpora prohibit standard MCMC inference algorithms to be applied directly. For example, in the experiments, we consider the *RCV1-v2* and the *Wikipedia* corpora, which contain about 800K and

10M documents, respectively. Therefore, fast algorithms for big Bayesian learning are essential. While parallel algorithms based on distributed architectures such as the *parameter server* [21], [22] are popular choices, in the work presented here, we focus on another direction for scaling up inference by stochastic algorithms, where mini-batches instead of the whole dataset are utilized in each iteration of the algorithms. Specifically, we develop two stochastic Bayesian inference algorithms based on Bayesian conditional density filtering [18] and stochastic gradient thermostats [19], both of which have theoretical guarantees in the sense of asymptotical convergence to the true posterior distribution.

G. Bayesian conditional density filtering

Bayesian conditional density filtering (BCDF) is a recently proposed stochastic algorithm for Bayesian online learning [18], that extends Markov chain Monte Carlo (MCMC) sampling to streaming data. Sampling in BCDF proceeds by drawing from the conditional posterior distributions of model parameters, obtained by propagating surrogate conditional sufficient statistics (SCSS). In practice, we repeatedly update the SCSS using the current mini-batch and draw S samples from the conditional densities using, for example, a Gibbs sampler. This eliminates the need to load the entire dataset into memory, and provides computationally cheaper Gibbs updates. More importantly, it can be proved that BCDF leads to an approximation of the conditional distributions that produce samples from the correct target posterior asymptotically, once the entire dataset is seen [18].

Algorithm 1 BCDF algorithm for DPFA.

Input: text documents, *i.e.*, a count matrix $\mathbf{X}$.
Initialize $\mathbf{\Psi}_g^{(0)}$ randomly and set $\mathbf{S}_g^{(0)}$ all to zero.
for $t = 1$ **to** ∞ **do**
  Get one mini-batch $\mathbf{X}^{(t)}$.
  Initialize $\mathbf{\Psi}_g^{(t)} = \mathbf{\Psi}_g^{(t-1)}$, and $\mathbf{S}_g^{(t)} = \mathbf{S}_g^{(t-1)}$.
  Initialize $\mathbf{\Psi}_l^{(t)}$ randomly.
  for $s = 1$ **to** S **do**
    Gibbs sampling for DPFA on $\mathbf{X}^{(t)}$.
    Collect samples $\mathbf{\Psi}_g^{1:S}, \mathbf{\Psi}_l^{1:S}$ and $\mathbf{S}_g^{1:S}$.
  end for
  Set $\mathbf{\Psi}_g^{(t)} = \mathrm{mean}(\mathbf{\Psi}_g^{1:S})$, and $\mathbf{S}_g^{(t)} = \mathrm{mean}(\mathbf{S}_g^{1:S})$.
end for

In the learning phase, we are interested in learning the global parameters $\mathbf{\Psi}_g = (\{\boldsymbol{\phi}_k\}, \{r_k\}, \gamma_0, \{\mathbf{W}^{(\ell)}, \boldsymbol{c}^{(\ell)}\})$. Denote local variables as $\mathbf{\Psi}_l = (\mathbf{\Theta}, \mathbf{H}^{(\ell)})$, and let $\mathbf{S}_g$ represent the SCSS for $\mathbf{\Psi}_g$, the BCDF algorithm can be summarized in Algorithm 1. Specifically, we need to obtain the conditional densities, which can be readily derived granted the full local conjugacy of the proposed model. Using *dot notation* to represent marginal sums, *e.g.*, $x_{\cdot nk} \triangleq \sum_p x_{pnk}$, we can write the key conditional densities for (2) as [20]

$$
\begin{aligned}
x_{pnk}|- &\sim \text{Multi}(x_{pn}; \zeta_{pn1}, \ldots, \zeta_{pnK}) \,, \\
\boldsymbol{\phi}_k|- &\sim \text{Dir}(a_\phi + x_{1\cdot k}, \ldots, a_\phi + x_{P\cdot k}) \,, \\
\theta_{kn}|- &\sim \text{Gamma}(r_k h_{kn}^{(1)} + x_{\cdot nk}, p_n) \,, \\
h_{kn}^{(1)}|- &\sim \delta(x_{\cdot nk} = 0)\text{Ber}\left(\tfrac{\tilde{\pi}_{kn}}{\tilde{\pi}_{kn} + (1-\pi_{kn})}\right) + \delta(x_{\cdot nk} > 0) \,,
\end{aligned}
$$

where $\tilde{\pi}_{kn} = \pi_{kn}(1-p_n)^{r_k}$ and $\pi_{kn} = \sigma((\boldsymbol{w}_k^{(1)})^\top \boldsymbol{h}_n^{(2)} + c_k^{(1)})$. Additional details are provided in the Supplementary Material. For the conditional distributions of $\mathbf{W}^{(\ell)}$ and $\mathbf{H}^{(\ell)}$, we use the same data augmentation technique as in [14], where Pólya-Gamma (PG) random variables $\gamma_{k_\ell n}^{(\ell)}$ [23] are introduced for hidden unit k_ℓ in layer ℓ corresponding to observation $\boldsymbol{v}_n$. Specifically, each $\gamma_{k_\ell n}^{(\ell)}$ has conditional posterior $\text{PG}(1, (\boldsymbol{w}_{k_\ell}^{(\ell)})^\top \boldsymbol{h}_n^{(\ell+1)} + c_{k_\ell}^{(\ell)})$. If we place a Gaussian prior $N(0, \sigma^2\mathbf{I})$ on $\boldsymbol{w}_{k_\ell}^{(\ell)}$, the posterior will still be Gaussian with covariance matrix $\boldsymbol{\Sigma}_{k_\ell}^{(\ell)} = [\sum_n \gamma_{k_\ell n}^{(\ell)} \boldsymbol{h}_n^{(\ell+1)} (\boldsymbol{h}_n^{(\ell+1)})^\top + \sigma^{-2}\mathbf{I}]^{-1}$ and mean $\boldsymbol{\mu}_{k_\ell}^{(\ell)} = \boldsymbol{\Sigma}_{k_\ell}^{(\ell)} [\sum_n (h_{k_\ell n}^{(\ell)} - 1/2 - c_{k_\ell}^{(\ell)} \gamma_{k_\ell n}^{(\ell)}) \boldsymbol{h}_n^{(\ell+1)}]$. Furthermore, for any $\ell > 1$, the conditional posterior distribution of $h_{k_\ell n}^{(\ell)}$ can be obtained as[1]

$$
h_{k_\ell n}^{(\ell)} \sim \text{Bernoulli}\left(\sigma(d_{k_\ell n})\right) \,, \tag{8}
$$

where

$$
\begin{aligned}
d_{k_\ell n} &= (\boldsymbol{w}_{\cdot, k_\ell}^{(\ell-1)})^\top \boldsymbol{h}_n^{(\ell-1)} + (\boldsymbol{w}_{k_\ell}^{(\ell)})^\top \boldsymbol{h}_n^{(\ell+1)} + c_{k_\ell}^{(\ell)} \\
&- \frac{1}{2} \sum_{k_{\ell-1}} \left(w_{k_{\ell-1} k_\ell}^{(\ell-1)} + \gamma_{k_{\ell-1} n}^{(\ell-1)} (2\psi_{k_{\ell-1} n}^{\backslash k_\ell} w_{k_{\ell-1} k_\ell}^{(\ell-1)} + (w_{k_{\ell-1} k_\ell}^{(\ell-1)})^2) \right) \,,
\end{aligned}
$$

and $\psi_{k_{\ell-1} n}^{\backslash k_\ell} = \sum_{k'_\ell \neq k_\ell} w_{k_{\ell-1} k'_\ell}^{(\ell-1)} h_{k'_\ell n}^{(\ell)} + c_{k_{\ell-1}}^{(\ell-1)}$. Note that $\boldsymbol{w}_{\cdot, k_{\ell+1}}^{(\ell)}$ and $\boldsymbol{w}_{k_\ell}^{(\ell)}$ represents the $k_{\ell+1}$th column and the transpose of the k_ℓth row of $\mathbf{W}^{(\ell)}$, respectively. As can be seen, the conditional posterior distribution of $h_{k_\ell n}^{(\ell)}$ is both related to $\boldsymbol{h}_n^{(\ell-1)}$ and $\boldsymbol{h}_n^{(\ell+1)}$.

H. Stochastic gradient thermostats

Our second algorithm adopts the recently proposed SGNHT for large scale Bayesian sampling [19], which is more scalable and accurate than the previous BCDF algorithm. SGNHT generalizes the *stochastic gradient Langevin dynamics* (SGLD) [24] and the *stochastic gradient Hamiltonian Monte Carlo* (SGHMC) [25] by introducing momentum variables into the system, which is adaptively damped using a thermostat. The thermostat exchanges energy with the target system (*e.g.*,

[1]Here and in the rest of the paper, whenever $\ell > L$, $\boldsymbol{h}_n^{(\ell)}$ is defined as a zero vector, for conciseness.

a Bayesian model) to maintain a constant temperature; this has the potential advantage of making the system jump out of local modes easier and reach the equilibrium state faster [19].

Specifically, let $\boldsymbol{\Psi}_g \in \mathbb{R}^M$ be model parameters[2] which corresponds to the location of particles in a physical system, $\boldsymbol{v} \in \mathbb{R}^M$ be the momentum of these particles, which are driven by stochastic forces $\tilde{f}$ defined as the negative stochastic gradient (evaluated on a subset of data) of a Bayesian posterior, *e.g.*, $\tilde{f}(\boldsymbol{\Psi}_g) \triangleq -\nabla_{\boldsymbol{\Psi}_g}\tilde{U}(\boldsymbol{\Psi}_g)$, where $\tilde{U}(\boldsymbol{\Psi}_g)$ is the negative log-posterior of a Bayesian model. The motion of the particles in the system are then defined by the following stochastic different equations:

$$\begin{aligned} \mathrm{d}\boldsymbol{\Psi}_g = \boldsymbol{v}\mathrm{d}t\,, \quad \mathrm{d}\boldsymbol{v} &= \tilde{f}(\boldsymbol{\Psi}_g)\mathrm{d}t - \xi\boldsymbol{v}\mathrm{d}t + \sqrt{D}\mathrm{d}\mathcal{W}\,, \\ \mathrm{d}\xi &= \left(\tfrac{1}{M}\boldsymbol{v}^T\boldsymbol{v} - 1\right)\mathrm{d}t\ , \end{aligned} \tag{9}$$

where t indexes time, $\mathcal{W}$ is the standard Wiener process, ξ is called the thermostat variable which ensures the system temperature to be constant, and D is the variance of the total noise injected into the system and is assumed to be constant.

It can be shown that under certain assumptions, the equilibrium distribution of system (9) corresponds to the model posterior [19]. As a result, the SDE (9) can be solved by using the Euler-Maruyama scheme [26], where a mini-batch of the whole data is used to evaluate the stochastic gradient $\tilde{f}$. Note only one thermostat variable ξ is used in the SDE system (9), this is not robust enough to control the system temperature well because of the high dimensionality of $\boldsymbol{\Psi}_g$. Based on the techniques in [19], we extend the SGNHT by introducing multiple thermostat variables $(\xi_1, \cdots, \xi_M)$ into the system such that each ξ_i controls one degree of the particle momentum. Intuitively, this allows energy to be exchanged between particles and thermostats more efficiently, thus driving the system to equilibrium states more rapidly. Empirically we have also verified the superiority of the proposed modification over the original SGNHT. Formally, let $\boldsymbol{\Xi} = \mathrm{diag}(\xi_1, \xi_2, \cdots, \xi_M)$, $\mathbf{q} = \mathrm{diag}(v_1^2, \cdots, v_M^2)$, we define our proposed SGNHT using the following SDEs

$$\begin{aligned} \mathrm{d}\boldsymbol{\Psi}_g = \boldsymbol{v}\mathrm{d}t\,, \quad \mathrm{d}\boldsymbol{v} &= \tilde{f}(\boldsymbol{\Psi}_g)\mathrm{d}t - \boldsymbol{\Xi}\boldsymbol{v}\mathrm{d}t + \sqrt{D}\mathrm{d}\mathcal{W}\,, \\ \mathrm{d}\boldsymbol{\Xi} &= (\mathbf{q} - \mathbf{I})\,\mathrm{d}t\,, \end{aligned} \tag{10}$$

where $\mathbf{I}$ is the identity matrix. Interestingly, we are still able to prove that the equilibrium distribution of the above system corresponds to the model posterior.

Theorem 1: The equilibrium distribution of the SDE system in (10) is $p(\boldsymbol{\Psi}_g, \boldsymbol{v}, \boldsymbol{\Xi})$

[2]With a little abuse of notation but for conciseness, we use $\boldsymbol{\Psi}_g$ to denote the reparameterized version of the parameters (such that $\boldsymbol{\Psi}_g \in \mathbb{R}^M$) if any, required in SGNHT.

$$\propto \exp\left(-\frac{1}{2}\boldsymbol{v}^T\boldsymbol{v} - U(\boldsymbol{\Psi}_g) - \frac{1}{2}\text{tr}\left\{(\Xi - D)^T(\Xi - D)\right\}\right).$$

The proof of the theorem is provided in the Supplementary Material. By Theorem 1, it is straightforward to see that the marginal distribution $p(\boldsymbol{\Psi}_g)$ of $p(\boldsymbol{\Psi}_g, \boldsymbol{v}, \Xi)$ is exactly the posterior of our Bayesian model. As a result, again we can generate approximate samples from $p(\boldsymbol{\Psi}_g, \boldsymbol{v}, \Xi)$ using the Euler-Maruyama scheme and discard the auxiliary variables $\boldsymbol{v}$ and Ξ.

d) Learning for the SBN based model: Our SBN based model is illustrated in Figure 1. In the learning phase we are interested in learning the global parameters $\boldsymbol{\Psi}_g$, the same as in BCDF. The constraints inside the parameters $\{\boldsymbol{\phi}_k\}$, *i.e.*, $\sum_p \phi_{pk} = 1$, prevent the SGNHT from being applied directly. Although we can overcome this problem by using some re-parameterization methods as those used in [27], we find it converges better when considering information geometry for these parameters. As a result, we use stochastic gradient Riemannian Langevin dynamics (SGRLD) [27] to sample the topic-word distributions $\{\boldsymbol{\phi}_k\}$, and use the SGNHT to sample the remaining parameters. Based on the data augmentation for x_{pn} above, Section II-G shows that the posteriors of $\{\boldsymbol{\phi}_k\}$'s are Dirichlet distributions. This enables us to apply the same scheme as the SGRLD for LDA [27] to sample $\{\boldsymbol{\phi}_k\}$'s. More details are provided in the Supplementary Material.

The rest of the parameters can be straightforwardly sampled using the SGNHT algorithm. Specifically we need to calculate the stochastic gradients of $\mathbf{W}^{(\ell)}$ and $\boldsymbol{c}^{(\ell)}$ evaluated on a mini-batch of data (denote $\mathcal{D}$ as the index set of a mini-batch). Based on the model definition in (6), these can be calculated as

$$\frac{\partial \tilde{U}}{\partial \boldsymbol{w}_{k_\ell}^{(\ell)}} = \frac{N}{|\mathcal{D}|}\sum_{n\in\mathcal{D}} \mathbb{E}_{\boldsymbol{h}_n^{(\ell)}, \boldsymbol{h}_n^{(\ell+1)}}\left[\left(\tilde{\sigma}_{k_\ell n}^{(\ell)} - h_{k_\ell n}^{(\ell)}\right)\boldsymbol{h}_n^{(\ell+1)}\right],$$

$$\frac{\partial \tilde{U}}{\partial c_{k_\ell}^{(\ell)}} = \frac{N}{|\mathcal{D}|}\sum_{n\in\mathcal{D}} \mathbb{E}_{\boldsymbol{h}_n^{(\ell)}, \boldsymbol{h}_n^{(\ell+1)}}\left[\tilde{\sigma}_{k_\ell n}^{(\ell)} - h_{k_\ell n}^{(\ell)}\right],$$

where $\tilde{\sigma}_{k_\ell n}^{(\ell)} = \sigma((\boldsymbol{w}_{k_\ell}^{(\ell)})^\top \boldsymbol{h}_n^{(\ell+1)} + c_{k_\ell}^{(\ell)})$, and the expectation is taken over posteriors. As in the case of LDA [27], no closed-form integrations can be obtained for the above gradients, we thus use Monte Carlo integration to approximate the quantity. Specially, given $\{\boldsymbol{w}_{k_\ell}^{(\ell)}, c_{k_\ell}^{(\ell)}\}$, we are able to collect samples of the local binary variables $(\boldsymbol{h}_n^{(\ell)})_{n\in\mathcal{D}}$ by running a few Gibbs steps and then using these samples to approximate the intractable integrations. A direct variable cancelation approach results in exact conditional distributions for $h_{kn}^{(\ell)}$, however, we found that this approach does not mix well due to the highly correlated structure of hidden variables. Instead, we sample $h_{k_\ell n}^{(\ell)}$ based on the same augmentation used in BCDF, given in (8).

e) Learning for the RBM based model: As mentioned above, our RBM based model is recovered when replacing the SBN with the RBM in Figure 1. Despite minor changes in the construction, the unnormalized distribution of the RBM prohibits exact MCMC sampling from being applied. As a result, we develop an approximate learning algorithm that alternates between sampling $(\{\boldsymbol{\phi}_k\}, \{\gamma_k\}, \gamma_0\})$ and $(\{\mathbf{W}^{(\ell)}, \boldsymbol{c}^{(\ell)}\})$. Specifically, we use the same conditional posteriors as the SBN based model to sample the former, but use the *contrastive divergence* algorithm (CD-1) [7] for the latter. One main difference of our CD-1 algorithm *w.r.t* the original one is that the inputs (*i.e.*, $\boldsymbol{h}_n^{(1)}$) are actually latent variables. To make the CD-1 work, conditioned on other model parameters, we first sample $\boldsymbol{h}_n^{(1)}$ using the posterior given in Section II-G, then conditioned on $\boldsymbol{h}_n^{(1)}$, we apply the original CD-1 algorithm to calculate the approximate gradients for $(\{\mathbf{W}^{(\ell)}, \boldsymbol{c}^{(\ell)}\})$, which are then used for a gradient descent step. In fact, the CD-1 also makes part of to the stochastic approximate algorithms in [28], making it naturally fit into our SGNHT framework.

I. Discussion

Both the BCDF and SGNHT are stochastic inference algorithms, allowing the models be applied to large-scale data. In terms of ease of implementation, BCDF beats SGNHT in most cases, especially when the model is conjugate and the domain of parameters is constrained (*e.g.*, variables on a simplex). However, in general terms BCDF is more restrictive than SGNHT. For example, BCDF requires the conditional densities for all the parameters, which is unavailable in some cases. Furthermore, BCDF has the limitation of being unable to deal with some *big models* where the number of model parameters is large, for instance, when the dimension of the hidden variables from the SBN in our model is huge. Finally, the conditions for the BCDF to converge to the true posterior are more restricted. Altogether, these reasons make the SGNHT more robust than the BCDF.

J. Related Work

In traditional Bayesian topic models, topic correlations are typically modeled with shallow structures, *e.g.*, the correlated topic model [29] with correlation between topic proportions imposed via the logistic normal distribution. There exist also some work on hierarchical ("deep") correlation modeling, *e.g.*, the hierarchical Dirichlet process [30], which models topic proportions hierarchically via a stack of DPs. The nested Chinese restaurant process [4] (nCRP) models topic hierarchies by defining a tree structure prior based on the Chinese restaurant process, and the nested hierarchical Dirichlet process [5] extends the nCRP by allowing each document to be able to access all the paths in the tree. One major difference between these models and ours is that they focus on discovering topic hierarchies instead of modeling general topic correlations.

In the deep learning community, topic models are mostly built using the RBM as building block. For example, [15] and [31] extended the DBN for topic modeling, while a deep version of the RSM was proposed by [16]. More recent work focuses on employing deep directed generative models for topic modeling, *e.g.*, deep exponential families [32], a class of latent variable models extending the DBN by defining the distribution of hidden variables in each layer using the exponential family, instead of the restricted Bernoulli distribution.

In terms of learning and inference algorithms, most of existing Bayesian topic models rely on MCMC methods or variational Bayes algorithms, which are impractical when dealing with large scale data. Therefore, stochastic variational inference algorithms have been developed [33], [34], [35], [36]. Although scalable and usually fast converging, one unfavorable shortcoming of stochastic variational inference algorithms is the mean-field assumption on the approximate posterior.

Another direction for scalable Bayesian learning relies on the theory from stochastic differential equations (SDE). Specifically, [24] proposed the first stochastic MCMC algorithm, called *stochastic gradient Langevin dynamics* (SGLD), for large scale Bayesian learning. In order to make the learning faster, [27] generalized SGLD by considering information geometry [37], [38] of model posteriors. Furthermore, [25] generalized the SGLD by a second-order Langevin dynamic, called *stochastic gradient Hamiltonian Monte Carlo* (SGHMC). This is the stochastic version of the well known Hamiltonian MCMC sampler. One problem with SGHMC is that the unknown stochastic noise needs to be estimated to make the sampler correct, which is impractical. *Stochastic gradient thermostats* algorithms (SGNHT) overcome this problem by introducing the thermostat into the algorithm, such that the unknown stochastic noise could be adaptively absorbed into the thermostat, making the sampler asymptotically exact. Given the advantages of the SGNHT, in this paper we extend it to a multiple thermostats setting, where each thermostat exchanges energy with a degree of freedom of the system. Empirically we show thatour extension improves on the original algorithm.

K. Experiments

1) Datasets and Setups: We present experimental results on three publicly available corpora: a relatively small, *20 Newsgroups*, a moderately large, Reuters Corpus Volume I (*RCV1-v2*), and a large one, *Wikipedia*. The first two corpora are the same as those used in [16]. Specifically, the *20 Newsgroups* corpus contains 18,845 documents with a total of 0.7M words and a vocabulary size of 2K. The data was partitioned chronologically into 11,314 training and 7,531 test documents. The *RCV1-v2* corpus contains 804,414 newswire articles. There are 103 topics that form a tree hierarchy. After preprocessing, we are left with about 75M words, with a vocabulary size of 10K. We randomly select 794,414 documents for training and 10,000 for testing. Finally, we downloaded

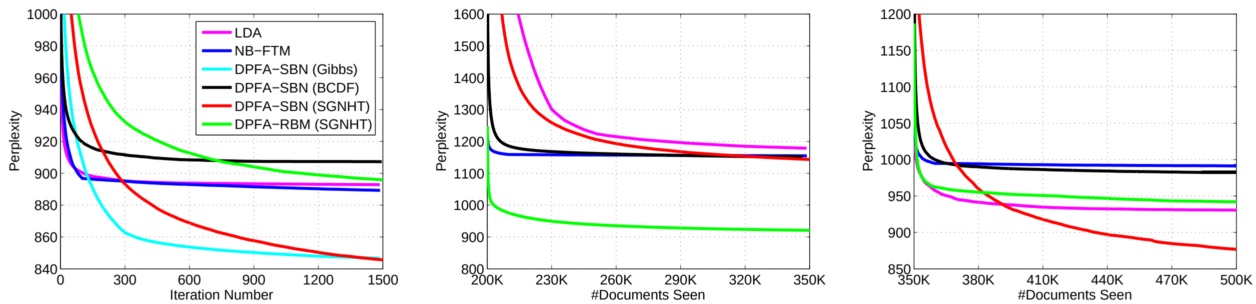

Fig. 2. Predictive perplexities on a held-out test set as a function of training documents seen. The number of hidden unites in each layer is 128,64,32, respectively. (Left) *20 Newsgroups*. (Middle) *RCV1-v2*. (Right) *Wikipedia*.

10M random documents from *Wikipedia* using scripts provided in [33] and randomly selected 1K documents for testing. As in [33], [27], a vocabulary size of 7,702 was taken from the top 10K words in Project Gutenberg texts.

The DPFA model consisting of SBN is denoted as DPFA-SBN, while its RBM counterpart is denoted DPFA-RBM. The performance of DPFA is compared to that of the following models: LDA [1], NB-FTM [20], nHDP [5] and RSM [6].

For all the models considered, we calculate the predictive perplexities on the test set as follows: holding the global model parameters fixed, for each test document we randomly partition the words into a 80/20% split. We learn document-specific parameters using the 80% portion, and then calculate the predictive perplexities on the remaining 20% subset. Evaluation details are provided in the Supplementary Material.

For *20 Newsgroups* and *RCV1-v2* corpora, we use 2,000 mini-batches for burn-in followed by 1,500 collection samples to calculate test perplexities; while for the *Wikipedia* dataset, 3,500 mini-batches are used for burn-in. The mini-batch size for all stochastic algorithms is set to 100. To choose good parameters for SGNHT, *e.g.*, the step size and the variance of the injected noise, we randomly choose about 10% documents from the training data as validation set. For BCDF, 100 MCMC iterations are evaluated for each mini-batch, with the first 60 samples discarded. We set the hyperparameters of DPFA as $a_\phi = 1.01, c_0 = e_0 = 1, f_0 = 0.01$. The RSM is trained using convergence-divergence with step size 5 and a maximum of 10,000 iterations. For nHDP, we use the publicly available code from [5], in which stochastic variational Bayes (sVB) inference is implemented.

2) Quantitative Evaluation:

a) 20 Newsgroups: The results for the *20 Newsgroups* corpus are shown in Table I. Perplexities are reported for our implementation of Gibbs sampling, BCDF and SGNHT, and the four considered

TABLE I

TEST PERPLEXITIES FOR *20 Newsgroups*. "DIM" REPRESENTS THE NUMBER OF HIDDEN UNITS IN EACH LAYER, STARTING FROM THE BOTTOM. DPFA-SBN-*t* REPRESENTS THE DPFA-SBN MODEL WITH STUDENT'S *t* PRIOR ON $\mathbf{W}^{(\ell)}$. ($\diamond$) REPRESENTS THE *base tree* SIZE IN NHDP.

MODEL	METHOD	DIM	PERP.
DPFA-SBN-*t*	GIBBS	128-64-32	**827**
DPFA-SBN	GIBBS	128-64-32	**846**
DPFA-SBN	SGNHT	128-64-32	**846**
DPFA-RBM	SGNHT	128-64-32	896
DPFA-SBN	BCDF	128-64-32	905
DPFA-SBN	GIBBS	128-64	851
DPFA-SBN	SGNHT	128-64	850
DPFA-RBM	SGNHT	128-64	893
DPFA-SBN	BCDF	128-64	896
LDA	GIBBS	128	893
NB-FTM	GIBBS	128	887
RSM	CD5	128	877
NHDP	SVB	$(10,10,5)^{\diamond}$	889

competing methods. First, we examine the performance of different inference algorithms. As can be seen, for the same size model, *e.g.*, 128-64-32, SGNHT can achieve essentially the same performance as Gibbs sampling, while BCDF is more likely to get trapped in a local mode. Next, we explore the advantage of employing deep models. Using three layers instead of two gives performance improvements in almost all the algorithms. In Gibbs sampling, there is an improvement of 36 units for the DPFA-SBN model, when a second layer is learned (NB-FTM being the one-hidden-layer DPFA). Adding the third hidden layer further improves the test perplexity.

Adding a sparsity-encouraging prior on $\mathbf{W}^{(\ell)}$ acts as a more stringent regularization that prevents overfitting, when compared with the commonly used L_2 norm (Gaussian prior). Furthermore, shrinkage priors have the effect of being able to effectively switch off the elements of $\mathbf{W}^{(\ell)}$, which benefits interpretability and helps to infer the number of units needed to represent the data. In our experiment, we observe that the DPFA-SBN model with the Student's *t* prior on $\mathbf{W}^{(\ell)}$ achieves a better test perplexity when compared with its counterpart without shrinkage.

b) RCV1-v2 & Wiki: We present results for the *RCV1-v2* and *Wikipedia* corpora in Table III. Gibbs sampling in such setting is prohibitive, thus not discussed. First, we explore the effect of utilizing a larger deep network. For our DPFA-SBN model using the SGNHT algorithm, we can see

T1	T3	T8	T9	T10	T14	T15	T19	T21	T24
year	people	group	world	evidence	game	israel	software	files	team
hit	real	groups	country	claim	games	israeli	modem	file	players
runs	simply	reading	countries	people	win	jews	port	ftp	player
good	world	newsgroup	germany	argument	cup	arab	mac	program	play
season	things	pro	nazi	agree	hockey	jewish	serial	format	teams
T25	**T26**	**T29**	**T40**	**T41**	**T43**	**T50**	**T54**	**T55**	**T64**
god	fire	people	wrong	image	boston	problem	card	windows	turkish
existence	fbi	life	doesn	program	toronto	work	video	dos	armenian
exist	koresh	death	jim	application	montreal	problems	memory	file	armenians
human	children	kill	agree	widget	chicago	system	mhz	win	turks
atheism	batf	killing	quote	color	pittsburgh	fine	bit	ms	armenia
T65	**T69**	**T78**	**T81**	**T91**	**T94**	**T112**	**T118**	**T120**	**T126**
truth	window	drive	makes	question	code	children	people	men	sex
true	server	disk	power	answer	mit	father	make	women	sexual
point	display	scsi	make	means	comp	child	person	man	cramer
fact	manager	hard	doesn	true	unix	mother	things	hand	gay
body	client	drives	part	people	source	son	feel	world	homosexual

TABLE II

TOP WORDS FROM THE 30 TOPICS CORRESPONDING TO THE GRAPH IN FIGURE 3, LEARNED BY DPFA-SBN FROM THE *20Newsgroup* CORPUS.

that making the network 8 time larger decreases the test perplexities by 155 and 84 units on *RCV1-v2* and *Wikipedia*, respectively. This demonstrates the ability of our stochastic inference algorithm to scale up both in terms of model and corpus size.

TABLE III

TEST PERPLEXITIES ON *RCV1-v2* AND *Wikipedia*. "DIM" REPRESENTS THE NUMBER OF HIDDEN UNITS IN EACH LAYER, STARTING FROM THE BOTTOM. ($\diamond$) REPRESENTS THE *base tree* SIZE IN NHDP.

MODEL	METHOD	DIM	RCV	WIKI
DPFA-SBN	SGNHT	1024-512-256	964	**770**
DPFA-SBN	SGNHT	512-256-128	1073	799
DPFA-SBN	SGNHT	128-64-32	1143	876
DPFA-RBM	SGNHT	128-64-32	**920**	942
DPFA-SBN	BCDF	128-64-32	1152	986
LDA	BCDF	128	1179	931
NB-FTM	BCDF	128	1155	991
RSM	CD5	128	1171	1001
NHDP	SVB	$(10,5,5)^{\diamond}$	1041	–

Both SBN and RBM can be utilized as the building block in our deep specification. For the *RCV1-v2* corpus, our best result is obtained by utilizing a three-layer deep Boltzmann machine.

However, for the *20 Newsgroups* and *Wikipedia* corpora, with the same size model, we found empirically that the deep SBN achieves better performance.

Compared with nHDP, our DPFA models define a more flexible prior on topic interactions, and therefore in practice we also consistently achieve better perplexity results. We further show test perplexities as a function of documents processed during model learning in Figure 2. As can be seen, performance smoothly improves as the amount of data processed increases.

c) Sensitivity analysis: We examined the sensitivity of the model performance with respect to batch sizes in SGNHT on the three corpora considered. We found that overall performance, both convergence speed and test perplexity, suffer considerably when the batch size is smaller than 10 documents. However, for batch sizes larger than 50 (100 for *RCV1-v2*) we can obtain performances comparable to those shown in Tables II and III. Additional details including test perplexity traces as a function of documents seen by the model are presented in the Supplementary Material.

3) Visualization: We can obtain a visual representation of the topic structure implied by the deep component of our DPFA model by computing correlations between topics using the weight matrices, $\mathbf{W}^{(\ell)}$, learned by DPFA-SBN, *i.e*, we evaluate the covariance $\mathbf{W}^{(1)}\mathbf{W}^{(2)}(\mathbf{W}^{(1)}\mathbf{W}^{(2)})^{\top}$, then scale it accordingly. Figure 3 shows a graph for a subset of 30 topics (nodes), where edge thickness encodes correlation coefficients and we have chosen, to ease visualization, to show only coefficients larger than 0.85. In addition, Table II shows the top words for each topic depicted in Figure 3. We see three very interesting subgraphs representing different categories, namely, sports, computers and politics/law.

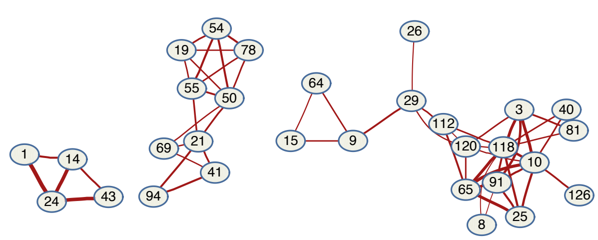

Fig. 3. Graphs induced by the correlation structure learned by DPFA-SBN for the *20 Newsgroups*. Each node represents a topic with top words shown in Table II.

Complete tables of topics' top words and graphs for the three corpora considered are presented in the Supplementary Material.

III. BAE SYSTEMS EVALUATION REPORT

Evaluation Cycle 3 Report

1 Overview

In MSEE Evaluation Cycle 3, BAE Systems worked with the performers to evaluate their performance on the Phase 3 Testing Data. The evaluation was performed entirely on-line, with the EES hosted on BAE Systems servers and the MSEE performer SUTs accessing the EES over the Internet. Based on the nature of the evaluation and implementation of the EES and client interfaces, travel to performer sites for Evaluation Cycle 3 was deemed unnecessary by DARPA.

Only the UCLA performer team participated in Evaluation Cycle 3, and this report provides details on their performance. MIT and Brown did not participate in Evaluation Cycle 3.

2 EES Architecture

The EES (Evaluation Execution System) is implemented as a web API (application programming interface). Communication between the EES web server (located at BAE Systems) and the SUT clients (located at the performer sites) is conducted over the Internet using HTTP (hypertext transfer protocol) to send and receive XML (extensible markup language) documents.

This web-based architecture provides two key benefits compared to traditional APIs that require linking to a binary library:

1. The API is agnostic to operating system and programming language. Since the API is built around two standards – HTTP and XML – with widespread support, the performers are free to use whatever OS or programming language they wish.
2. The SUT and EES are not required to run on the same computer, which enables performer evaluations over the Internet.

The expected interaction between the EES server and the SUT client during an evaluation is as follows:

1. SUT performs POST request to EES server to create a new session
   - ➔ EES responds with a session description document
2. SUT performs a GET request to EES server to get the next SOC (scene observation collection)
   - ➔ If EES responds that there are no more SOCs: exit (end of evaluation)
   - ➔ Else if EES responds with an SOC description document: continue
3. SUT performs a GET request to EES server to get the next storyline
   - ➔ If EES responds that there are no more storylines: goto 2.
   - ➔ Else if EES responds with a storyline description document: continue
4. SUT performs a GET request to EES server to get the next query
   - ➔ If EES responds that there are no more queries: goto 3.

→ Else if EES responds with a query document: continue

5. SUT performs a PUT request to EES server containing its answer to the query
   → EES acknowledges receipt of answer.
6. SUT performs a GET request to EES server to get the assessor response
   → EES responds with assessor response document.
7. Goto 4.

The evaluation session is complete when the SUT has responded to all queries for all SOCs. If the SUT sends a request that the EES does not understand, the EES will respond with an HTTP error code. If the SUT sends an XML document that is invalid according to the XML schema, the EES will respond with an HTTP error code.

More details about the operation of the EES and the EES-SUT API may be found in the *MSEE Interface Control Document*.

3 Phase 3 Testing Data

The Phase 3 Testing Data consists of four scene observation collections (SOCs), described below. These SOCs were designed to cover a variety of scenes, activities, and sensor types. By request of UCLA, one SOC from the Phase 2.2 testing data (the SIG Office) was also included in the Phase 3 evaluation. New queries were written for the SIG Office SOC for use in the Phase 3 evaluation.

3.1 SIG Parking Lot #1 (4 January 2014)

This is the first of two Phase 3 SOCs staged in the SIG Parking Lot. Approximately 20 actors participated in this SOC.

Activities in this SOC mostly center around sports and games, which created a variety of interesting motions and relationships with a great deal of occlusion and dynamism. Examples of scripted or quasi-scripted activities in this SOC include:

- Selection of teams by two captains
- Actors line up at a “concession stand” to buy various items.
- A Jeep is disassembled (roof, doors, etc. removed) and reassembled.
- An actor pushes a shopping cart around the parking lot, collecting items and trying to sell them to people.
- The actors play a short game of dodgeball.
- The actors play a short game of kickball.
- The actors participate in a relay bicycle race.

More detail about the activities in this SOC may be found in the script document included in the reference data that was distributed with the SOC data.

The duration of the first SIG Parking Lot SOC was 47 minutes. A total of ten EO cameras and one IR camera were used to record the SOC events. Three stationary EO cameras were located on the roof of the building looking down at the parking lot. Five stationary EO cameras and one stationary IR camera were located at ground level. Two mobile EO cameras were also utilized, both attached to handle bars of bicycles.

In addition to the video data, scene descriptive text (SDT) and Ground Moving Target Indicator (GMTI) radar data were also provided for this SOC.

Distinguishing challenges of this dataset include:

- **Lighting conditions**: The time of year (winter) and day of the collection presented harsh lighting conditions manifested as specularities, high contrast, and saturation.
- **Sparse sensor coverage**: the large collection area combined with a limited number of sensors resulted in sparse sensor coverage of the AOR. Most activities are only visible at reasonable resolution in one or two cameras.

Figure 1 shows the area of interest and approximate camera placement of stationary cameras for the first SIG Parking Lot SOC. The area outlined in red is the area of interest of the SOC. Compared to the parking lot SOC from Phase 2.2, the area of interest for this SOC is greatly compressed, covering about half the area used in the Phase 2.2 parking lot SOC.

Figure 1: Area of interest and camera placement for the first SIG Parking Lot SOC. "GR1" denotes the position of the GMTI radar.

3.2 SIG Parking Lot #2 (18 October 2014)

This is the second of two Phase3 SOCs staged in the SIG Parking Lot. Approximately 18 actors participated in this SOC.

Activities in this SOC include a mix of scripted activities with unscripted "background" activities. Background activities consist of actors driving and parking their cars, walking and/or biking through the parking lot, entering and exiting the building, walking a dog, etc. Examples of scripted or quasi-scripted activities include:

- Two individuals rendezvous in the parking lot to exchange packages. At different times packages are exchanged by passing them through car windows, by meeting outside the cars and exchanging packages hand-to-hand, and by allowing person A to retrieve a package from the trunk of person B's unattended car.
- An automobile "breaks down" and actors help to perform maintenance on the automobile.
- An actor "steals" an item from an unlocked car.
- An actor is forcibly escorted from the building and into a car, which drives away.
- An object is buried.

- Luggage is left unattended.
- An actor has a brief, non-physical, altercation with another actor.

More detail about the activities in this SOC may be found in the script document included in the reference data that was distributed with the SOC data.

The duration of the second SIG Parking Lot SOC was 22 minutes. A total of ten EO cameras and one IR camera were used to record the SOC events. Two stationary EO cameras were located on the roof of the building looking down at the parking lot. Seven stationary EO cameras and one stationary IR camera were located at ground level. One mobile EO camera was also utilized; the mobile camera was hand-carried by an actor through the scene.

In addition to the video data, scene descriptive text (SDT) and GMTI radar data were also provided for this SOC.

Distinguishing challenges of this dataset include:

- **Sensor data time synchronization**: Most objects come into the AOR already moving, and they keep moving while they appear on the cameras. Object definitions use a single pixel point. If the time syncing used by SUT is different than that of the SIG system, the SUT may miss identifying objects, causing all the subsequent queries dependent on those objects to fail.

Figure 2 shows the area of interest, approximate placement of stationary cameras, and approximate placement of GMTI radar for the second SIG Parking Lot SOC. The area outlined in red is the area of interest of the SOC. Compared to the first parking lot SOC from Phase 3, the area of interest for this SOC is somewhat smaller.

Figure 2: Area of interest and camera and GMTI radar placement for the second SIG Parking Lot SOC.

3.3 Duke Pratt Garden (20 September 2014)

This SOC was collected in a small garden outside a building at Duke University. Approximately 17 actors participated in this SOC.

Activities in this SOC mostly involve three quasi-scripted vignettes:

- Exercise class: an actor leads the other actors in an “exercise class,” including stretching, jogging, calisthenics, etc.
- Fashion show: actors show off various outfits while other actors look on.
- Sports activities: actors engage in various sports activities, including bike riding, disc golf, baseball, and parkour.

More detail about the activities in this SOC may be found in the document “MSEE_PrattGarden_20140920_Scripts.docx,” which is included in the reference data that was distributed with the SOC data.

The duration of this SOC is 34 minutes. A total of seven EO and one IR cameras were used to record the SOC events. One stationary EO camera was placed on the second story of a nearby building, looking down at the area of interest. Five stationary EO cameras and one stationary IR camera were positioned at ground level around the area of interest to capture the scene from multiple angles. One mobile EO camera was carried through the scene by an actor.

In addition to the cameras, scene descriptive text was also provided for this SOC.

Distinguishing challenges of this dataset include:

- **Lighting conditions**: There is poor lighting and shade in one corner of the AOR which creates contrast issues for several sensors.
- **Sparse sensor coverage**: many occluding features combined with a limited number of sensors resulted in sparse sensor coverage of the AOR. Many corners of the AOR and activities are only visible at reasonable resolution in a single sensor.
- **Single overhead sensor**: There is a single overhead sensor, HC4, which provides context for the overall AOR. Unfortunately because this camera was stationed indoors looking outside, it suffers from slight glare and blurriness from the window.
- **Dropped frames**: Two IP cameras intermittently drop frames in the second temporal half of the collection.

Figure 3 shows the area of interest and approximate placement of stationary cameras for the Pratt Garden SOC. The area outlined in red is the area of interest of the SOC.

Figure 3: Area of interest and camera placement for the Pratt Garden SOC.

3.4 Duke Schiciano Auditorium (22 February 2014)

This SOC was collected inside Duke's Schiciano Auditorium, and the associated lobby area. Approximately 21 actors participated in this SOC.

Activities in the SOC center around a simulated academic conference. Among the quasi-scripted activities for this SOC are:

- Registration: actors approach a registration desk in the lobby, check in, and proceed into the auditorium.
- Presentation: two actors give a presentation, while the other actors listen and perform various activities.
- Simulated fire alarm: during the presentation, a fire alarm is simulated. All actors leave the auditorium, then return later.
- Simulated panic: everyone exits the building in a panicked manner.

More detail about the activities in this SOC may be found in the document "MSEE_Schiciano20140222_Scripts.docx," which is included in the reference data that was distributed with the SOC data.

The duration of this SOC is 40 minutes. A total of eleven cameras (ten EO and one IR) were used to record the SOC events. Four stationary EO cameras were placed in the lobby area. Three stationary EO cameras were placed in the auditorium. A stationary IR camera was placed in a hallway adjacent to the auditorium so that it had a view of the lobby. A stationary EO camera was also placed in the hallway, with a view down the hallway. Two mobile EO cameras were utilized. One EO camera was attached to a cart used at the registration table. Another mobile EO camera was hand carried by an actor through the scene.

In addition to the cameras, scene descriptive text was also provided for this SOC.

Distinguishing challenges of this dataset include:

- **Novelty**: This AOR was not used for any Phase 2 collection and as such, may present novel challenges to the SUT.
- **Occlusions**: The SOCs dense collection area and environment, including walls and columns, creates more occlusions as compared to other SOCs. People and chairs in the auditorium are partially occluded by tables.
- **Segmented AOR:** The multiple rooms within the AOR can make tracking and spatial awareness more difficult for the SUT.

Figure 4 shows the area of interest and approximate camera placement for the SOC. The area of interest for this SOC includes the "Fitzpatrick Center Lobby," the "Schiciano Auditorium Side B," and the "Access Hallway." "Auditorium Side A" was not included in the area of interest and was not recorded.

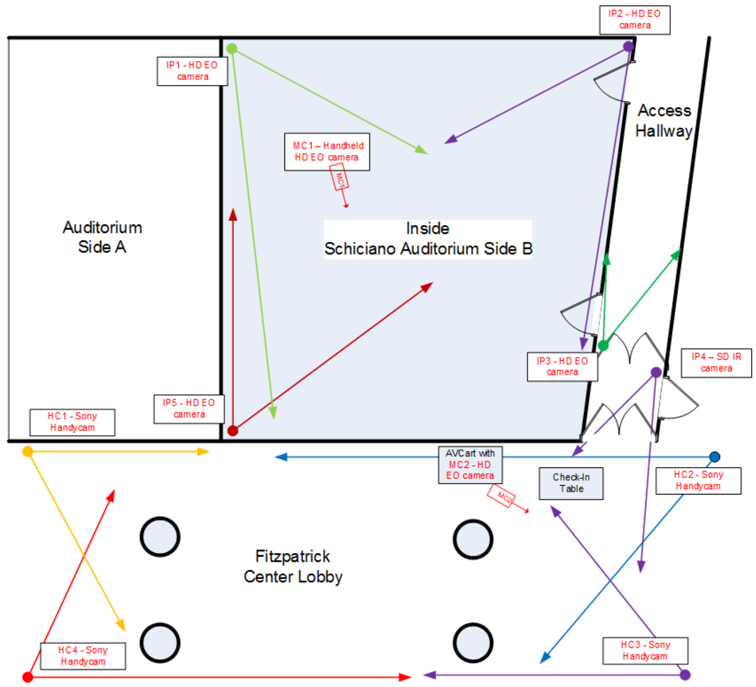

Figure 4: Area of interest and camera placement for Duke Schiciano Auditorium SOC.

3.5 SIG Office (4 September 2013)

The SIG Office SOC was created during Phase 2 and used in the Phase 2.2 evaluation. For the Phase 3 evaluation, UCLA requested that this SOC be included again, as they claimed to have done quite a bit of work to improve their performance on this SOC in particular. BAE Systems developed a new set of queries for use with the SIG Office SOC for the Phase 3 evaluation. Because all Phase 2 performers have previously "seen" the SIG Office SOC in the Phase 2 evaluation, the SIG Office SOC is evaluated separately from the Phase 3 SOCs in this report.

The SIG Office SOC consists of video collected inside SIG's office suite in Durham, NC. The SOC contains video from three main rooms in the office (the reception area, the break room, and the conference room) as well as the hallways connecting these areas. A total of 23 actors participated in this SOC.

Activities in the SOC include a mix of activity appropriate to an office environment, as well as activities designed to be DoD-relevant. Some activities were scripted; others, such as the pizza lunch, were not. Examples of the activities in the SIG Office SOC include:

- A pizza lunch in the conference room, where approximately 15 people eat pizza, mingle, chat, and play games.
- An actor leaves a package in a room, which is later recovered by another actor.
- An actor surreptitiously "steals" another actor's backpack.

- Packages are carried throughout the three rooms, changing hands several times.
- A small meeting between three people takes place in the break room.
- A larger, more formal meeting takes place in the conference room, including a presentation using a projector and several actors interacting at a whiteboard.

More detail about the activities in this SOC can be found in the document "MSEE_SIG_Office_Scripts.docx," which is included in the reference data that was distributed with the SIG Office SOC.

The duration of the SIG Office SOC was 1 hour, 32 minutes, 40 seconds. Twelve EO cameras, all stationary, were used to record the events. Figure 5 shows the layout of the rooms and approximate locations of all cameras.

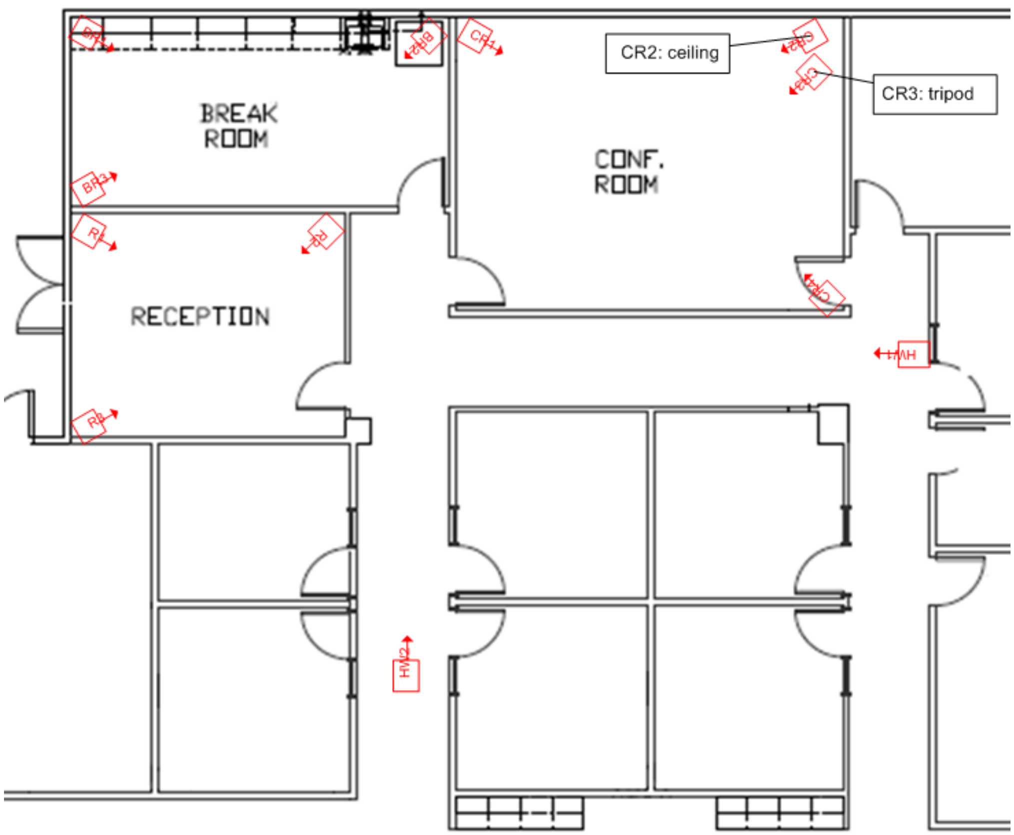

Figure 5: Layout of rooms and cameras used in the SIG Office SOC

4 Phase 3 Testing Queries

4.1 Query Strategy for Phase 3

Based on lessons learned during the Phase 2.2 evaluation, a number of changes were made to the strategy used to develop queries for Phase 3.

New approach to defining objects used in queries. In Phase 2.2, the only approach for selecting objects (people, vehicles, etc.) to be used in queries was to define a set of objects. This frequently required stringing together multiple conditions to select an object or objects we wanted to ask about. If even one condition was not understood correctly by the SUT, the SUT might use a different object or set of objects in its response to the query, or even not be able to respond to the query at all.

In Phase 3, we introduced the concept of an object definition. Phase 3 queries may specify objects by giving the SUT the coordinates of a pixel where the object appears at a particular time in a particular camera. The SUT can then specify whether or not it can detect an object of the specified type at that time and place. If it can, then the queries can continue with confidence that the SUT and EES are "discussing" the same object. If the SUT cannot identify the object, the EES can skip all queries concerning that object.

Simpler queries. Phase 2.2 queries frequently involved a conjunction of many predicates. If the SUT's knowledge was incorrect about any one predicate, it could get the answer to the entire query wrong. If the SUT was wrong on a query, it would be difficult for the evaluator to know which of the multiple predicates it was wrong about. Conversely, it is difficult to understand why an SUT may have correctly answered a complex, multi-predicate query.

In Phase 3, queries have been made much simpler to address these challenges of performance interpretation. The use of defined objects (described above) has helped achieve the goal of simpler queries. Most Phase 3 queries involve only a single predicate, operating on one or more defined objects.

Larger number of queries. Because the Phase 3 queries are simpler, we can produce and test on a greater number of queries. This should allow us to test the SUTs on each concept under a greater range of operating conditions.

Reduced set of predicates emphasized in queries. There are 148 predicates in the MSEE Formal Language Specification. It is not feasible to test SUT performance on all 148 predicates across multiple operating conditions. For the Phase 3 evaluation, we chose to emphasize some predicates and de-emphasize others (meaning the de-emphasized predicates were rarely used or not used at all). Combined with a larger number of queries, this allows us to exercise most of the emphasized predicates multiple times. Table 1 shows the emphasized predicates, broken down by category.

Table 1: Phase 3 emphasized predicates, by category

Categories	Description
Classification	Predicates related to classes of objects: person, male, female, animal, vehicle, two-wheeled-vehicle, automobile, small-object, luggage, package, ball, disc, clothing, hat, top-wear, building, room, table, chair.
Part Of	Predicates related to "part-of" hierarchies: part-of, building, door, room, wall, floor, person, head, arm, hand, lower-body, vehicle, door, trunk, hood, wheel.
Spatial	Predicates related to spatial reasoning: clear-line-of-sight, occluding, closer, father, facing, facing-opposite.
Attributes	Predicates related to attributes of single objects: open, closed, sitting, standing, pointing, crawling, walking, running, talking.
Relationships	Predicates related to relationships between two or more objects: same-object, on, together, touching, inside, outside, below, driving, entering, exiting, carrying, loading, unloading, mounting, dismounting, donning, doffing, throwing, catching, putting-down, picking-up, dropping.
Tracking	Predicates related to tracking: starting, stopping, moving, stationary, turning, turning-right, turning-left, u-turn, same-motion, opposite-motion, following, passing.

4.2 Phase 3 Query Summary

For the five SOCs used in the Phase 3 evaluation, BAE Systems developed a total of 1,060 queries. Natural language versions of all queries, the "correct" responses as determined by a human assessor, and the responses given by UCLA's SUT can be found in Appendix B: Phase 3 Testing Queries and Answers.

Table 2 shows the breakdown of number of queries by SOC:

Table 2: Number of queries by SOC

SOC	# Queries
SIG Office	108
SIG Parking Lot #1	247
SIG Parking Lot #2	236
Duke Pratt Garden	215
Duke Schiciano Auditorium	254
Total	**1,060**

Table 3 shows the number of queries that use predicates from each category. The "object definition" category consists of queries that are used only to verify that the SUT was able to identify a defined object that will be used in subsequent queries. Object definition queries are of the general form "Is object x an object?" For example: "Is obj-person-1 a person?" or "Is obj-vehicle-1 a vehicle?" If the SUT could not successfully identify the object, it should respond with an "UnknownObject" flag, which will let the EES know to skip any future queries about that object. The SUT will not be penalized for failing to answer these queries, as unanswered queries are not counted in the results.

(It is possible that the SUT could respond “true” to all object definition queries without even trying to process the data, since the answer will always be true by definition. However, an SUT that tries to “cheat” by answering “true” to an object definition query when it cannot actually identify the object will then be faced with a sequences of queries about that object – which it will likely perform poorly on.)

It is not possible to formulate a query without using at least one predicate from the classification category – classification predicates are the only way to define the object or objects involved in the query. Therefore, if a query uses *only* classification predicates, it is counted in the classification category. If the query uses classification predicates, plus predicates from one other category X, it is counted in category X. If the query uses classification predicates, plus predicates from two or more additional categories, it is not assigned a category. With the emphasis on simpler queries for Phase 3, there are only five such queries that use multiple categories of predicates.

Table 3: Number of queries by category

Category	# Queries
Object definition	243
Classification	71
Part Of	93
Spatial	58
Attributes	165
Relationships	291
Tracking	134
Multiple predicate categories	5

The number of predicates used in each query can serve as a proxy for complexity of the query. Number of predicates is not a perfect measure of the complexity of a query because not all predicates are equally complex, and other factors affect query complexity (such as the size of the temporal and spatial windows that must be considered in answering the query). Nevertheless, in a later section we will evaluate SUT performance relative to the number of predicates in each query. Figure 6 is a histogram of the number of predicates used in queries; Table 4 shows the same data in tabular form.

Most queries have either 1, 2, or 3 predicates. This is a natural result of the choice in Phase 3 to simplify the queries. The queries with 1, 2, or 3 predicates can mostly be explained as follows:

- 1 predicate: These are queries that deal only with the predicates for the various types of objects (people, automobiles, etc.). Most of these queries (243) are object definition queries; the others deal with counting objects (e.g. “how many people are in the scene?”).

- 2 predicates: These queries are mostly queries involving unary predicates operating on an object. One predicate is used to define the object (usually person or automobile), and the unary predicate is the second predicate involved.
- 3 predicates: These queries are mostly queries involving binary predicates operating on two objects. Two predicates are used to define the operands, and the binary predicate is the third predicate involved.

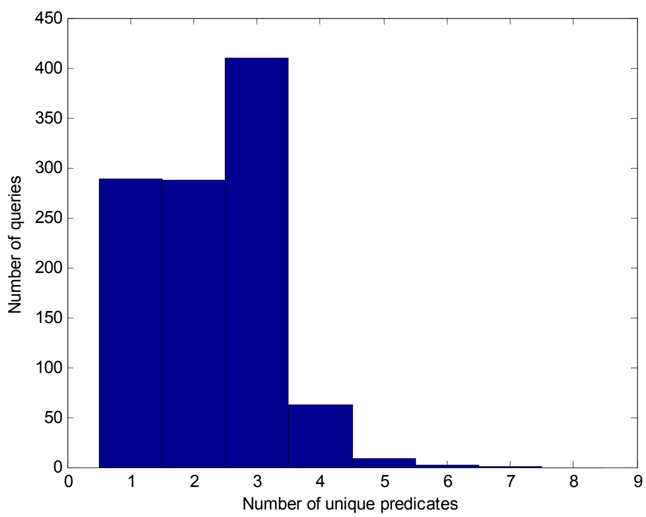

Figure 6: Histogram of number of predicates used in queries

Number of predicates	Number of queries
1	289
2	287
3	410
4	62
5	9
6	2
7	1

Table 4: Number of predicates used in queries

5 Evaluation Timeline

On April 3rd, 2015 the Phase 3.0 Testing Data stored on an external hard disk drive was shipped to UCLA, DARPA, and AFRL, with an expected arrival date no later than April 6th. UCLA was allotted two weeks to perform data preprocessing required by their SUT. Table 6 provides details on the preprocessing performed and the associated time durations, as reported by UCLA. BAE Systems exposed the EES interface for Phase 3 Evaluation on 12:01 AM EDT April 20, 2015.

UCLA started its lone EES session at 5:30 PM EDT on April 21, and completed it at 8:55 PM EDT on April 28. After processing the queries associated with the first SOC, “soc-sig-office-2013-09-04-testing”, UCLA paused the evaluation at 11:52 PM EDT on 4/22/15 to address interfacing and other SUT issues. Note that the “soc-sig-office-2013-09-04-testing” SOC is not part of the Phase 3 testing data sets. UCLA resumed the evaluation at 3:27 AM EDT 4/26/15 and completed the evaluation at 8:55 AM EDT 4/28/15.

Table 5: Summary of UCLA SUT Data Preprocessing (as reported by UCLA)

Reported Data Preprocessing Metric	SIG Office 2013-09-04	SIG Parking Lot 2014-01-04	SIG Parking Lot 2014-10-18	Pratt Garden 2014-09-20	Schiciano Auditorium 2014-02-22
Video duration	17:35:36	8:14:42	4:27:44	4:15:56	8:53:24
Total # of frames	2,486,289	888,053	481,414	458,629	959,216
Detection					
Human bounding boxes	1,341,704	1,885,106	487,808	2,718,738	2,433,349
Car bounding boxes	N/A	204,238	1,212,573	N/A	N/A
Bicycle bounding boxes	N/A	4,801	13,192	43,291	N/A
Processing time	unreported	~16 hours	~19 hours	~13 hours	~15 hours
Tracking					
Generated human tracks	2227	17,547	3,061	16,964	11,860
Generated car tracks	N/A	437	1,490	N/A	N/A
Generated bicycle tracks	N/A	92	186	321	N/A
Processing time	unreported	~34 hours	~9 hours	~33 hours	~22 hours
Attributes					
Processed human bounding boxes	1,340,395	1,236,245	486,879	2,270,448	2,429,963
Generated attribute boxes	4,057,955	4,442,954	1,574,064	8,194,540	7,667,580
Processing time	~18 hours	~34 hours	~6 hours	~32 hours	~22 hours
Action					
Processed bounding boxes	1,330,240	1,884,974	487,614	2,718,600	2,433,209
Processing time	~20 hours	~34 hours	~6.5 hours	~33 hours	~23 hours
Behavior					
Processed bounding boxes	1,217,572	1,993,022	1,664,075	2,594,942	2,246,625
Processing time	~13 min.	~27 min.	~17 min.	~33 min.	~33 min.

6 Evaluation Performance

Because the Phase 3 evaluation included not only the Phase 3 data (which had never been seen by the performers before) but also one repurposed Phase 2.2 SOC (which *had* been seen by the performers before), in presenting the results we will distinguish between performance on Phase 3 data only, performance on the Phase 2.2 SIG Office SOC only, and overall performance on *all* data.

Table 6 summarizes the performance of UCLA's SUT on Phase 3 data, Phase 2.2 data, and both datasets combined. Note that object definition queries are excluded from the results in this table (see section 6.5 for more on the SUT performance on object definition queries).

Table 6: UCLA SUT performance metrics for all queries (excluding object definition queries)

Metric	Phase 3 SOCs	Phase 2.2 SIG Office	Overall
Number of queries	709	108	817
Number of responses	459 (65%)	79 (73%)	538 (66%)
Error rate	0.370 (170/459)	0.215 (17/79)	0.348 (187/538)
Confidence error	0.089	0.087	0.089
Brier score	0.323	0.211	0.307

Considering the Phase 3 data only, 709 queries (not including object definition queries) were available. The SUT responded to 459 of these queries (65%). Of the queries the SUT responded to, 289 responses were correct (63%). An additional 243 object definition queries were presented; the SUT was able to identify 197 of the objects (81%).

Considering the Phase 2.2 SIG Office SOC only, 108 queries were available. The SUT responded to 79 of these queries (73%). Of the queries the SUT responded to, 62 responses were correct (78%). There were no object definition queries for this SOC.

Considering all SOCs (both Phase 3 and Phase 2 together), 817 queries were available (not including object definition queries). The SUT responded to 538 of these queries (66%). Of the queries the SUT responded to, 351 responses were correct (65%). An additional 243 object definition queries were presented; the SUT was able to identify 197 of the objects (81%).

The following sections go into more details about the performance of UCLA's system on the Phase 3 evaluation. The metrics examined in these sections are based on the document "Evaluation Metrics for the MSEE Program" version 0.1 dated 14 February 2013, produced by the AFRL COMPASE Center.

6.1 SUT Confidence in Answers

It is important to evaluate not just the accuracy of the SUT answers, but also the accuracy of the SUT confidences in their answers. In general, the SUT reported very high confidences in its answers. For UCLA's SUT, 46.0% of answers had a confidence above 0.9, and 78.6% of answers had a confidence above 0.6.

Intuitively, SUT answers with higher confidences should be more likely to be correct. Figure 7 explores this concept for the UCLA SUT. The horizontal axis represents SUT answer confidence. The vertical axis shows the error rate for answers having a confidence greater than or equal to the specified level. Initially, error rate actually increases as confidence increases, but for the very highest confidence values error rate does go down significantly.

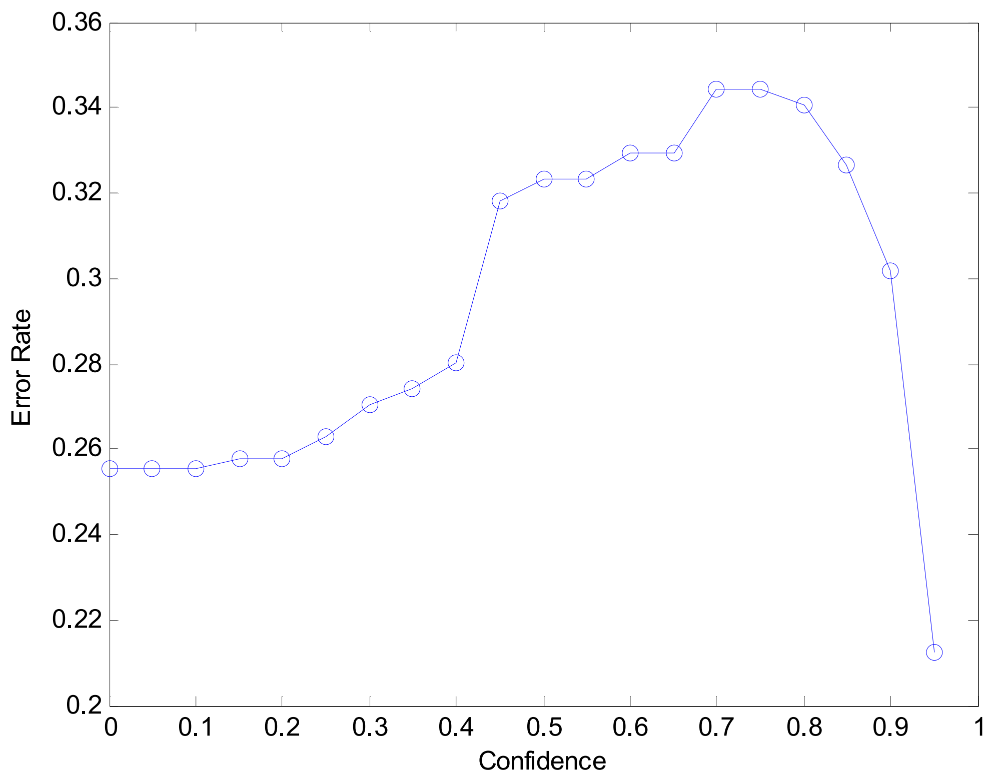

Figure 7: Answer confidence vs. error rate for the UCLA SUT

To help explore the relationship between confidence and accuracy more formally, "Evaluation Metrics for the MSEE Program" introduces the concept of a "declaration," which is defined as follows:

- Given a confidence threshold C, the SUT declares an answer when both of the following are true:

- o The SUT responded to the query.
- o The SUT's confidence is >= C.

For the metric results presented below, each metric is computed at three different confidence levels, C = 0.0, C = 0.6, and C= 0.9. For each confidence level, the metric is computed only for the queries where the SUT declared a response at that confidence level. Intuitively, one expects the metrics to improve as the confidence level rises.

6.2 Performance by SOC

Table 7 shows the performance of the UCLA SUT broken down by SOC. "Object definition" queries are excluded from the metrics reported in this table. Note that the "SIG-Office 2013-09-04" SOC was used in the Phase 2.2 evaluation, though the queries presented in the Phase 3 evaluation are new. Note that object definition queries are not included in these results.

Table 7: Performance metrics by SOC

Metric	SIG Parking Lot 2014-01-04	SIG Parking Lot 2014-10-18	Pratt Garden 2014-09-20	Schiciano Auditorium 2014-02-22	SIG Office 2013-09-04
Number of queries	184	165	161	199	108
Number of responses	96 (52.2%)	99 (60.0%)	128 (79.5%)	136 (68.3%)	79 (73.1%)
Confidence >= 0.0					
Number of declarations	96	99	128	136	79
Declaration rate	1	1	1	1	1
Error rate	0.385	0.374	0.414	0.316	0.215
Confidence error	0.121	0.076	0.051	0.112	0.087
Brier score	0.332	0.340	0.307	0.320	0.211
Confidence >= 0.6					
Number of declarations	67	81	114	99	59
Declaration rate	0.698	0.818	0.891	0.728	0.747
Error rate	0.433	0.395	0.421	0.364	0.203
Confidence error	0.016	0.009	0.026	0.009	0.008
Brier score	0.343	0.343	0.314	0.312	0.183
Confidence >= 0.9					
Number of declarations	28	63	24	96	50
Declaration rate	0.292	0.636	0.188	0.706	0.633
Error rate	0.429	0.381	0.333	0.365	0.200
Confidence error	0.010	0.004	0.002	0.006	0.003
Brier score	0.362	0.340	0.304	0.315	0.185

6.3 Performance by Number of Predicates

The following tables summarize performance metrics for sets of queries based on the number of predicates used in the query. Intuitively, we expect queries with more predicates to be more complex and therefore to have higher error rates. Separate tables are presented for results using Phase 3 data only (Table 8), Phase 2.2 SIG Office SOC data only (Table 9), and all data (Table 10). Once again, object definition queries are excluded from these results.

Table 8: Performance by number of predicates (Phase 3 datasets only)

Metric	N=1	N = 2	N = 3	N = 4	5 <= N < 10
Number of queries	31	255	373	45	5
Number of responses	31	191	218	15	4
Confidence >= 0.0					
Number of declarations	31	191	218	15	4
Declaration rate	1	1	1	1	1
Error rate	0.387	0.298	0.408	0.478	0.750
Confidence error	0.129	0.113	0.064	0.061	0.077
Brier score	0.255	0.298	0.348	0.391	0.464
Confidence >= 0.6					
Number of declarations	21	139	185	13	3
Declaration rate	0.677	0.728	0.849	0.867	0.750
Error rate	0.333	0.345	0.438	0.538	0.667
Confidence error	0.016	0.016	0.016	0.017	0.018
Brier score	0.249	0.278	0.359	0.425	0.538
Confidence >= 0.9					
Number of declarations	14	80	109	7	1
Declaration rate	0.452	0.419	0.500	0.467	0.250
Error rate	0.286	0.312	0.413	0.571	1
Confidence error	0.003	0.006	0.006	0.004	0.007
Brier score	0.251	0.270	0.363	0.500	0.836

Table 9: Performance by number of predicates (Phase 2.2 SIG Office SOC only)

Metric	N=1	N = 2	N = 3	N = 4	5 <= N < 10
Number of queries	15	32	37	17	7
Number of responses	15	24	25	8	7
Confidence >= 0.0					
Number of declarations	15	24	25	8	7
Declaration rate	1	1	1	1	1
Error rate	0.067	0.167	0.280	0.250	0.429
Confidence error	0.036	0.142	0.085	0.010	0.103
Brier score	0.100	0.179	0.220	0.239	0.492
Confidence >= 0.6					
Number of declarations	13	13	20	8	5
Declaration rate	0.867	0.542	0.800	1	0.714
Error rate	0.077	0.077	0.250	0.250	0.600
Confidence error	0.002	0.002	0.017	0.010	0.002
Brier score	0.076	0.072	0.211	0.239	0.547
Confidence >= 0.9					
Number of declarations	13	13	13	6	5

Declaration rate	0.867	0.542	0.520	0.750	0.714
Error rate	0.077	0.077	0.231	0.333	0.600
Confidence error	0.002	0.002	0.002	0.006	0.002
Brier score	0.076	0.072	0.211	0.304	0.547

Table 10: Performance by number of predicates (all data)

Metric	N=1	N = 2	N = 3	N = 4	5 <= N < 10
Number of queries	46	287	410	62	12
Number of responses	46	215	243	23	11
Confidence >= 0.0					
Number of declarations	46	215	243	23	11
Declaration rate	1	1	1	1	1
Error rate	0.283	0.284	0.395	0.478	0.545
Confidence error	0.099	0.116	0.066	0.043	0.094
Brier score	0.205	0.284	0.335	0.338	0.482
Confidence >= 0.6					
Number of declarations	34	152	205	21	8
Declaration rate	0.739	0.707	0.844	0.913	0.727
Error rate	0.235	0.322	0.42	0.429	0.625
Confidence error	0.01	0.014	0.016	0.014	0.008
Brier score	0.183	0.261	0.344	0.355	0.543
Confidence >= 0.9					
Number of declarations	27	93	122	13	6
Declaration rate	0.587	0.433	0.502	0.565	0.545
Error rate	0.185	0.28	0.393	0.462	0.667
Confidence error	0.003	0.005	0.006	0.005	0.003
Brier score	0.167	0.243	0.347	0.41	0.595

6.4 Performance by Predicate Category

SUT performance may be evaluated based on the predicates used within the queries, which may indicate an SUT's general capability strengths and weaknesses. The predicate categories evaluated in the Phase 3 evaluation are described in Table 1.

Separate tables are presented for results using Phase 3 data only (Table 11), Phase 2.2 SIG Office SOC data only (Table 12), and all data (Table 13).

Table 11: Performance by predicate category (Phase 3 datasets only)

Metric	classification	part of	spatial	attributes	relationships	tracking
Number of queries	56	93	50	146	246	114
Number of responses	53	65	25	93	119	101
	Confidence >= 0.0					
Number of declarations	53	65	25	93	119	101
Declaration rate	1	1	1	1	1	1
Error rate	0.283	0.292	0.560	0.290	0.471	0.366
Confidence error	0.150	0.102	0.039	0.082	0.058	0.102

Brier score	0.264	0.292	0.463	0.265	0.367	0.341
	Confidence >= 0.6					
Number of declarations	34	48	23	76	103	75
Declaration rate	0.642	0.738	0.920	0.817	0.866	0.743
Error rate	0.294	0.333	0.609	0.316	0.476	0.413
Confidence error	0.015	0.011	0.023	0.016	0.017	0.015
Brier score	0.223	0.291	0.474	0.252	0.384	0.341
	Confidence >= 0.9					
Number of declarations	22	35	14	42	47	51
Declaration rate	0.415	0.538	0.560	0.452	0.395	0.505
Error rate	0.227	0.286	0.500	0.286	0.532	0.392
Confidence error	0.004	0.003	0.007	0.006	0.007	0.008
Brier score	0.200	0.269	0.426	0.249	0.457	0.340

Table 12: Performance by predicate category (Phase 2.2 SIG Office SOC only)

Metric	classification	part of	spatial	attributes	relationships	tracking
Number of queries	8	19	45	20	0	0
Number of responses	4	16	35	8	N/A	N/A
	Confidence >= 0.0					
Number of declarations	4	16	35	8	N/A	N/A
Declaration rate	1	1	1	1	N/A	N/A
Error rate	0	0.250	0.314	0	N/A	N/A
Confidence error	0.012	0.113	0.082	0.201	N/A	N/A
Brier score	0.023	0.168	0.282	0.201	N/A	N/A
	Confidence >= 0.6					
Number of declarations	4	9	23	4	N/A	N/A
Declaration rate	1	0.562	0.800	0.500	N/A	N/A
Error rate	0	0.111	0.321	0	N/A	N/A
Confidence error	0.012	0.002	0.013	0.002	N/A	N/A
Brier score	0.023	0.103	0.281	0.002	N/A	N/A
	Confidence >= 0.9					
Number of declarations	2	9	21	4	N/A	N/A
Declaration rate	0.500	0.562	0.600	0.500	N/A	N/A
Error rate	0	0.111	0.333	0	N/A	N/A
Confidence error	0.002	0.002	0.003	0.002	N/A	N/A
Brier score	0.002	0.103	0.305	0.002	N/A	N/A

Table 13: Performance by predicate category (all data)

Metric	classification	part of	spatial	attributes	relationships	tracking
Number of queries	71	93	58	165	291	134
Number of responses	68	65	29	109	154	109
	Confidence >= 0.0					
Number of declarations	68	65	29	109	154	109
Declaration rate	1	1	1	1	1	1
Error rate	0.235	0.292	0.483	0.284	0.435	0.339
Confidence error	0.124	0.102	0.036	0.087	0.063	0.109
Brier score	0.228	0.292	0.403	0.251	0.348	0.331
	Confidence >= 0.6					
Number of declarations	47	48	27	85	131	79
Declaration rate	0.691	0.738	0.931	0.78	0.851	0.725
Error rate	0.234	0.333	0.519	0.294	0.443	0.392
Confidence error	0.012	0.011	0.021	0.015	0.016	0.014
Brier score	0.182	0.291	0.407	0.236	0.362	0.324
	Confidence >= 0.9					
Number of declarations	35	35	16	51	68	55
Declaration rate	0.515	0.538	0.552	0.468	0.442	0.505
Error rate	0.171	0.286	0.438	0.255	0.471	0.364
Confidence error	0.004	0.003	0.006	0.005	0.006	0.007
Brier score	0.154	0.269	0.373	0.224	0.41	0.315

6.5 Performance on Object Definition Queries

Object definition queries constitute a special class of queries. As described in section 4.1, object definitions were added for Phase 3 as an approach to simplify queries and allow for more precision in specifying which object is being asked about.

Objects are defined by specifying an object type and a single pixel that is part of the object. The SUT then determines if it knows of an object of the given type containing the given pixel. Objects are first introduced with a very simple query that we term an "object definition" query. The query itself is a tautology. (Specifically, it is of the form "is the object, which is defined to be of type X, of type X?") Therefore, if the SUT can identify the object at all it should always return "true" in answer to an object definition query. The EES may then proceed to ask additional questions about the object. If the SUT cannot identify the object, it should return the "unknown object" response. In that case, the EES will skip all queries related to that object.

UCLA has stated that their Phase 3 SUT will respond with "unknown object" when either:

- no object is found to match the queries (most cases); or
- multiple objects are found and the system cannot resolve which is the best match.

UCLA also reported that the following objects are not supported by their Phase 3 SUT:

- clothing-footwear (Note: this predicate was not used in the Phase 3 evaluation)

- building – wall – door (Note: this predicate was used in only five queries in the Phase 3 evaluation)
- room – wall – switch (Note: this predicate was not used in the Phase 3 evaluation)
- vehicle – fender (Note: this predicate was not used in the Phase 3 evaluation)
- room – wall – art (Note: this predicate was not used in the Phase 3 evaluation)

Though not explicitly listed as unsupported by UCLA, UCLA implied that the “ball” and “disc” object types were “too small and move too fast to detect”.

The Phase 3 testing queries included 243 object definition queries. Of these 243, UCLA’s SUT gave a “true” response (indicating a confidence that it could correctly identify the described object) in 197 cases (81%). UCLA’s SUT responded with “unknown object” in 45 cases (19%).

Curiously, UCLA’s SUT responded with “false” in one case. Since a “false” response is never appropriate for an object definition query, we suspect this might have been the result of a software logic error in the SUT. Figure 8 shows the object (in this case a person) in question.

Figure 8: UCLA's SUT responded "false" to an object definition query involving this woman.

Table 14 shows SUT performance on object definition queries by the type of object. Almost half of all objects used were people, and the SUT had a high success rate at detecting people (87.5%). About one-eighth of all objects were automobiles, which the SUT detected at a rate of 77.4%.

Table 14: Summary of performance on object definitions queries, by object type.

Object type	# definitions	# detected	Detection rate
Person	120	105	0.875
Automobile	31	24	0.774
Head	17	16	0.941
Arm	12	6	0.500
Lower-body	12	11	0.917
Luggage	10	7	0.700
Door	7	4	0.571
Hand	5	3	0.600

Object type	# definitions	# detected	Detection rate
Trunk	5	5	1.000
Small-object	4	4	1.000
Wheel	4	2	0.500
Hat	3	2	0.667
Hood	3	1	0.333
Disc	2	0	0.000
Two-wheeled-vehicle	2	2	1.000
Animal	1	1	1.000
Ball	1	0	0.000
Room	1	1	1.000
Table	1	1	1.000
Tool	1	1	1.000
Wall	1	1	1.000
Total	**243**	**197**	**0.811**

Examples of undetected object definitions follow, where the green circle denotes the pixel location of the object definition.

Figure 9: Object definition of obj-jeep in storyline-Tracking-Automobiles, SIG Parking Lot 2014-01-04

Figure 10: Object definition of obj-person3 in storyline-Person-Attributes, SIG Parking Lot 2014-01-04. obj-person3 is in the center if the image.

Figure 11: Object definition of obj-person4 in storyline-attributes, SIG Parking Lot 2014-10-18

Figure 12: Object definition of obj-ball1 in storyline-sports, Pratt Garden 2014-09-20. obj-ball1 is held by the individual in the center of the FOV.

Figure 13: Object definition of obj-presenter1 in storyline-presentation2, Schiciano Auditorium 2014-02-22.

6.6 Unanswered Queries

In the Phase 3 evaluation, the UCLA SUT was unable to respond to a total of 279 queries – 29 for the "SIG Office 2013-09-04" SOC and 250 for the Phase 3 SOCs. All of these queries are omitted from the computation of overall SUT performance metrics.

An SUT may fail to respond to a query in one of two ways:

- The SUT received the query and sent one of the "unable to respond" codes described below. This happened 107 times in the Phase 3 evaluation (38% of non-responses).
- The SUT indicated it could not identify an object involved in the query (by responding "unknown object" to an object definition query), and the EES did not send the query. Of the 279 non-responses, 172 (62%) were queries skipped because a prerequisite object definition was not detected by the SUT.

The format of the answer documents that the SUT sends to the EES contains a provision for the SUT to indicate that it was unable to respond to the query, and why. The options for the reason the SUT could not respond, as defined in the ICD, are:

- Unknown Predicate: could not respond because the query used predicates the SUT doesn't understand. The list of offending predicates may be given in an optional "unknown predicates" string.
- Cannot Identify Single Frame: could not respond because the SUT only works on single frames and it could not figure out which frame to use.
- UnsupportedDataType: could not respond because the query involves a data type (e.g. mobile cameras) which the SUT does not support.
- SoftwareError: could not respond due to an unexpected software error (e.g. SUT code throws an exception).
- Other: could not respond for some other reason. Details may be given in an optional comment string.

Post-evaluation, UCLA provided a list of unsupported predicates and predicate combinations by their Phase 3 SUT:

- **color**: color predicates are only supported with respect to the following object types: automobile, top-wear, bottom-wear, table, chair. Note that the Phase 3 SOCs' evaluation queries did not use color predicates.
- **action/behavior**: the following predicates are not supported "due to the lack of reliable visual clues or training samples, or high ambiguity in definition":
  - talking - [the UCLA SUT] "needs subtle cues, e.g. gesture, facial motion to recognize who is really talking at a certain time instance."
  - touching - [the UCLA SUT] "needs accurate 3D hand position which is not available"
  - catching – "Disc and balls are too small and move too fast to detect"
  - swinging – [the UCLA SUT] "needs accurate 3D arm motion, also [the predicate is] ambiguous."
  - occluding – "Ambiguous definition and [the UCLA SUT] needs reliable 3D information."
  - donning, doffing – due to unreliable results, these predicates were not supported.
- **indoor scenes**: the following predicates are not supported in indoor scenes because "the objects involved are too small in loading/unloading; the 3D projected human positions are not accurate in indoor; or the predicates are not well-defined in indoor": driving, crossing, mounting, dismounting, loading, unloading, same_motion, opposite_motion, following, passing, turning, turning_left, turning_right, u-turn.
- **open/closed attributes**: the "open" and "closed" attributes are only supported for the following objects: vehicle – door, vehicle – hood, and vehicle – trunk.
- **predicate combinations**: the following predicates were supported only when applied to the "person" or "vehicle" object types (or sub-classes): clear_line_of_sight, driving, together, loading, unloading, passing, following, same_motion, opposite_motion, closer, further, mounting, dismounting.

There were 107 total SUT "unable to respond" responses reported by the UCLA SUT during the Phase 3 evaluation, all of which utilized the "Other" option. UCLA's SUT provided an explanation for the non-response in the comment string. In 105 of the 107 cases, the comment indicates that the SUT was unable to respond because the combination of predicate and argument used in the query was not supported by the SUT. Table 15 details the predicate-argument combinations the SUT did not support.

Table 15: Reasons for Unable to Answer Responses with an "Other" code

predicate-argument combinations	SIG Parking Lot 2014-01-04	SIG Parking Lot 2014-10-18	Pratt Garden 2014-09-20	Schiciano Auditorium 2014-02-22	SIG Office 2013-09-04	Totals
passing INDOOR				2	2	**4**
opposite-motion INDOOR				1	1	**2**
same-motion INDOOR				2	2	**4**
following INDOOR				2	2	**4**
turning INDOOR					2	**2**
turning-right INDOOR					1	**1**
turning-left INDOOR					1	**1**
u-turn INDOOR					1	**1**
pointing			1	7	1	**9**
clear-line-of-sight(person,person) clear-line-of-sight(person,package) clear-line-of-sight(wheel)	1				4	**5**
together(person,small-object) together(person,package)					2	**2**
taking-down(person,small-object)					1	**1**
loading(person,small-object) loading(person,package) loading(person,trunk) loading(person,luggage)		2	2		2	**6**
unloading(person,package) unloading(person,luggage) unloading(person,small-object) unloading(person,trunk)	1	3			1	**5**
putting-up(person,small-object)					1	**1**
crawling					2	**2**
dropping	5	2	2	2	1	**12**
talking	2	2	3	3		**10**
touching	2	1	1			**4**
catching	2			1	1	**4**
swinging					1	**1**
occluding	2		1			**3**
donning	1	2	3*			**6**
doffing	2	1	4*			**7**
facing(automobile)		1				**1**
mounting(person,two-wheeled-vehicle)			3			**3**
dismounting(person,two-wheeled-vehicle)		1	1			**2**
exiting(person,room)				1		**1**
entering(person,room)				1		**1**
TOTALS	**18**	**15**	**21**	**22**	**29**	**105**

During post-evaluation analysis, we discovered errors in the formal language specification of two queries (denoted by a * in Table 15 above):

- Query ID 640 - doffing(topwear,obj-person3) – the order of operands is wrong
- Query ID 641 - donning(hat,female) – the order of operands is wrong

These queries do not follow the Formal Language Specification for the donning and doffing predicates.

Two queries were answered by the SUT with "Unable_to_respond"-"Other" responses in a manner different from the responses reported in the above table:

- Query ID 200: the UCLA SUT responded with a "Service Error". For reference, the natural language version of query ID 200 is: "Is person3 closer to person1 than to person2?"
- Query ID 1025: the UCLA SUT claimed the "identifier is not defined" in the comments and references the "pointing" predicate. The natural language version of query ID 1025 is "Is obj-student3 pointing?". Curiously, the defined object (obj-student-3) and temporal window were defined and understood by the SUT in previously answered queries.

The following objects were not detected by the SUT leading to a number of skipped queries that depended on these objects:

Table 16: Queries skipped due to unidentified objects.

object	SOC	storyline	# skipped queries
obj-jeep	SIG Parking Lot 2014-01-04	storyline-Tracking-Automobiles	8
obj-person3	SIG Parking Lot 2014-01-04	storyline-Person-Attributes	8
obj-person4	SIG Parking Lot 2014-01-04	storyline-Person-Attributes	7
obj-jeep	SIG Parking Lot 2014-01-04	storyline-Vehicle-Attributes	4
obj-door2	SIG Parking Lot 2014-01-04	storyline-Vehicle-Attributes	4
obj-pitcher	SIG Parking Lot 2014-01-04	storyline-Geometry	5
obj-head	SIG Parking Lot 2014-01-04	storyline-People-Parts	3
obj-hand	SIG Parking Lot 2014-01-04	storyline-People-Parts	3
obj-arm1	SIG Parking Lot 2014-01-04	storyline-People-Parts	3
abj-arm2	SIG Parking Lot 2014-01-04	storyline-People-Parts	3
obj-person1	SIG Parking Lot 2014-01-04	storyline-People-Car-Interactions-1	5
obj-jeep	SIG Parking Lot 2014-01-04	storyline-People-Car-Interactions-1	9
obj-jeep	SIG Parking Lot 2014-01-04	storyline-People-Car-Interactions-2	10
obj-person1	SIG Parking Lot 2014-01-04	storyline-People-Car-Interactions-2	5
obj-person2	SIG Parking Lot 2014-01-04	storyline-People-Car-Interactions-2	4
obj-person2	SIG Parking Lot 2014-01-04	storyline-People-Object-Interactions-1-Dodgeball	10

object	SOC	storyline	# skipped queries
obj-person2	SIG Parking Lot 2014-01-04	storyline-People-Object-Interactions-2-Kickball	7
obj-arm1	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	2
obj-arm2	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	2
obj-lowerbody2	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	2
obj-hood1	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	2
obj-wheel1	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	3
obj-wheel4	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	4
obj-hood4	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	3
obj-person3	SIG Parking Lot 2014-10-18	storyline-attributes	8
obj-person4	SIG Parking Lot 2014-10-18	storyline-attributes	8
obj-door1	SIG Parking Lot 2014-10-18	storyline-attributes	5
obj-car3	SIG Parking Lot 2014-10-18	storyline-spatial-relationships	6
obj-car4	SIG Parking Lot 2014-10-18	storyline-spatial-relationships	4
obj-car6	SIG Parking Lot 2014-10-18	storyline-spatial-relationships	3
obj-person3	SIG Parking Lot 2014-10-18	storyline-relationships	3
obj-disc1	SIG Parking Lot 2014-10-18	storyline-relationships	3
obj-disc2	SIG Parking Lot 2014-10-18	storyline-relationships	6
obj-person7	SIG Parking Lot 2014-10-18	storyline-relationships	6
obj-person2	Pratt Garden 2014-09-20	storyline-fashion-show	5
obj-arm2	Pratt Garden 2014-09-20	storyline-fashion-show	3
obj-ball1	Pratt Garden 2014-09-20	storyline-sports	6
obj-arm1	Pratt Garden 2014-09-20	storyline-sports	2
obj-hand2	Schiciano Auditorium 2014-02-22	storyline-part-of-relationships	3
obj-backpack	Schiciano Auditorium 2014-02-22	storyline-registration	8
obj-bag	Schiciano Auditorium 2014-02-22	storyline-presentation	8
obj-hat	Schiciano Auditorium 2014-02-22	storyline-presentation2	10
obj-presenter1	Schiciano Auditorium 2014-02-22	storyline-presentation2	10
obj-door2	Schiciano Auditorium 2014-02-22	storyline-panic	2
obj-bag	Schiciano Auditorium 2014-02-22	storyline-panic	6

Note that some skipped queries in the above table were skipped for multiple undetected objects, so the overall query total in the above table is higher than the actual skipped query total.

6.7 Response Time

Two metrics related to response time were evaluated:

- SOC startup time, measured as the time between when the EES sends the SOC description to the SUT and when the EES receives the request for the first query from the SUT. This is the time required by the system to prepare to answer queries in the SOC. Since the pre-processing of video data was done before the evaluation process with the EES started, SOC startup times are expected to be relatively low.
- Query response time, measured as the time between when the EES sends the query to the SUT and when the EES receives the answer to the query from the SUT.

All times include a certain amount of overhead due to time required to transmit data over the Internet. This overhead cannot be measured directly and will vary depending on a variety of factors. However, the overhead should be a small percentage of the total time.

Table 17 shows the SOC startup times and mean/min/max response times for UCLA's SUT. We know that UCLA paused processing for several days between the first and second SOCs. We believe that processing was also suspended before the start of the third through fifth SOCs. Therefore, the SOC startup times (aside from the first) are not informative.

Table 17: Response time data for UCLA. All times are reported in seconds

SOC	Startup Time	Mean Query Response Time	Min. Query Response Time	Max. Query Response Time
SIG Office 2013-09-04	5.619	26.911	1.941	206.629
SIG Parking Lot 2014-01-04	272130.567	27.437	7.104	204.483
SIG Parking Lot 2014-10-18	31603.248	29.629	0.375	622.537
Pratt Garden 2014-09-20	26752.901	84.098	2.973	2965.791
Schiciano Auditorium 2014-02-22	42324.831	12.825	3.304	201.038
Entire Session	N/A	39.381	0.375	2965.791

Figure 14 shows a histogram of UCLA's query response times (excluding response times above 1000 seconds for better scaling). The minimum response time was 0.375 seconds.

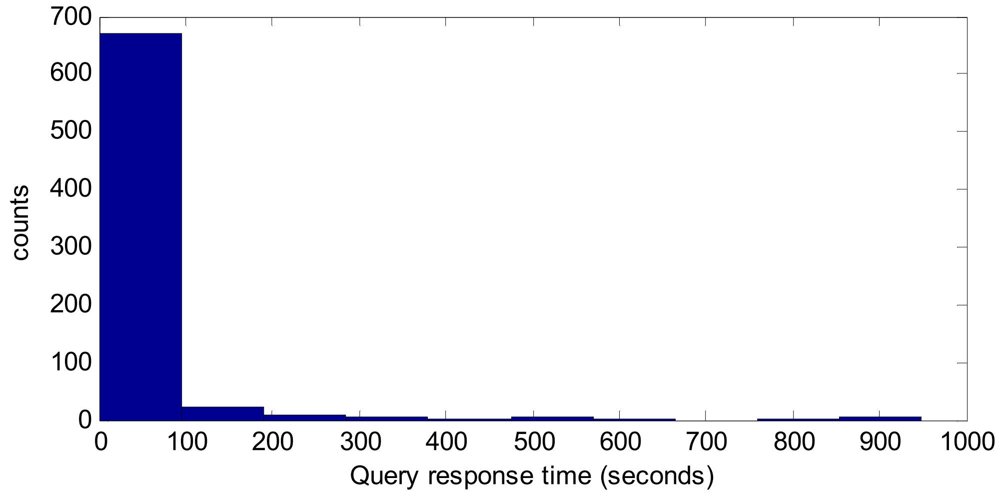

Figure 14: Histogram of UCLA query response times (excluding times greater than 1000 seconds)

Figure 15 shows a plot of the Query response times for each query. A total of 18 queries had a response time greater than 1000 seconds.

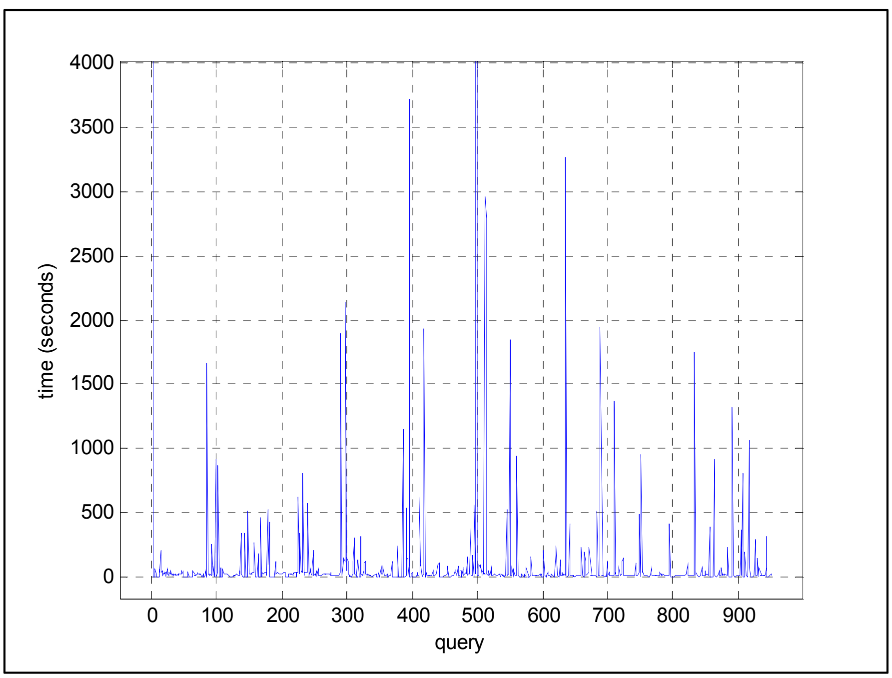

Figure 15: SUT Response Time (seconds) by Phase3 query number.

Table 18 provides further details on the 18 queries with response times greater than 1000 seconds. Of the 18 queries, 13 are object definition queries. Query ID 109 is the first query of the "SIG Parking Lot 2014-01-04" SOC and the response time is attributable to the pause between SOCs taken by UCLA to address SUT issues and the SOC startup time. The other 5 queries use one of the following predicates: clear-line-of-sight, same-motion, following, entering, and pointing.

Table 18: Queries with response times greater than 1000 seconds.

ID	SOC	Storyline	Name	Description	Query Response Time
20	SIG Office 2013-09-04	storyline-additional-reception	query-p-so-relationships-person-clear-line-of-sight-small-object	Is the person clear-line-of-sight the small-object?	1431.8
109	SIG Parking Lot 2014-01-04	storyline-Tracking-Automobiles	query-1	Is obj-jeep detected?	81966.6
193	SIG Parking Lot 2014-01-04	storyline-Geometry	query-1	Is person1 detected?	1667.9
398	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	query-1	Is there a person <person1> at pixel(1078,410) in the FOV of sensor GL4?	1900.4
406	SIG Parking Lot 2014-10-18	storyline-part-of-relationships	query-9	Is there an arm <arm3> at pixel (1444,370) in the FOV of sensor GL4?	2143.2
493	SIG Parking Lot 2014-10-18	storyline-attributes	query-33	Is there a door <door1> at pixel(234,541) in the FOV of GL1?	1147.5
504	SIG Parking Lot 2014-10-18	storyline-spatial-relationships	query-7	Is there an automobile <car7> at pixel (1671,504) in the FOV of sensor RT1?	3720.7
526	SIG Parking Lot 2014-10-18	storyline-relationships	query-12	Is there a tool <tool1> at pixel (337,378) in the FOV of sensor GL1?	1931.6
607	Pratt Garden 2014-09-20	storyline-exercise-class	query-16	Is there a head <head1> in the FOV of HC2 at pixel (634,363)?	8900.7
620	Pratt Garden 2014-09-20	storyline-exercise-class	query-29	Are there two-people moving in the same direction (same-motion)?	2965.8
621	Pratt Garden 2014-09-20	storyline-exercise-class	query-30	Is there a person following another person?	2788.7
658	Pratt Garden 2014-09-20	storyline-exercise-class	query-67	Is there a lower-body <lb1> in the FOV of HC3 at pixel (150,532)?	1844.8
744	Pratt Garden 2014-09-20	storyline-sports	query-10	Is there a ball <ball1> in the FOV of HC2 at pixel (933,414)?	3260.6

ID	SOC	Storyline	Name	Description	Query Response Time
797	Pratt Garden 2014-09-20	storyline-sports	query-63	Is there a head <head2> in the FOV of IP5 at pixel (288,250)?	1947.9
818	Schiciano Auditorium 2014-02-22	storyline-part-of-relationships	query-4	Identify obj-head	1367.6
942	Schiciano Auditorium 2014-02-22	storyline-presentation	query-2	Are there at least 5 people who enter the auditorium during time-enter?	1749.1
1000	Schiciano Auditorium 2014-02-22	storyline-presentation2	query-16	Identify person as obj-student2.	1325.2
1025	Schiciano Auditorium 2014-02-22	storyline-presentation2	query-41	Is obj-student3 pointing?	1068.7

6.8 Utilization of Assessor Responses

Performers were given the option in Phase 3 to use assessor responses to adapt their SUT between queries. UCLA has stated that their system did not use the assessor responses to adapt their SUT.

7 Conclusions

The Phase 3 evaluation included five SOCs: four new SOCs developed for Phase 3, and one SOC – the SIG Office SOC – that was reused from Phase 2.2. For the SIG Office SOC, new queries were developed for use in the Phase 3 evaluation.

Based on lessons learned during the Phase 2.2 evaluation, a number of changes were made to the strategy used to develop queries for Phase 3. Phase 3 queries were greatly simplified, with most using only one to three predicates. Although queries were simpler, there were far more queries in Phase 3 (1,060 compared to 276 in Phase 2.2). The set of predicates used was pruned back to those judged most important. In addition, adding object definitions allowed query developers to more precisely specify the objects involved in the queries.

Considering the Phase 3 data only, 709 queries (not including object definition queries) were available. The UCLA SUT responded to 459 of these queries (65%). Of the queries the SUT responded to, 289 responses were correct (63%). An additional 243 object definition queries were presented; the SUT was able to identify 197 of the objects (81%).

Considering the Phase 2.2 SIG Office SOC only, 108 queries were available. The SUT responded to 79 of these queries (73%). Of the queries the SUT responded to, 62 responses were correct (78%). There were no object definition queries for this SOC.

Considering all SOCs (both Phase 3 and Phase 2 together), 817 queries were available (not including object definition queries). The SUT responded to 538 of these queries (66%). Of the queries the SUT responded to, 351 responses were correct (65%). An additional 243 object definition queries were presented; the SUT was able to identify 197 of the objects (81%).

Query response accuracy degrades as the number of query predicates increases (increasing complexity): 28.3% error rate for a 1-predicate query to 58.5% error rate for a query with 5 or more predicates.

The UCLA SUT was most accurate when answering queries in the “object definition”, “classification”, “part of”, and “attributes” categories (error rates less than 30% at a declaration confidence >= 0). The UCLA SUT performed relatively poorly when answering queries in the “spatial” (48.3% error rate at a declaration confidence >= 0) and “relationships” (43.5% error rate at a declaration confidence >= 0) query categories.

Appendix A: SUT Hardware and Software Configuration

The UCLA performer team provided the following information about their Phase 3 SUT hardware and software configuration.

A.1 Hardware

UCLA's SUT ran on the following set of computers:

- 1 deployment machine:
  - Processor: Intel i7-3770, x64, 4 cores, 3.40GHz
  - RAM: 32 GB
  - Hard Disk: 3TB (for both software and data storage)
- 12 cluster nodes:
  - Processor: Intel i7-3820, x64, 4 cores, 3.60GHz
  - RAM: 32 GB
  - Hard Disk: 2TB
- 1 query engine node:
  - Processor: Intel i7-3770, x64, 4 cores, 3.40GHz
  - RAM: 16 GB
  - Hard Disk: 1TB

Network connection between all machines was 1 gbps Ethernet.

A.2 Software

Operating systems:

- Windows 7 (x64) on the query engine node
- Ubuntu 14.04 (x64) on all other machines

The following software packages are needed for building the code from source and deploying the system.

- Python 2.7
- Thrift 0.9.2
- Django 1.7
- Python Pillow 2.3
- numpy
- Boost 1.55.0
- Cmake 2.8
- G++ 4.8
- OpenMP
- OpenCV 2.4.10
- Libfftw-dev
- Libeigen3-dev
- Libmatio 1.5.2
- MPICH2
- FFMPEG
- Java 7
- Apache Jena
- Eclipse IDE 3.8
- MATLAB 2014b with Parallel Computing Toolbox

Appendix B: Phase 3 Testing Queries and Answers

The following is a complete list of all 1025 queries used in the Phase 3 evaluation. Queries are divided into sections by SOC and storyline. The two columns under "Assessor" give, respectively, the human assessor's (i.e. correct) answer to the query and the human assessor's confidence in his or her answer (as 'H' for high, 'M' for medium, or 'L' for low).

Under the headings for "ucla" performer, the three columns are, respectively:

- The SUT's answer
- The SUT's confidence
- The SUT's query response time in seconds

SUT answers that match the assessor's response are shaded green; those that do not match are shaded red. If an SUT did not respond to a query, the two SUT answer and confidence columns are instead used to display the reason for the non-response. "UnknownObject" means that the SUT responded that it was not able to identify an object used in the query. "Skipped" means that the EES did not present the query because it depended on an object that the SUT had previously indicated it could not identify. In all other cases where UCLA's SUT was unable to respond to a query, UCLA sent the "Other" response code.

soc-sig-office-2013-09-04-testing							
storyline-additional-reception							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Do two people enter the reception?	relationships	T	H	F	0.79	66.55
query-2	Do two people enter the reception?	relationships	F	H	F	0.79	74.37
query-3	Do two people exit the reception?	relationships	F	H	T	0.38	2.26
query-4	Do two people enter the reception?	relationships	F	H	T	0.38	1.94
query-5	Do two people exit the reception?	relationships	F	H	F	0.79	24.41
query-relationships-person-facing-opposite-person	Is the person facing-opposite the person?	spatial	F	H	F	0.79	100.86
query-relationships-person-passing-person	Is the person passing the person?	tracking	T	H	Other		96.24
query-relationships-person-opposite-motion-person	Is the person opposite-motion the person?	tracking	F	H	Other		83.26
query-relationships-	Is the person following	tracking	T	H	Other		20.42

person-following-person	the person?						
query-relationships-person-touching-person	Is the person touching the person?	relationships	F	H	T	0.95	16.75
query-relationships-person-facing-person	Is the person facing the person?	spatial	F	M	F	0.79	20.23
query-relationships-person-clear-line-of-sight-person	Is the person clear-line-of-sight the person?	spatial	T	H	Other		15.73
query-relationships-person-together-person	Is the person together the person?	relationships	T	M	T	0.95	206.63
query-relationships-person-same-motion-person	Is the person same-motion the person?	tracking	T	H	Other		197.50
query-6	Is there a person carrying luggage in the reception?	relationships	F	H	T	0.95	52.38
query-p-so-relationships-person-swinging-small-object	Is the person swinging the small-object?	relationships	T	H	F	0.79	26.10
query-p-so-relationships-person-catching-small-object	Is the person catching the small-object?	relationships	F	H	F	0.79	20.93
query-p-so-relationships-person-dropping-small-object	Is the person dropping the small-object?	relationships	F	H	F	0.79	13.42
query-p-so-relationships-person-touching-small-object	Is the person touching the small-object?	relationships	T	H	T	0.95	36.04
query-p-so-relationships-person-clear-line-of-sight-small-object	Is the person clear-line-of-sight the small-object?	spatial	T	H	Other		1431.82
query-p-so-relationships-person-together-small-object	Is the person together the small-object?	relationships	T	H	Other		105.50
query-p-so-relationships-person-picking-up-small-object	Is the person picking-up the small-object?	relationships	T	H	T	0.95	17.19
query-p-so-relationships-person-carrying-small-object	Is the person carrying the small-object?	relationships	T	H	T	0.95	17.84
query-p-so-relationships-person-throwing-small-object	Is the person throwing the small-object?	relationships	F	H	F	0.79	79.93
query-p-so-relationships-person-taking-down-small-object	Is the person taking-down the small-object?	relationships	F	H	Other		203.27

query-p-so-relationships-person-putting-down-small-object	Is the person putting-down the small-object?	relationships	T	H	T	0.95	58.93
query-p-so-relationships-person-loading-small-object	Is the person loading the small-object?	relationships	F	H	Other		46.94
query-p-so-relationships-person-putting-up-small-object	Is the person putting-up the small-object?	relationships	F	H	Other		29.51
query-7	Do more than 5 people enter the room?	relationships	T	H	T	0.38	24.88
query-8	Do less than 6 people enter the room?	relationships	F	H	T	0.95	46.35
storyline-additional-package-exchange							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there at least one person in the AOR?	classification	T	H	T	0.95	68.87
query-2	Is there at least one person in the reception room?	classification	T	H	T	0.95	30.55
query-3	Is there at least one person in the breakroom?	classification	T	H	T	0.95	15.35
query-4	Is there at least one person in the conference room?	classification	T	H	T	0.49	9.26
query-5	Is there at least one female in the long hallway?	classification	F	H	F	0.97	12.82
query-6	Is there at least one person in the long hallway?	classification	F	H	F	0.97	11.16
query-7	Is there at least one female in the reception room?	classification	T	H	T	0.95	22.94
query-8	Is there at least one female in the breakroom?	classification	T	H	F	0.98	22.98
query-9	Is there at least one female in the conference room?	classification	F	H	F	0.95	19.88
query-confroom-person-stationary	Is the person stationary?	tracking	T	H	T	0.49	29.27
query-confroom-person-	Is the person reading?	attributes	T	H	T	0.49	19.00

reading							
query-confroom-person-eating	Is the person eating?	attributes	F	H	T	0.49	18.48
query-confroom-person-stopping	Is the person stopping?	tracking	T	H	T	0.33	13.35
query-confroom-person-crawling	Is the person crawling?	attributes	F	H	Other		22.09
query-confroom-person-writing	Is the person writing?	attributes	F	H	T	0.49	16.45
query-confroom-person-moving	Is the person moving?	tracking	T	H	T	0.33	16.97
query-confroom-person-running	Is the person running?	attributes	F	H	F	0.95	16.56
query-confroom-person-talking	Is the person talking?	attributes	F	H	T	0.49	10.01
query-confroom-person-turning	Is the person turning?	tracking	T	H	Other		17.14
query-confroom-person-turning-right	Is the person turning-right?	tracking	T	H	Other		9.65
query-confroom-person-turning-left	Is the person turning-left?	tracking	T	H	Other		9.11
query-confroom-person-u-turn	Is the person u-turn?	tracking	T	M	Other		12.92
query-confroom-person-starting	Is the person starting?	tracking	T	H	T	0.33	9.49
query-confroom-person-sitting	Is the person sitting?	attributes	T	H	T	0.49	15.80
query-confroom-person-pointing	Is the person pointing?	attributes	F	H	F	0.95	10.47
query-confroom-person-walking	Is the person walking?	attributes	T	H	T	0.49	9.18
query-confroom-person-standing	Is the person standing?	attributes	T	H	T	0.49	9.08
query-10	Does the same person enter the breakroom and the conference room?	relationships	T	H	F	0.95	22.51
query-11	Does the same female enter the breakroom and the conference room?	relationships	F	H	F	0.95	17.59
query-12	Is there a package in the breakroom?	classification	T	H	T	0.95	11.54
query-relationships-person-touching-package	Is the person touching the package?	relationships	T	H	T	0.95	15.89

query-relationships-person-loading-package	Is the person loading the package?	relationships	F	H	Other		27.42
query-relationships-person-unloading-package	Is the person unloading the package?	relationships	F	H	Other		12.74
query-relationships-person-swinging-package	Is the person swinging the package?	relationships	F	H	Other		7.64
query-relationships-person-catching-package	Is the person catching the package?	relationships	F	H	Other		6.43
query-relationships-person-dropping-package	Is the person dropping the package?	relationships	F	H	Other		9.75
query-relationships-person-facing-package	Is the person facing the package?	spatial	T	H	T	0.95	13.75
query-relationships-person-clear-line-of-sight-package	Is the person clear-line-of-sight the package?	spatial	T	H	Other		8.43
query-relationships-person-together-package	Is the person together the package?	relationships	T	H	Other		9.99
query-relationships-person-picking-up-package	Is the person picking-up the package?	relationships	T	H	T	0.95	13.74
query-relationships-person-carrying-package	Is the person carrying the package?	relationships	T	H	T	0.95	18.30
query-relationships-person-throwing-package	Is the person throwing the package?	relationships	F	H	F	1.00	9.63
query-relationships-person-putting-down-package	Is the person putting-down the package?	relationships	T	H	T	0.40	19.30
query-relationships-person-on-package	Is the person on the package?	relationships	F	H	T	0.95	14.44
query-13	Does a person touch a package in the conference room?	relationships	F	H	F	0.95	20.53
query-14	Does a person touch a table in the conference room?	relationships	T	H	T	0.49	10.75
query-15	Does a person's foot touch a table in the conference room?		T	H	F	0.95	13.78
storyline-additional-bag-switch							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there at least one person in the AOR?	classification	T	H	T	0.95	118.75
query-2	Is there at least one	classification	T	H	T	0.95	13.27

	person in the reception room?						
query-3	Is there at least one person in the breakroom?	classification	T	H	T	0.95	7.67
query-4	Is there at least one person in the conference room?	classification	F	H	F	0.95	14.16
query-5	Is there at least one person in the long hallway?	classification	T	H	T	0.49	14.44
query-6	Do two different people enter the breakroom?	relationships	T	H	T	0.40	36.61
query-7	Do two people enter the breakroom together?	relationships	T	H	T	0.40	115.86
query-interpersonal-person-same-motion-person	Is the person same-motion the person?	tracking	T	H	Other		15.48
query-interpersonal-person-passing-person	Is the person passing the person?	tracking	T	H	Other		27.64
query-interpersonal-person-following-person	Is the person following the person?	tracking	T	H	Other		21.69
query-interpersonal-person-touching-person	Is the person touching the person?	relationships	T	H	T	0.95	26.43
query-interpersonal-person-facing-person	Is the person facing the person?	spatial	T	H	T	0.95	11.68
query-interpersonal-person-clear-line-of-sight-person	Is the person clear-line-of-sight the person?	spatial	T	H	Other		56.44
query-interpersonal-person-together-person	Is the person together the person?	relationships	T	H	T	0.95	76.06
query-interpersonal-person-carrying-person	Is the person carrying the person?	relationships	F	H	T	0.95	15.74
query-unary-person-stationary	Is the person stationary?	tracking	T	H	T	0.95	16.68
query-unary-person-reading	Is the person reading?	attributes	T	H	T	0.95	12.23
query-unary-person-stopping	Is the person stopping?	tracking	T	H	T	0.95	11.52
query-unary-person-crawling	Is the person crawling?	attributes	F	H	Other		17.00
query-unary-person-moving	Is the person moving?	tracking	T	H	T	0.95	6.08

query-unary-person-running	Is the person running?	attributes	F	H	F	0.98	15.09
query-unary-person-talking	Is the person talking?	attributes	T	H	T	0.95	13.44
query-unary-person-writing	Is the person writing?	attributes	F	H	T	0.95	10.97
query-unary-person-starting	Is the person starting?	tracking	T	H	T	0.95	11.80
query-unary-person-sitting	Is the person sitting?	attributes	T	H	T	0.95	7.37
query-unary-person-pointing	Is the person pointing?	attributes	T	H	Other		12.02
query-unary-person-turning	Is the person turning?	tracking	T	H	Other		10.00
query-unary-person-walking	Is the person walking?	attributes	T	H	T	0.95	13.74
query-unary-person-standing	Is the person standing?	attributes	T	H	T	0.95	9.38
query-8	Does the same person enter the reception room and exit the hallway?	relationships	T	H	F	0.97	14.56
query-9	Do more than 2 people enter the reception room and exit the hallway?	relationships	F	H	F	0.97	14.85
soc-SIGParkingLot-2014-01-04-Testing							
storyline-Tracking-Automobiles							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is obj-jeep detected?	object_definition	T	H	UnknownObject		81966.65
query-jeep1-obj-jeep-moving	Is obj-jeep moving?	tracking	T	H	Skipped		0.01
query-jeep1-obj-jeep-stationary	Is obj-jeep stationary?	tracking	T	H	Skipped		0.01
query-jeep1-obj-jeep-starting	Is obj-jeep starting?	tracking	F	H	Skipped		0.01
query-2	Is obj-suv detected?	object_definition	T	H	T	0.69	65.14
query-suv-obj-suv-stationary	Is obj-suv stationary?	tracking	T	H	F	0.87	25.35
query-suv-obj-suv-moving	Is obj-suv moving?	tracking	F	H	F	0.87	21.11
query-jeep2-obj-jeep-moving	Is obj-jeep moving?	tracking	T	H	Skipped		0.02
query-jeep2-obj-jeep-	Is obj-jeep starting?	tracking	F	H	Skipped		0.03

starting							
query-jeep2-obj-jeep-turning-right	Is obj-jeep turning-right?	tracking	F	H	Skipped		0.01
query-jeep3-person-driving-obj-jeep	Is person driving obj-jeep?	relationships	T	H	Skipped		0.01
storyline-Tracking-People							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.44	54.92
query-person1-obj-person1-u-turn	Is obj-person1 u-turn?	tracking	T	H	F	0.87	30.49
query-person1-obj-person1-turning-left	Is obj-person1 turning-left?	tracking	F	H	F	0.87	204.48
query-2	Is person2 detected?	object_definition	T	H	T	0.44	36.17
query-person2-obj-person2-turning	Is obj-person2 turning?	tracking	T	H	T	0.44	46.24
query-person2-obj-person2-turning-right	Is obj-person2 turning-right?	tracking	F	H	T	0.44	13.32
query-3	Is person3 detected?	object_definition	T	H	T	0.44	38.94
query-person3-obj-person2-opposite-motion-obj-person3	Is obj-person2 opposite-motion obj-person3?	tracking	F	H	F	0.87	27.78
query-person3-obj-person2-following-obj-person3	Is obj-person2 following obj-person3?	tracking	T	H	F	0.87	23.79
query-4	Is person4 detected?	object_definition	T	H	T	0.44	33.64
query-person4-obj-person4-turning	Is obj-person4 turning?	tracking	T	H	F	0.87	13.81
query-person4-obj-person4-turning-left	Is obj-person4 turning-left?	tracking	T	H	F	0.87	57.50
query-person4-obj-person4-stopping	Is obj-person4 stopping?	tracking	T	H	F	0.87	25.30
query-5	Is person5 detected?	object_definition	T	H	T	0.44	19.07
query-person5-obj-person5-starting	Is obj-person5 starting?	tracking	T	H	T	0.44	15.22
query-6	Is person6 detected?	object_definition	T	H	T	0.44	38.19
query-person6-obj-person6-stopping	Is obj-person6 stopping?	tracking	F	H	F	0.91	20.27
query-person6-obj-person6-starting	Is obj-person6 starting?	tracking	T	H	F	0.91	47.87
query-person56-obj-person5-same-motion-obj-person6	Is obj-person5 same-motion obj-person6?	tracking	T	H	T	0.95	14.57

query-person56-obj-person5-following-obj-person6	Is obj-person5 following obj-person6?	tracking	F	L	F	0.91	19.84
storyline-Person-Attributes							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.44	29.25
query-person1-male	Is person1 male?	classification	T	H	T	0.44	9.13
query-person1-obj-person1-crawling	Is obj-person1 crawling?	attributes	F	H	F	0.87	17.88
query-person1-obj-person1-walking	Is obj-person1 walking?	attributes	F	H	F	0.87	16.16
query-person1-obj-person1-running	Is obj-person1 running?	attributes	F	H	F	0.87	20.40
query-person1-obj-person1-talking	Is obj-person1 talking?	attributes	F	H	Other		25.73
query-person1-obj-person1-standing	Is obj-person1 standing?	attributes	T	H	F	0.87	16.17
query-person1-obj-person1-pointing	Is obj-person1 pointing?	attributes	F	H	F	0.87	13.01
query-2	Is person2 detected?	object_definition	T	H	T	0.44	26.68
query-person2-male	Is person2 male?	classification	F	H	F	0.91	13.37
query-person2-obj-person2-crawling	Is obj-person2 crawling?	attributes	F	H	F	0.87	27.30
query-person2-obj-person2-walking	Is obj-person2 walking?	attributes	F	H	F	0.87	28.26
query-person2-obj-person2-running	Is obj-person2 running?	attributes	T	H	F	0.87	29.03
query-person2-obj-person2-talking	Is obj-person2 talking?	attributes	F	H	Other		22.31
query-person2-obj-person2-standing	Is obj-person2 standing?	attributes	T	H	T	0.44	54.31
query-person2-obj-person2-pointing	Is obj-person2 pointing?	attributes	F	H	F	0.87	22.44
query-3	Is person3 detected?	object_definition	T	H	UnknownObject		45.07
query-person3-male	Is person3 male?	classification	T	H	Skipped		0.00
query-person3-obj-person3-crawling	Is obj-person3 crawling?	attributes	F	H	Skipped		0.00
query-person3-obj-person3-walking	Is obj-person3 walking?	attributes	T	H	Skipped		0.00
query-person3-obj-person3-running	Is obj-person3 running?	attributes	T	H	Skipped		0.00
query-person3-obj-	Is obj-person3 talking?	attributes	F	H	Skipped		0.00

person3-talking							
query-person3-obj-person3-standing	Is obj-person3 standing?	attributes	T	H	Skipped		0.00
query-person3-obj-person3-pointing	Is obj-person3 pointing?	attributes	T	H	Skipped		0.01
query-4	Is person4 detected?	object_definition	T	H	UnknownObject		32.75
query-person4-male	Is person4 male?	classification	T	H	Skipped		0.01
query-person4-obj-person4-crawling	Is obj-person4 crawling?	attributes	F	H	Skipped		0.00
query-person4-obj-person4-walking	Is obj-person4 walking?	attributes	F	H	Skipped		0.00
query-person4-obj-person4-running	Is obj-person4 running?	attributes	F	H	Skipped		0.01
query-person4-obj-person4-standing	Is obj-person4 standing?	attributes	T	H	Skipped		0.00
query-person4-obj-person4-pointing	Is obj-person4 pointing?	attributes	T	H	Skipped		0.01
query-5	Is person5 detected?	object_definition	T	H	T	0.35	42.39
query-person5-male	Is person5 male?	classification	T	H	T	0.35	22.99
query-person5-obj-person5-crawling	Is obj-person5 crawling?	attributes	F	H	F	0.87	12.73
query-person5-obj-person5-walking	Is obj-person5 walking?	attributes	F	H	T	0.35	14.22
query-person5-obj-person5-running	Is obj-person5 running?	attributes	F	H	F	0.87	11.18
query-person5-obj-person5-sitting	Is obj-person5 sitting?	attributes	T	H	F	0.87	11.80
query-person5-obj-person5-pointing	Is obj-person5 pointing?	attributes	F	H	F	0.87	13.17
storyline-Vehicle-Attributes							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is obj-suv detected?	object_definition	T	H	T	0.40	35.63
query-2	Is the door detected?	object_definition	T	H	T	0.40	17.84
query-3	Is the door part of the SUV?	part_of	T	H	T	0.40	10.48
query-4	Is the door open?	attributes	F	H	F	0.87	18.79
query-5	Is the jeep detected?	object_definition	T	H	UnknownObject		21.48
query-6	Is door2 detected?	object_definition	T	H	UnknownObject		20.15
query-7	Is door2 part of the SUV?	part_of	F	H	Skipped		0.01
query-8	Is door2 part of the	part_of	T	H	Skipped		0.01

	jeep?						
query-9	Is door2 open?	attributes	T	H	Skipped		0.01
query-10	Is the wheel detected?	object_definition	T	H	T	0.40	24.88
query-11	Is the wheel part of the jeep?	part_of	F	H	Skipped		0.01
query-12	Is the wheel part of the SUV?	part_of	T	H	F	0.91	16.53
query-13	Is there a clear line-of-sight from RT3 to the wheel?	spatial	F	H	Other		15.70
query-14	Is the hood detected?	object_definition	T	H	T	0.95	47.89
query-15	Is the hood part of the jeep?	part_of	T	H	Skipped		0.01
storyline-Geometry							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.35	1667.95
query-2	Is person1 occluding at least one other person from view of camera GL5?	spatial	T	H	Other		19.31
query-3	Is person2 detected?	object_definition	T	H	T	0.35	25.65
query-4	Is person2 occluding at least one other person from view of camera GL5?	spatial	F	H	Other		8.32
query-5	Is there a clear-line-of-sight from person1 to person2?	spatial	T	M	T	0.35	13.49
query-6	Is person3 detected?	object_definition	T	H	T	0.35	22.09
query-7	Is person3 facing-opposite person1?	spatial	T	H	F	0.87	16.89
query-8	Is person3 closer to person1 than to person2?	spatial	F	H	Other		254.29
query-9	Is person4 detected?	object_definition	T	H	T	0.44	43.34
query-10	Is there a clear line-of-sight from RT1 to person4?	spatial	T	H	F	0.87	12.26
query-11	Is the pitcher detected?	object_definition	T	H	UnknownObject		86.92
query-12	Is the pitcher facing person4?	spatial	T	H	Skipped		0.01
query-13	Is there a clear line-of-	spatial	T	H	Skipped		0.01

	sight from the pitcher to person4?						
query-14	Is person5 detected?	object_definition	T	H	T	0.44	916.68
query-15	Is person5 farther from person4 than the pitcher?	spatial	F	H	Skipped		0.00
query-16	Is person5 occluding the pitcher from person4's perspective?	spatial	F	H	Skipped		0.00

storyline-People-Parts

					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.44	31.38
query-2	Is person2 detected?	object_definition	T	H	T	0.44	874.04
query-3	Is the head detected?	object_definition	T	H	UnknownObject		75.96
query-4	Is head part of person2?	part_of	F	H	Skipped		0.01
query-5	Is head part of person1?	part_of	T	H	Skipped		0.01
query-6	Is the hand detected?	object_definition	T	H	UnknownObject		78.43
query-7	Is the hand part of person2?	part_of	F	H	Skipped		0.01
query-8	Is the hand part of person1?	part_of	T	H	Skipped		0.01
query-9	Is person3 detected?	object_definition	T	H	T	0.44	60.55
query-10	Is person4 detected?	object_definition	T	H	T	0.44	21.91
query-11	Is head2 detected?	object_definition	T	H	T	0.44	17.87
query-12	Is head3 detected?	object_definition	T	H	T	0.44	22.35
query-13	Is arm1 detected?	object_definition	T	H	UnknownObject		21.85
query-14	Is arm2 detected?	object_definition	T	H	UnknownObject		18.73
query-15	Is lowerbody1 detected?	object_definition	T	H	T	0.44	22.13
query-16	Is lowerbody2 detected?	object_definition	T	H	T	0.44	28.69
query-17	Is head2 part of person4?	part_of	F	H	F	0.91	15.79
query-18	Is head2 part of person3?	part_of	T	H	T	0.44	15.92
query-19	Is head3 part of person3?	part_of	F	H	F	0.91	16.27
query-20	Is head3 part of person4?	part_of	T	H	T	0.44	8.08

query-21	Is arm1 part of person3?	part_of	F	H	Skipped		0.01
query-22	Is arm1 part of person4?	part_of	T	H	Skipped		0.01
query-23	Is arm2 part of person3?	part_of	T	H	Skipped		0.01
query-24	Is arm2 part of person4?	part_of	F	H	Skipped		0.01
query-25	Is lowerbody1 part of person3?	part_of	F	H	F	0.91	18.19
query-26	Is lowerbody1 part of person4?	part_of	T	H	T	0.44	7.42
query-27	Is lowerbody2 part of person3?	part_of	T	H	T	0.44	15.27
query-28	Is lowerbody2 part of person4?	part_of	F	H	F	0.91	11.24
query-29	Is person5 detected?	object_definition	T	H	T	0.35	20.68
query-30	Is person6 detected?	object_definition	T	H	T	0.35	21.11
query-31	Is head4 detected?	object_definition	T	H	T	0.35	23.40
query-32	Is head5 detected?	object_definition	T	H	T	0.35	21.16
query-33	Is head4 part of person5?	part_of	F	H	T	0.35	16.58
query-34	Is head4 part of person6?	part_of	T	H	T	0.35	15.77
query-35	Is head5 part of person5?	part_of	T	H	T	0.35	44.96
query-36	Is head5 part of person6?	part_of	F	H	T	0.35	20.64
storyline-People-Car-Interactions-1							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	UnknownObject		345.59
query-2	Is jeep detected?	object_definition	T	H	UnknownObject		26.77
query-3	Is person1 inside the jeep?	relationships	F	H	Skipped		0.01
query-4	Is person1 together with the jeep?	relationships	T	H	Skipped		0.01
query-5	Is person1 dismounting the jeep?	relationships	T	H	Skipped		0.01
query-6	Is person1 driving the jeep?	relationships	F	H	Skipped		0.01
query-7	Is person2 detected?	object_definition	T	H	T	0.35	340.98

query-8	Is person2 inside the jeep?	relationships	F	H	Skipped		0.01
query-9	Is person2 together with the jeep?	relationships	T	H	Skipped		0.01
query-10	Is person2 dismounting the jeep?	relationships	T	H	Skipped		0.01
query-11	Is person2 driving the jeep?	relationships	T	H	Skipped		0.03
query-12	Is person3 detected?	object_definition	T	H	T	0.40	508.88
query-13	Is the SUV detected?	object_definition	T	H	T	0.40	42.70
query-14	Is person3 together with the suv?	relationships	T	H	T	0.40	24.91
query-15	Is person3 touching the SUV?	relationships	T	H	Other		20.94
query-16	Is person3 mounting the SUV?	relationships	F	H	F	0.87	39.76
query-17	Is person3 outside the SUV?	relationships	T	H	T	0.40	28.32
query-18	Is person3 unloading the SUV?	relationships	T	H	F	0.87	21.78
query-19	Is person3 loading the suv?	relationships	F	H	F	0.87	35.24
query-20	Is person3 driving the SUV?	relationships	F	H	F	0.87	36.01
storyline-People-Car-Interactions-2							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	UnknownObject		267.88
query-2	Is the jeep detected?	object_definition	T	H	UnknownObject		28.72
query-3	Is person1 inside the jeep?	relationships	T	H	Skipped		0.01
query-4	Is person1 dismounting the jeep?	relationships	F	H	Skipped		0.01
query-5	Is person1 mounting the jeep?	relationships	T	H	Skipped		0.01
query-6	Is person1 driving the jeep?	relationships	F	H	Skipped		0.01
query-7	Is person2 detected?	object_definition	T	H	UnknownObject		185.06
query-8	Is person2 outside the jeep?	relationships	T	H	Skipped		0.01
query-9	Is person2 dismounting the jeep?	relationships	T	H	Skipped		0.01
query-10	Is person2 driving the	relationships	T	H	Skipped		0.01

	jeep?						
query-11	Is person3 detected?	object_definition	T	H	T	0.35	460.17
query-12	Is person3 inside the jeep?	relationships	F	H	Skipped		0.00
query-13	Is person3 mounting the jeep?	relationships	F	H	Skipped		0.01
storyline-People-Object-Interactions-1-Dodgeball							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.40	31.13
query-A-obj-person1-touching-ball	Is obj-person1 touching ball?	relationships	T	H	Other		19.86
query-A-obj-person1-putting-down-ball	Is obj-person1 putting-down ball?	relationships	T	H	F	0.87	27.82
query-A-obj-person1-picking-up-ball	Is obj-person1 picking-up ball?	relationships	F	H	T	0.38	27.41
query-A-obj-person1-throwing-ball	Is obj-person1 throwing ball?	relationships	F	H	F	0.87	40.37
query-A-obj-person1-catching-ball	Is obj-person1 catching ball?	relationships	T	H	Other		34.96
query-A-obj-person1-dropping-ball	Is obj-person1 dropping ball?	relationships	F	H	Other		16.67
query-2	Is person1's foot on a ball?		T	H	F	0.87	13.30
query-B-obj-person1-donning-clothing	Is obj-person1 donning clothing?	relationships	F	H	Other		530.49
query-B-obj-person1-doffing-clothing	Is obj-person1 doffing clothing?	relationships	F	H	Other		91.45
query-3	Is person2 detected?	object_definition	T	H	UnknownObject		429.23
query-C-obj-person2-touching-ball	Is obj-person2 touching ball?	relationships	T	H	Skipped		0.01
query-C-obj-person2-putting-down-ball	Is obj-person2 putting-down ball?	relationships	F	H	Skipped		0.01
query-C-obj-person2-picking-up-ball	Is obj-person2 picking-up ball?	relationships	F	H	Skipped		0.01
query-C-obj-person2-throwing-ball	Is obj-person2 throwing ball?	relationships	T	H	Skipped		0.00
query-C-obj-person2-catching-ball	Is obj-person2 catching ball?	relationships	T	H	Skipped		0.01
query-C-obj-person2-dropping-ball	Is obj-person2 dropping ball?	relationships	F	H	Skipped		0.01
query-4	Is person2's foot on a ball?		F	H	Skipped		0.01

query-D-obj-person2-donning-clothing	Is obj-person2 donning clothing?	relationships	F	H	Skipped		0.00
query-D-obj-person2-doffing-clothing	Is obj-person2 doffing clothing?	relationships	F	H	Skipped		0.00
query-5	Is person3 detected?	object_definition	T	H	T	0.44	121.81
query-E-obj-person3-throwing-ball	Is obj-person3 throwing ball?	relationships	F	H	F	0.87	23.70
query-E-obj-person3-dropping-ball	Is obj-person3 dropping ball?	relationships	F	H	Other		17.93
query-E-obj-person3-catching-ball	Is obj-person3 catching ball?	relationships	F	H	Other		19.43
query-E-obj-person3-putting-down-ball	Is obj-person3 putting-down ball?	relationships	F	H	F	0.87	31.45
query-F-obj-person3-carrying-luggage	Is obj-person3 carrying luggage?	relationships	T	H	F	0.87	25.14
query-F-obj-person3-dropping-luggage	Is obj-person3 dropping luggage?	relationships	T	H	Other		20.35
query-F-obj-person3-unloading-luggage	Is obj-person3 unloading luggage?	relationships	F	H	Other		19.33
storyline-People-Object-Interactions-2-Kickball							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.44	22.23
query-2	Is person1 female?	classification	F	H	F	0.91	28.92
query-3	Is person1 carrying luggage?	relationships	F	H	F	0.91	26.70
query-4	Is person1 carrying a ball?	relationships	T	H	F	0.91	28.16
query-5	Is person1 picking up luggage?	relationships	F	H	T	0.44	24.91
query-6	Is person1 putting down a ball?	relationships	T	H	F	0.91	31.03
query-7	Is person2 detected?	object_definition	T	H	UnknownObject		37.53
query-8	Is person2 female?	classification	F	H	Skipped		0.01
query-9	Is person2 carrying luggage?	relationships	T	H	Skipped		0.01
query-10	Is person2 carrying a ball?	relationships	F	H	Skipped		0.01
query-11	Is person2 picking up luggage?	relationships	T	H	Skipped		0.01
query-12	Is person2 putting down a ball?	relationships	F	H	Skipped		0.01
query-13	Are person1 and	relationships	F	H	Skipped		0.01

	person2 together?						
query-14	Is person3 detected?	object_definition	T	H	T	0.44	30.04
query-15	Is person3 female?	classification	T	H	T	0.44	7.10
query-15b	Is person3 on the ground?	relationships	T	H	T	0.44	22.40
query-16	Is person3 together with at least one other person?	relationships	F	H	T	0.44	20.03
query-17	Is person3 doffing an article of clothing?	relationships	T	H	Other		24.10
query-18	Is person3 dropping clothing?	relationships	T	H	Other		10.46
query-19	Is person4 detected?	object_definition	T	H	T	0.44	39.07
query-20	Is person4 female?	classification	F	H	F	0.91	29.15
query-21	Does person4 pick up a ball?	relationships	T	H	F	0.91	36.87
query-22	Does person4 carry a ball?	relationships	T	H	F	0.91	30.69
query-23	Does person4 throw a ball?	relationships	T	H	F	0.91	31.63
query-24	Does person4 drop a ball?	relationships	F	H	Other		3.88
storyline-People-Object-Interactions-3							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.35	34.23
query-2	Is person1 female?	classification	F	H	F	0.91	31.12
query-3	Is person2 detected?	object_definition	T	H	T	0.35	617.74
query-4	Is person2 female?	classification	F	H	F	0.91	28.58
query-5	Is person3 detected?	object_definition	T	H	T	0.35	336.82
query-6	Is person3 female?	classification	T	H	F	0.91	30.94
query-7	Are person1 and person2 together?	relationships	T	H	F	0.87	52.41
query-8	Are person1 and person3 together?	relationships	F	H	F	0.87	50.03
query-9	Are person1 and person2 touching?	relationships	F	H	F	0.87	36.57
storyline-Bike-Cam							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is person1 detected?	object_definition	T	H	T	0.30	803.11

query-2	Is person1 male?	classification	T	H	T	0.30	11.28
query-3	Is person1 sitting?	attributes	F	H	F	0.91	39.57
query-4	Is person1 standing?	attributes	T	H	F	0.91	35.33
query-5	Is person2 detected?	object_definition	T	H	T	0.30	38.87
query-6	Is person2 male?	classification	F	H	F	0.91	34.06
query-7	Is person2 sitting?	attributes	T	H	F	0.91	37.28
query-8	Is person2 standing?	attributes	F	H	F	0.91	41.94
query-9	Is person3 detected?	object_definition	T	H	T	0.30	576.87
query-10	Is person3 male?	classification	T	H	T	0.30	42.79
query-11	Is person3 moving?	tracking	T	H	F	0.91	12.23
query-12	Is person3 walking?	attributes	T	H	T	0.30	61.52
query-13	Is person3 together with at least one other person?	relationships	T	H	F	0.91	41.93
storyline-Georeferencing							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Are there more than two automobiles in [geo bounding box]?	classification	F	H	T	0.69	7.67
query-2	Is there exactly one person in right field?	classification	T	H	T	0.40	13.16
query-3	Is there at least one person in the middle of the parking lot?	classification	F	H	T	0.31	77.61
soc-SIGParkingLot-2014-10-18-Testing							
storyline-tracking-vehicles							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there is an automobile <car1> at pixel (1094,325) in the FOV of sensor RT1?	object_definition	T	H	T	0.44	212.03
query-2	Is obj-car1 moving?	tracking	F	H	T	0.33	8.78
query-3	Is obj-car1 stationary?	tracking	T	H	T	0.38	21.30
query-4	Is there is an automobile<car2> at pixel (246,692) in the FOV of sensor RT1?	object_definition	T	H	T	0.44	17.87
query-5	Is obj-car2 stopping?	tracking	F	H	F	0.91	13.39
query-6	Is obj-car2 moving?	tracking	T	H	F	0.91	42.34

query-7	Is obj-car2 turning?	tracking	T	H	F	0.91	16.28
query-8	Is obj-car2 turning-right?	tracking	F	H	F	0.91	63.21
query-9	Is obj-car2 starting?	tracking	F	H	F	0.91	11.10
query-10	Is there is an automobile<car3> at pixel (1554,623) in the FOV of sensor RT2?	object_definition	T	H	T	0.69	11.51
query-11	Is obj-car3 stopping?	tracking	T	H	T	0.35	12.66
query-12	Is obj-car3 starting?	tracking	T	H	T	0.40	6.95
query-13	Is obj-car3 turning-right?	tracking	F	H	T	0.35	7.95
query-14	Is there is an automobile<car4> at pixel (785,621) in the FOV of sensor C1?	object_definition	T	H	T	0.57	10.81
query-15	Is there is an automobile<car5> at pixel (941,425) in the FOV of sensor GL1?	object_definition	T	H	T	0.57	8.03
query-16	Is there is an automobile<car6> at pixel (496,705) in the FOV of sensor GL6?	object_definition	T	H	T	0.69	8.95
query-17	Is there is an automobile<car7> at pixel (1070,364) in the FOV of sensor GL4?	object_definition	T	H	T	0.44	21.55
query-18	Is there is an automobile<car8> at pixel (143,485) in the FOV of sensor RT1?	object_definition	T	H	T	0.27	17.35
query-19	Is obj-car4 passing obj-car7?	tracking	F	H	F	0.91	28.83
query-20	Is obj-car4 same-motion obj-car7?	tracking	F	H	F	0.91	23.71
query-21	Is obj-car5 following obj-car7?	tracking	F	H	F	0.91	28.86
query-22	Is obj-car5 opposite-motion obj-car7?	tracking	F	H	F	0.91	23.86
query-23	Is obj-car5 same-motion obj-car7?	tracking	T	H	F	0.91	22.99
query-24	Is obj-car7 following obj-car5?	tracking	T	H	F	0.91	24.45
query-25	Is obj-car7 passing	tracking	F	H	F	0.91	26.44

	obj-car5?						
query-26	Is obj-car7 same-motion obj-car8?	tracking	T	H	F	0.91	28.56
query-27	Is obj-car6 passing obj-car7?	tracking	F	H	T	0.95	16.60
query-28	Is obj-car6 passing obj-car8?	tracking	F	H	F	0.91	35.90
query-29	Is obj-car8 turning?	tracking	T	H	T	0.35	8.69
query-30	Is obj-car8 turning-left?	tracking	T	H	F	0.91	10.29
query-31	Does obj-car8 u-turn?	tracking	F	H	F	0.91	13.85
query-32	Is obj-car8 starting?	tracking	T	H	T	0.40	11.92
query-33	Is obj-car8 turning-right?	tracking	T	H	T	0.35	9.37
query-34	Is obj-car8 moving?	tracking	T	H	T	0.40	10.47
query-35	Is obj-car8 stopping?	tracking	T	H	T	0.40	8.80
query-36	Is obj-car8 stationary?	tracking	T	H	T	0.40	8.01
query-37	Is obj-car6 stopping?	tracking	F	H	T	0.57	5.75
query-38	Is obj-car6 u-turn?	tracking	F	H	F	0.91	13.72
query-39	Is obj-car6 moving?	tracking	T	H	T	0.69	6.16
query-40	Is obj-car6 following obj-car4?	tracking	F	H	T	0.95	7.87
query-41	Is obj-car6 opposite-motion obj-car4?	tracking	F	H	T	0.95	14.97
query-42	Is obj-car6 passing obj-car4?	tracking	F	H	F	0.91	32.95

storyline-part-of-relationships

					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there a person <person1> at pixel(1078,410) in the FOV of sensor GL4?	object_definition	T	H	T	0.69	1900.39
query-2	Is there a person <person2> at pixel(1200,456) in the FOV of sensor GL4?	object_definition	T	H	T	0.29	30.48
query-3	Is there a person <person3> at pixel (1390,374) in the FOV sensor GL4?	object_definition	T	H	T	0.29	29.23
query-4	Is there a head <head1> at pixel	object_definition	T	H	T	0.69	7.69

	(1098, 346) in the FOV of sensor GL4						
query-5	Is there a head <head2> at pixel (1204,360) in the FOV of sensor GL4?	object_definition	T	H	T	0.29	117.02
query-6	Is there a head <head3> at pixel (1386,308) in the FOV of sensor GL4?	object_definition	T	H	T	0.29	145.62
query-7	Is there an arm <arm1> at pixel (1052,402) in the FOV of sensor GL4?	object_definition	T	H	UnknownObject		117.98
query-8	Is there an arm <arm2> at pixel (1256,454) in the FOV of sensor GL4?	object_definition	T	H	UnknownObject		118.65
query-9	Is there an arm <arm3> at pixel (1444,370) in the FOV of sensor GL4?	object_definition	T	H	T	0.29	2143.22
query-10	Is there a lower-body <lowerbody1> at pixel (1080,494) in the FOV of sensor GL4?	object_definition	T	H	T	0.69	47.72
query-11	Is there a lower-body <lowerbody2> at pixel (1184,580) in the FOV of sensor GL4?	object_definition	T	H	UnknownObject		145.86
query-12	Is there a lower-body <lowerbody3> at pixel (1408,454) in the FOV of sensor GL4?	object_definition	T	H	T	0.29	116.61
query-13	Is obj-head2 part-of obj-person1?	part_of	F	H	F	0.91	18.64
query-14	Is obj-head1 part-of obj-person3?	part_of	F	H	F	0.91	15.81
query-15	Is obj-arm2 part-of obj-person2?	part_of	T	H	Skipped		0.01
query-16	Is obj-lowerbody3 part-of obj-person1?	part_of	F	H	F	0.91	12.10
query-17	Is obj-lowerbody2 part-of obj-person3?	part_of	F	H	Skipped		0.01
query-18	Is obj-arm1 part-of obj-person1?	part_of	T	H	Skipped		0.01
query-19	Is there an automobile	object_definition	T	H	T	0.69	18.17

	<car1> at pixel (1200,654) in the FOV of sensor C1?						
query-20	Is there a trunk <trunk1> at pixel(1140,676) in the FOV of sensor C1?	object_definition	T	H	T	0.69	123.54
query-21	Is there a hood <hood1> at pixel(878,354) in the FOV of sensor GL3?	object_definition	T	H	UnknownObject		123.18
query-22	Is there a wheel <wheel1> at pixel (810,382) in the FOV of sensor GL3?	object_definition	T	H	UnknownObject		310.81
query-23	Is there an automobile <car2> at pixel (570,578) in the FOV of sensor GL2?	object_definition	T	H	T	0.69	22.80
query-24	Is there a door <door1> at pixel (758,632) in the FOV of sensor GL2?	object_definition	T	H	T	0.69	8.54
query-25	Is there a trunk <trunk2> at pixel (1612, 466) in the FOV of sensor GL1?	object_definition	T	H	T	0.69	4.61
query-26	Is obj-hood1 part-of obj-car1?	part_of	T	H	Skipped		0.01
query-27	Is obj-trunk1 part-of obj-car2?	part_of	F	H	F	0.99	16.97
query-28	Is obj-door1 part-of obj-car1?	part_of	F	H	F	0.99	131.65
query-29	Is obj-trunk2 part-of obj-car2?	part_of	T	H	F	0.99	20.05
query-30	Is obj-wheel1 part-of obj-car2?	part_of	F	H	Skipped		0.01
query-31	Is obj-wheel1 part-of obj-car1?	part_of	T	H	Skipped		0.01
query-32	Is obj-door1 part-of obj-car2?	part_of	T	H	T	0.69	323.33
query-33	Is there an automobile <car3> at pixel (856, 372) in the FOV of sensor RT1?	object_definition	T	H	T	0.27	22.76
query-34	Is there an automobile <car4> at pixel	object_definition	T	H	T	0.27	14.78

	(800,476) in the FOV of sensor RT1?						
query-35	Is there an automobile <car5> at pixel (1030,566) in the FOV of sensor RT1?	object_definition	T	H	T	0.95	10.09
query-36	Is there a wheel <wheel3> at pixel (290,712) in the FOV of sensor GL5?	object_definition	T	H	T	0.49	48.48
query-37	Is there a door <door4> at pixel (618,426) in the FOV of sensor GL4?	object_definition	T	H	T	0.30	11.66
query-38	Is there a wheel <wheel4> at pixel (1274,486) in the FOV of sensor GL2?	object_definition	T	H	UnknownObject		113.74
query-39	Is there a hood <hood4> at pixel (440, 678) in the FOV of sensor GL5?	object_definition	T	H	UnknownObject		117.60
query-40	Is there a trunk <trunk5> at pixel (1120, 730) in the FOV of sensor GL6?	object_definition	T	H	T	0.95	18.42
query-41	Is obj-wheel3 part-of obj-car3?	part_of	T	H	F	0.99	12.23
query-42	Is obj-wheel3 part-of obj-car4?	part_of	F	H	F	0.99	29.92
query-43	Is obj-hood4 part-of obj-car3?	part_of	F	H	Skipped		0.01
query-44	Is obj-hood4 part-of obj-car4?	part_of	T	H	Skipped		0.01
query-45	Is obj-door4 part-of obj-car3?	part_of	F	H	F	0.99	14.30
query-46	Is obj-door4 part-of obj-car4?	part_of	T	H	F	0.99	15.95
query-47	Is obj-door4 part-of obj-car5?	part_of	F	H	F	0.99	10.21
query-48	Is obj-wheel4 part-of obj-car3?	part_of	F	H	Skipped		0.01
query-49	Is obj-wheel4 part-of obj-car4?	part_of	T	H	Skipped		0.01
query-50	Is obj-wheel4 part-of obj-car5?	part_of	F	H	Skipped		0.01

query-51	Is obj-trunk5 part-of obj-car3?	part_of	F	H	F	0.99	14.82
query-52	Is obj-trunk5 part-of obj-car4?	part_of	F	H	F	0.99	10.66
query-53	Is obj-trunk5 part-of obj-car5?	part_of	T	H	T	0.95	30.56
storyline-classification							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Are there less than 4 people in the AOR during time <time-151000-151200>?	classification	F	H	F	0.89	26.99
query-2	Are there at least 4 animals in the AOR during time <time-151000-151200>?	classification	F	H	F	0.89	22.19
query-3	Are there less than 3 cars in the AOR during time <time-151000-151200>?	classification	F	H	F	0.89	11.50
query-4	Are there at least 5 cars in the AOR during time <time-151000-151200>?	classification	F	H	T	0.69	4.66
query-5	Are there at least 2 males in the AOR during time <time-151000-151200>?	classification	T	H	T	0.95	0.38
query-6	Is there at least one disc in the AOR during time <time-151330-151500>?	classification	T	H	T	0.95	41.78
query-7	Is there at least one hat in the AOR during time <time-151330-151500>?	classification	T	H	T	0.95	0.61
query-8	Are there at least 4 items of luggage in the AOR during time <time-151330-151500>?	classification	T	H	T	0.95	17.45
query-9	Are there at least 3 females in the AOR during time <time-151330-151500>?	classification	T	H	T	0.95	8.79
query-10	Are there at least three	classification	F	H	T	0.44	75.39

	Query	Category	Assessor		ucla		Time
	bicycles in the AOR during time <time-151500-151510>?						
storyline-attributes							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there a person <person1> at pixel(302,520) in the FOV of GL1?	object_definition	T	H	T	0.69	24.48
query-2	Is obj-person1 standing?	attributes	T	H	T	0.95	20.42
query-3	Is obj-person1 walking?	attributes	F	H	T	0.95	78.05
query-4	Is obj-person1 pointing?	attributes	F	H	F	0.93	11.34
query-5	Is obj-person1 sitting?	attributes	T	H	F	0.93	9.98
query-6	Is obj-person1 talking?	attributes	F	H	Other		12.51
query-7	Is obj-person1 running?	attributes	F	H	F	0.93	10.19
query-8	Is obj-person1 crawling?	attributes	F	H	F	0.93	9.85
query-9	Is there a person <person2> at pixel(917,349) in the FOV of GL4	object_definition	T	H	T	0.30	18.62
query-10	Is obj-person2 standing?	attributes	T	H	T	0.30	5.44
query-11	Is obj-person2 walking?	attributes	T	H	T	0.30	3.03
query-12	Is obj-person2 pointing?	attributes	T	H	F	0.99	13.46
query-13	Is obj-person2 sitting?	attributes	F	H	F	0.99	9.50
query-14	Is obj-person2 talking?	attributes	T	H	Other		8.10
query-15	Is obj-person2 running?	attributes	F	H	F	0.99	12.93
query-16	Is obj-person2 crawling?	attributes	F	H	F	0.99	10.22
query-17	Is there a person <person3> at pixel (295,414) in the FOV of sensor RT1?	object_definition	T	H	UnknownObject		124.42
query-18	Is obj-person3 standing?	attributes	T	H	Skipped		0.01

query-19	Is obj-person3 walking?	attributes	T	H	Skipped		0.00
query-20	Is obj-person3 pointing?	attributes	F	H	Skipped		0.00
query-21	Is obj-person3 sitting?	attributes	F	H	Skipped		0.01
query-22	Is obj-person3 swinging a small-object?	relationships	T	H	Skipped		0.01
query-23	Is obj-person3 running?	attributes	T	H	Skipped		0.00
query-24	Is obj-person3 crawling?	attributes	F	H	Skipped		0.01
query-25	Is there a person <person4> at pixel(279,388) in the FOV of GL3	object_definition	T	H	UnknownObject		246.44
query-26	Is obj-person4 standing?	attributes	F	H	Skipped		0.01
query-27	Is obj-person4 walking?	attributes	F	H	Skipped		0.03
query-28	Is obj-person4 pointing?	attributes	F	H	Skipped		0.01
query-29	Is obj-person4 sitting?	attributes	T	H	Skipped		0.01
query-30	Is obj-person4 talking?	attributes	F	H	Skipped		0.01
query-31	Is obj-person4 running?	attributes	F	H	Skipped		0.01
query-32	Is obj-person4 crawling?	attributes	F	H	Skipped		0.01
query-33	Is there a door <door1> at pixel(234,541) in the FOV of GL1	object_definition	T	H	UnknownObject		1147.50
query-34	Is obj-door1 open?	attributes	F	H	Skipped		0.01
query-35	Is obj-door1 closed?	attributes	T	H	Skipped		0.01
query-36	Is obj-door1 open?	attributes	T	H	Skipped		0.01
query-37	Is obj-door1 closed?	attributes	F	H	Skipped		0.01
storyline-spatial-relationships							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there a automobile <car1> at pixel(131,486) in the FOV of sensor RT1?	object_definition	T	H	T	0.27	11.13

query-2	Is there an automobile <car2> at pixel(722,565) in the FOV of sensor RT1?	object_definition	T	H	T	0.69	539.96
query-3	Is there an automobile <car3> at pixel (1008,691) in the FOV of sensor RT1?	object_definition	T	H	UnknownObject		132.46
query-4	Is there an automobile <car4> at pixel (1089,373) in the FOV of sensor RT1?	object_definition	T	H	UnknownObject		147.08
query-5	Is there an automobile <car5> at pixel (1099,322) in the FOV of sensor RT1?	object_definition	T	H	T	0.44	15.41
query-6	Is there an automobile <car6> at pixel (1653,392) in the FOV of sensor RT1?	object_definition	T	H	UnknownObject		117.69
query-7	Is there an automobile <car7> at pixel (1671,504) in the FOV of sensor RT1?	object_definition	T	H	T	0.69	3720.65
query-8	Is obj-car6 facing obj-car7?	spatial	F	H	Skipped		0.01
query-9	Is obj-car6 facing-opposite obj-car7?	spatial	F	H	Skipped		0.01
query-10	Is obj-car5 facing obj-car7?	spatial	F	H	Other		18.49
query-11	Is obj-car5 facing-opposite obj-car7?	spatial	T	H	F	0.91	39.06
query-12	Is obj-car3 clear-line-of-sight obj-car4?	spatial	T	H	Skipped		0.01
query-13	Is physical distance between car2 and car1 less than the physical distance between car3 and car1?	spatial	T	H	Skipped		0.01
query-14	Is physical distance between car2 and car7 less than the physical distance between car2 and car3?	spatial	F	H	Skipped		0.01
query-15	Is physical distance between car1 and car4 greater than the	spatial	T	H	Skipped		0.01

	physical distance between car5 and car4?						
query-16	Is car2 blocking the view from car3 to car5?	spatial	F	H	Skipped		0.01
query-17	Is car4 blocking the view from car3 to car5?	spatial	T	H	Skipped		0.01
storyline-relationships							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there a person <person1> at pixel(1013,678) in the FOV of sensor GL5	object_definition	T	H	T	0.27	22.05
query-2	Is there a two-wheeled-vehicle<bike1> at pixel(1026,773) in the FOV of sensor GL5	object_definition	T	H	T	0.27	42.29
query-3	Is obj-person1 dismounting obj-bike1?	relationships	F	H	Other		16.69
query-4	Is obj-person1 on obj-bike1?	relationships	T	H	F	0.91	11.12
query-5	Is obj-person1 touching obj-bike1?	relationships	T	H	F	0.91	622.54
query-6	Is there a person <person2> at pixel (1248,312) in the FOV of sensor GL3	object_definition	T	H	T	0.44	90.17
query-7	Is there a person <person3> at pixel (1209,309) in the FOV of sensor GL3	object_definition	T	H	UnknownObject		96.34
query-8	Is obj-person2 touching obj-person3?	relationships	F	H	Skipped		0.01
query-9	Is obj-person2 together obj-person3?	relationships	T	H	Skipped		0.01
query-10	Is there a person <person4> at pixel (398,447) in the FOV of sensor GL1?	object_definition	T	H	T	0.95	6.65
query-11	Is there a person <person5> at pixel	object_definition	T	H	T	0.95	19.84

	(291,440) in the FOV of sensor GL1?						
query-12	Is there a tool <tool1> at pixel (337,378) in the FOV of sensor GL1?	object_definition	T	H	T	0.95	1931.56
query-13	Is there a luggage <backpack1> at pixel (1280,479) in the FOV of sensor GL1?	object_definition	T	H	T	0.95	11.59
query-14	Is there luggage <backpack2> at pixel (128,573) in the FOV of sensor GL1?	object_definition	T	H	T	0.95	18.45
query-15	Is there a hat <hat1> at pixel (320,340) in the FOV of sensor GL1?	object_definition	T	H	T	0.95	4.29
query-16	Is obj-person4 carrying obj-tool1?	relationships	F	H	F	0.89	31.41
query-17	Is obj-person5 carrying obj-tool1?	relationships	T	H	F	0.89	9.53
query-18	Is obj-person5 wearing obj-hat1?	relationships	F	H	F	0.89	19.92
query-19	Is obj-person4 donning obj-hat1?	relationships	F	H	Other		12.83
query-20	Is obj-person4 doffing obj-hat1?	relationships	F	H	Other		14.17
query-21	Is obj-person5 carrying obj-backpack1?	relationships	F	H	F	0.89	13.74
query-22	Is obj-person4 carrying obj-backpack1?	relationships	F	H	F	0.89	13.92
query-23	Is obj-person5 donning obj-hat1?	relationships	F	H	Other		7.58
query-24	Is there a small-object<small-object1> at pixel (84,592) in the FOV of sensor GL1?	object_definition	T	H	T	0.95	7.62
query-25	Is obj-person5 touching obj-smallobject1?	relationships	T	H	Other		47.47
query-26	Is obj-person5 carrying obj-smallobject1?	relationships	T	H	F	0.89	11.64
query-27	Is obj-person5	relationships	T	H	Other		10.23

	dropping obj-smallobject1?						
query-28	Is obj-person5 unloading obj-backpack2?	relationships	T	H	Other		10.31
query-29	Is obj-person4 picking-up obj-smallobject1?	relationships	F	H	F	0.89	11.73
query-30	Is obj-smallobject1 outside obj-backpack1?	relationships	T	H	F	0.89	13.13
query-31	Is obj-smallobject1 inside obj-backpack1?	relationships	T	H	F	0.89	13.76
query-32	Is there a person <person6> at pixel (1741,756) in the FOV of sensor C1?	object_definition	T	H	T	0.23	21.15
query-33	Is there a disc <disc1> at pixel (1436,1041) in the FOV of sensor C1?	object_definition	T	H	UnknownObject		99.47
query-34	Is there a disc <disc2> at pixel (1655,846) in the FOV of sensor C1?	object_definition	T	H	UnknownObject		96.14
query-35	Is there luggage <shoulderbag1> at pixel (1508,948) in the FOV of sensor C1?	object_definition	T	H	T	0.23	106.84
query-36	Is obj-person6 throwing obj-disc2?	relationships	T	H	Skipped		0.01
query-37	Is obj-person6 dropping obj-disc2?	relationships	F	H	Skipped		0.01
query-38	Is obj-person6 catching obj-disc2?	relationships	F	H	Skipped		0.01
query-39	Is obj-person6 putting-down obj-disc2?	relationships	F	H	Skipped		0.01
query-40	Is obj-person6 picking-up obj-disc2?	relationships	T	H	Skipped		0.01
query-41	Is obj-person6 throwing obj-disc1?	relationships	T	H	Skipped		0.01
query-42	Is obj-person6 picking-up obj-disc1?	relationships	T	H	Skipped		0.01
query-43	Is obj-person6 carrying obj-shoulderbag1?	relationships	T	H	F	0.89	8.50
query-44	Is obj-person6 dropping obj-	relationships	F	H	Other		17.77

	shoulderbag1?						
query-45	Is obj-person6 putting-down obj-shoulderbag1?	relationships	T	H	F	0.89	13.64
query-46	Is obj-person6 picking-up obj-shoulderbag1?	relationships	T	H	F	0.89	12.06
query-47	Is there a car <car1> at pixel (168,485) in the FOV of sensor RT1?	object_definition	T	H	T	0.44	23.91
query-48	Is there a person <person7> at pixel (570,526) in the FOV of sensor RT1?	object_definition	T	H	UnknownObject		86.73
query-49	Is obj-person7 inside obj-car1?	relationships	T	H	Skipped		0.01
query-50	Is obj-person7 outside obj-car1?	relationships	T	H	Skipped		0.01
query-51	Is obj-person7 mounting obj-car1?	relationships	T	H	Skipped		0.01
query-52	Is obj-person7 dismounting obj-car1?	relationships	F	H	Skipped		0.01
query-53	Is obj-person7 driving obj-car1?	relationships	T	H	Skipped		0.02
query-54	Is there a car <car2> at pixel (1463,413) in the FOV of sensor GL4?	object_definition	T	H	T	0.95	8.63
query-55	Is there a person <person8> at pixel (1153,491) in the FOV of sensor GL4?	object_definition	T	H	T	0.95	7.29
query-56	Is obj-person6 same-object obj-person8?	relationships	T	H	F	0.99	11.35
query-57	Is obj-person8 inside obj-car2?	relationships	T	H	T	0.44	8.72
query-58	Is obj-person8 outside obj-car2?	relationships	T	H	T	0.44	7.90
query-59	Is obj-person8 mounting obj-car2?	relationships	T	H	F	0.91	44.28
query-60	Is obj-person8 dismounting obj-car2?	relationships	F	H	F	0.91	37.02
query-61	Is obj-person8 driving obj-car2?	relationships	T	H	F	0.91	38.04
query-62	Is there a person <person9> at pixel	object_definition	T	H	T	0.25	18.14

	Query	Category	Assessor		ucla		Time
	(1543,515) in the FOV of sensor RT1?						
query-63	Is there a trunk <trunk1> at pixel (1590,525) in the FOV of sensor RT1?	object_definition	T	H	T	0.44	7.49
query-64	Is there a trunk <trunk2> at pixel (1590,397) in the FOV of sensor RT1?	object_definition	T	H	T	0.69	81.22
query-65	Is there an automobile <car3> at pixel (1674,507) in the FOV of sensor RT1?	object_definition	T	H	T	0.44	14.78
query-66	Is there an automobile <car4> at pixel (1664,392) in the FOV of sensor RT1?	object_definition	T	H	T	0.69	7.63
query-67	Is obj-person9 mounting obj-car3?	relationships	T	H	F	0.91	47.49
query-68	Is obj-person9 dismounting obj-car3?	relationships	F	H	F	0.91	46.79
query-69	Is obj-person9 driving obj-car3?	relationships	F	H	F	0.91	43.30
query-70	Is obj-person9 loading obj-trunk1?	relationships	T	H	Other		5.51
query-71	Is obj-person9 unloading obj-trunk1?	relationships	F	H	Other		57.57
query-72	Is obj-person9 loading obj-trunk2?	relationships	F	H	Other		7.00
query-73	Is obj-person9 unloading obj-trunk2?	relationships	T	H	Other		13.23
query-74	Is obj-person9 inside obj-car4?	relationships	F	H	T	0.38	9.69
query-75	Is obj-person9 mounting obj-car4?	relationships	F	H	T	0.95	18.70
query-76	Is obj-person9 dismounting obj-car4?	relationships	F	H	F	0.91	41.98
query-77	Is obj-person9 driving obj-car4?	relationships	F	H	T	0.95	23.98
soc-PrattGarden-2014-09-20-Testing							
storyline-exercise-class							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there a person	object_definition	T	H	T	0.57	159.20

	<person1> in the FOV of HC1 at pixel (1113,475)?						
query-2	Is <person1> on ground?	relationships	T	H	T	0.57	16.82
query-3	Is there a person <person2> in the FOV of HC2 at pixel (814,717)?	object_definition	T	H	F	0.87	24.33
query-4	Is <person2> touching ground?	relationships	T	H	Other		23.31
query-5	Is obj-person1 same-object obj-person2?	relationships	F	H	F	0.83	19.45
query-6	Is obj-person1 facing obj-person2?	spatial	T	H	F	0.83	372.75
query-7	Is obj-person1 touching obj-person2?	relationships	F	H	F	0.83	13.50
query-8	Is obj-person1 together obj-person2?	relationships	F	H	F	0.83	13.01
query-9	Is obj-person1 clear-line-of-sight obj-person2?	spatial	T	H	F	0.83	14.80
query-10	Does <person1> throw a small-object?	relationships	F	H	F	0.83	168.42
query-11	Does <person2> drop a small-object?	relationships	F	H	Other		10.37
query-12	Is there an animal in the AOR?	classification	T	H	T	0.95	3.73
query-13	Is there a person together with an animal?	relationships	T	H	F	0.83	568.13
query-14	Is there a hand <hand1> in the FOV of HC2 at pixel (681,201)?	object_definition	T	H	T	0.49	22.46
query-15	Is <hand1> part-of <person2>?	part_of	T	H	F	0.83	25.04
query-16	Is there a head <head1> in the FOV of HC2 at pixel (634,363)?	object_definition	T	H	T	0.12	8900.67
query-17	Is <head1> part-of <person2>?	part_of	T	H	F	0.83	14.77
query-18	Is <hand1> below <head1>?	relationships	F	H	F	0.83	14.23

query-19	Is <hand1> touching another person (not <person2>)?	relationships	F	H	T	0.49	15.28
query-20	Is there a person in the FOV of IP1 with hands below head?		F	H	F	0.83	98.53
query-21	Is there a person touching another person?	relationships	F	H	T	0.95	71.04
query-22	Is <person1> walking?	attributes	F	H	F	0.83	103.25
query-23	Is there a person walking?	attributes	T	H	T	0.57	24.90
query-24	Is there a person running?	attributes	T	H	T	0.95	13.27
query-25	Is <person1> running?	attributes	F	H	F	0.83	48.90
query-26	Is <person2> running?	attributes	T	H	F	0.83	35.61
query-27	Is <person1> crawling?	attributes	F	H	F	0.83	19.41
query-28	Is <person1> pointing?	attributes	T	H	F	0.83	22.39
query-29	Are there two-people moving in the same direction (same-motion)?	tracking	T	H	T	0.95	2965.79
query-30	Is there a person following another person?	tracking	T	H	T	0.95	2788.71
query-31	Is there a person running and turning right?		F	H	T	0.44	205.66
query-32	Are there at least 10 people moving?	tracking	T	H	T	0.95	28.68
query-33	Is there a stationary person?	tracking	T	H	T	0.95	17.33
query-34	Is there a clear line of sight between <person1> and <person2>?	spatial	T	H	F	0.83	49.96
query-35	Is there a person sitting?	attributes	T	H	T	0.95	9.07
query-36	Is there a chair?	classification	T	H	F	0.83	14.51
query-37	Is there a table?	classification	F	H	F	0.83	14.60
query-38	Is there a person on a chair?	relationships	T	H	F	0.83	70.10
query-39	Is there a person	object_definition	T	H	T	0.57	17.13

	<person3> in the FOV of HC1 at pixel (298,499)?						
query-40	Is <person1> female?	classification	F	H	F	0.83	12.40
query-41	Is <person2> female?	classification	T	H	F	0.83	9.80
query-42	Is <person3> female?	classification	F	H	F	0.83	12.89
query-43	Is there a moving vehicle?	tracking	F	H	F	0.83	12.80
query-44	Is there a person facing <person3>?	spatial	T	H	F	0.83	21.74
query-45	Is there a luggage <luggage1> in the FOV of IP5 at pixel (254,412)?	object_definition	T	H	T	0.69	18.24
query-46	Does <person3> pick up <luggage1>?	relationships	F	H	F	0.83	17.68
query-47	Does <person3> put down <luggage1>?	relationships	F	H	F	0.83	13.05
query-48	Does <person3> load <luggage1>?	relationships	T	H	Other		7.84
query-49	Does <person3> doff top-wear?	relationships	T	H	Other		8.23
query-50	Does <person3> don a hat?	relationships	F	H	Other		6.45
query-51	Does <person3> carry <luggage1>?	relationships	F	H	F	0.83	21.88
query-52	Is there a person <person4> in the FOV of IP1 at pixel (876,534)?	object_definition	T	H	T	0.95	2.53
query-53	Is there a small-object <small-object1> in the FOV of IP1 at pixel (965,523)?	object_definition	T	H	T	0.95	6.35
query-54	Does <person4> pick up <small-object1>?	relationships	T	H	F	0.83	22.61
query-55	Does <person4> carry <small-object1>?	relationships	T	H	T	0.95	3.98
query-56	Does <person4> drop <small-object1>?	relationships	T	H	Other		9.44
query-57	Is <small-object1> the same object as <luggage1>?	relationships	F	H	F	0.83	11.85
query-58	Is <person4> stationary?	tracking	T	H	T	0.95	4.88

query-59	Is there a person <person5> in the FOV of IP1 at pixel (172,573)?	object_definition	T	H	T	0.95	5.01
query-60	Is <person5> stationary?	tracking	T	H	T	0.95	3.31
query-61	Is <person5> the same-object as <person1>?	relationships	F	H	F	0.83	13.79
query-62	Does <person5> touch <small-object1>?	relationships	F	H	F	0.83	524.68
query-63	Does <person5> touch small-object?	relationships	F	H	T	0.95	17.32
query-64	Are there at least 2 people in the geodetic polygon?	classification	T	H	T	0.95	13.31
query-65	Is there a person <person6> in the FOV of HC3 at pixel (214,368)?	object_definition	T	H	T	0.44	20.29
query-66	Is there a person <person7> in the FOV of HC3 at pixel (920,354)?	object_definition	T	H	T	0.57	27.40
query-67	Is there a lower-body <lb1> in the FOV of HC3 at pixel (150,532)?	object_definition	T	H	T	0.44	1844.82
query-68	Is there a lower-body <lb2> in the FOV of HC3 at pixel (902,392)?	object_definition	T	H	T	0.57	17.20
query-69	Is there a lower-body <lb3> in the FOV of HC2 at pixel (1156,614)?	object_definition	T	H	T	0.49	15.29
query-70	Is <lb1> part of <person6>?	part_of	T	H	F	0.87	67.30
query-71	Is <lb2> part of <person7>?	part_of	T	H	T	0.57	17.36
query-72	Is <lb3> part of <person7>?	part_of	T	H	F	0.87	31.84
storyline-fashion-show							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there a person	object_definition	T	H	T	0.35	56.22

	<person1> in the FOV of MC1 at pixel (934,597)?						
query-2	Is <person1> turning?	tracking	F	H	F	0.83	14.48
query-3	Is <person1> moving?	tracking	F	H	F	0.83	16.55
query-4	Is <person1> sitting?	attributes	T	H	F	0.83	15.13
query-5	Is <person1> female?	classification	F	H	F	0.83	15.68
query-6	Is there a person <person2> in the FOV of Contour2 at pixel (716,472)?	object_definition	T	H	UnknownObject		942.07
query-7	Is there an animal <animal1> in the FOV of Contour2 at pixel (586,687)?	object_definition	T	H	T	0.38	21.74
query-8	Are <person1> and <animal1> together?	relationships	F	H	F	0.83	13.26
query-9	Is <person1> closer to <animal1> than <person2>?	spatial	F	H	Skipped		0.01
query-10	Is <animal1> farther from <person1> than <person2>?	spatial	T	H	Skipped		0.01
query-11	Are <person2> and <animal1> together?	relationships	T	H	Skipped		0.01
query-12	Are <person2> and <animal1> touching?	relationships	F	H	Skipped		0.01
query-13	Is <animal1> moving?	tracking	F	H	F	0.83	21.24
query-14	Is there at least one person moving?	tracking	T	H	T	0.95	9.35
query-15	Is there at least one animal moving?	tracking	F	H	F	0.83	15.27
query-16	Is there a person <person3> in the FOV of IP1 at pixel (727,547)?	object_definition	T	H	T	0.95	2.49
query-17	Is <person3> standing?	attributes	T	H	F	0.83	28.08
query-18	Is <person3> moving?	tracking	T	H	T	0.95	8.05
query-19	Is<person3> turning?	tracking	T	H	F	0.83	12.75
query-20	Is <person3> doing a u-turn?	tracking	F	H	F	0.83	67.23
query-21	Is <person3> turning-right?	tracking	F	H	F	0.83	15.22

query-22	Is <person3> running?	attributes	F	H	F	0.83	14.95
query-23	Is <person3> walking?	attributes	T	H	T	0.95	2.97
query-24	Is <person3> stopping?	tracking	F	H	T	0.95	3.78
query-25	Is <person3> stationary?	tracking	F	H	T	0.95	3.79
query-26	Is there a small-object <small-object1> in the FOV of IP1 at pixel (65,807)?	object_definition	T	H	T	0.95	6.62
query-27	Is there a person <person4> in the FOV of IP1 at pixel (1290,501)?	object_definition	T	H	T	0.13	154.47
query-28	Is <person4> talking?	attributes	T	H	Other		7.85
query-29	Is <person4> crawling?	attributes	F	H	F	0.83	13.78
query-30	Is <person4> picking up <small-object1>?	relationships	T	H	F	0.83	13.63
query-31	Is <person4> putting down <small-object1>?	relationships	F	H	F	0.83	16.76
query-32	Is <person4> carrying <small-object1>?	relationships	T	H	F	0.83	14.61
query-33	Does <person3> have a CLOS to <person4>?	spatial	T	H	T	0.95	15.55
query-34	Is there a person <person5> in the FOV of HC2 at pixel (373,433)?	object_definition	T	H	T	0.49	23.61
query-35	Is <person5> doffing top-wear?	relationships	F	H	Other		5.81
query-36	Is <person5> donning a hat?	relationships	F	H	Other		6.04
query-37	Is <person5> female?	classification	T	H	F	0.83	11.79
query-38	Is <person5> starting?	tracking	T	H	F	0.83	13.35
query-39	Is <person5> walking?	attributes	T	H	F	0.83	14.34
query-40	Is <person5> turning?	tracking	T	H	F	0.83	12.82
query-41	Is <person5> pointing?	attributes	T	H	Other		21.48
query-42	Is <person5> crawling?	attributes	F	H	F	0.83	13.14
query-43	Is there a person	object_definition	T	H	T	0.44	19.30

	<person6> in the FOV of HC3 at pixel (985,245)?						
query-44	Is <person6> doffing top-wear?	relationships	F	H	Other		4.34
query-45	Is <person6> donning a hat?	relationships	T	H	Other		2.17
query-46	Is there a person <person7> in the FOV of MC1 at pixel (801,485)?	object_definition	T	H	T	0.35	27.26
query-47	Is there a hat <hat1> in the FOV of MC1 at pixel 770,568)?	object_definition	T	H	T	0.11	205.13
query-48	Is <person7> picking up <hat1>?	relationships	T	H	F	0.83	19.80
query-49	Is <person7> donning <hat1>?	relationships	T	H	Other		7.10
query-50	Is there a person <person8> in the FOV of IP1 at pixel (1259,550)?	object_definition	T	H	T	0.95	5.37
query-51	Are <person8> and <person6> the same-object?	relationships	F	H	F	0.83	11.82
query-52	Are <person8> and <person7> the same-object?	relationships	T	H	F	0.83	15.85
query-53	Is there a person <person9> in the FOV of IP1 at pixel (1474,607)?	object_definition	T	H	T	0.95	39.73
query-54	Are <person9> and <person6> the same-object?	relationships	T	H	F	0.83	14.23
query-55	Are <person9> and <person5> the same-object?	relationships	F	H	F	0.83	15.29
query-56	Is there an arm <arm1> in the FOV of IP1 at pixel (1319,612)?	object_definition	T	H	T	0.95	6.87
query-57	Is <arm1> part of <person9>?	part_of	F	H	F	0.83	12.38
query-58	Are there at least 2 people in the geodetic polygon?	classification	F	H	T	0.95	3.81

query-59	Is there a person <person10> in the FOV of HC3 at pixel (76,428)?	object_definition	T	H	T	0.44	18.63
query-60	Is there a person <person11> in the FOV of HC3 at pixel (166,410)?	object_definition	T	H	T	0.95	15.21
query-61	Is there a person <person12> in the FOV of HC3 at pixel (272,486)?	object_definition	T	H	T	0.95	11.65
query-62	Is there a head <head1> in the FOV of HC3 at pixel (196,338)?	object_definition	T	H	T	0.95	16.13
query-63	Is <head1> part of <person10>?	part_of	F	H	F	0.87	12.61
query-64	Is <head1> part of <person11>?	part_of	T	H	T	0.95	4.13
query-65	Is there a head <head2> in the FOV of HC2 at pixel (126,838)?	object_definition	T	H	T	0.49	172.30
query-66	Is there an arm <arm2> in the FOV of HC2 at pixel (1266,836)?	object_definition	T	H	UnknownObject		239.68
query-67	Is <head2> part of <person12>?	part_of	F	H	F	0.87	14.93
query-68	Is <head2> part of <person11>?	part_of	F	H	F	0.87	13.75
query-69	Is <head2> part of <person10>?	part_of	T	H	F	0.87	11.77
query-70	Is <arm2> part of <person10>?	part_of	F	H	Skipped		0.01
query-71	Is <arm2> part of <person11>?	part_of	F	H	Skipped		0.01
storyline-sports							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there a person <person1> in the FOV of HC1 at pixel (719,522)?	object_definition	T	H	T	0.57	131.04
query-2	Is there a two-	object_definition	T	H	T	0.57	8.06

	wheeled-vehicle <bicycle1> in the FOV of HC1 at pixel (868,807)?						
query-3	Are <person1> and <bicycle1> together?	relationships	T	H	F	0.79	15.93
query-4	Is <person1> mounting <bicycle1>?	relationships	F	H	Other		32.50
query-5	Is <person1> dismounting <bicycle1>?	relationships	F	H	Other		6.25
query-6	Is <person1> driving <bicycle1>?	relationships	F	H	F	0.79	24.92
query-7	Is there a person <person2> in the FOV of HC2 at pixel (621,421)?	object_definition	T	H	T	0.49	16.35
query-8	Is there a person <person3> in the FOV of HC2 at pixel (950,386)?	object_definition	T	H	T	0.49	19.20
query-9	Is there a person <person4> in the FOV of HC2 at pixel (1492,460)?	object_definition	T	H	T	0.49	13.78
query-10	Is there a ball <ball1> in the FOV of HC2 at pixel (933,414)?	object_definition	T	H	UnknownObject		3260.55
query-11	Is <person3> throwing <ball1>?	relationships	T	H	Skipped		0.01
query-12	Is <person4> throwing <ball1>?	relationships	F	H	Skipped		0.01
query-13	Is <person2> catching <ball1>?	relationships	T	H	Skipped		0.01
query-14	Is <person4> running?	attributes	F	H	F	0.79	16.61
query-15	Is <person2> passing <person3>?	tracking	F	H	T	0.95	411.24
query-16	Is <person2> crawling?	attributes	F	H	F	0.79	12.93
query-17	Does <person4> have a clear-line-of-sight to <person3>?	spatial	T	H	T	0.49	8.92
query-18	Is there a person <person5> in the FOV of HC1 at pixel (711,564)?	object_definition	T	H	T	0.57	25.96

query-19	Is there a person <person6> in the FOV of HC1 at pixel (506,564)?	object_definition	T	H	T	0.57	9.46
query-20	Is <person5> walking?	attributes	F	H	F	0.79	17.05
query-21	Is <person6> stationary?	tracking	F	H	T	0.57	5.50
query-22	Is <person5> following <person6>?	tracking	F	H	F	0.79	13.26
query-23	Is <person6> following <person5>?	tracking	T	H	F	0.79	15.72
query-24	Is <person5> touching <ball1>?	relationships	F	H	Skipped		0.01
query-25	Is <person5> touching <ball1>?	relationships	T	H	Skipped		0.01
query-26	Is <person5> crawling?	attributes	T	H	F	0.79	12.80
query-27	Is <person6> crawling?	attributes	T	H	F	0.79	12.82
query-28	Do <person5> and <person6> have the same motion (direction)?	tracking	T	H	F	0.79	11.89
query-29	Is there a person <person7> in the FOV of HC1 at pixel (1477,495)?	object_definition	T	H	T	0.57	10.78
query-30	Are <person7> and <person1> the same object?	relationships	T	H	F	0.79	13.00
query-31	Is <person7> running?	attributes	F	H	F	0.79	13.72
query-32	Is <person7> stationary?	tracking	F	H	T	0.57	7.71
query-33	Is <person7> male?	classification	T	H	T	0.57	5.61
query-34	Is there a luggage <luggage1> in the FOV of HC1 at pixel (876,700)?	object_definition	T	H	T	0.11	236.55
query-35	Does <person7> load <luggage1>?	relationships	T	H	Other		6.91
query-36	Does <person7> carry <luggage1>?	relationships	T	H	F	0.79	14.71
query-37	Does <person7> mount a bicycle?	relationships	F	H	Other		6.12
query-38	Does <person7>	relationships	T	H	Other		3.84

	mount a bicycle?						
query-39	Is there a person <person8> in the FOV of HC3 at pixel (610,231)?	object_definition	T	H	T	0.13	191.71
query-40	Is <person8> throwing a small-object?	relationships	T	H	F	0.79	116.46
query-41	Is <person8> female?	classification	T	H	T	0.44	12.63
query-42	Is there a person <person9> in the FOV of HC3 at pixel (644,446)?	object_definition	T	H	T	0.44	24.38
query-43	Is there a person <person10> in the FOV of HC3 at pixel (519,392)?	object_definition	T	H	T	0.44	27.08
query-44	Is there a person <person11> in the FOV of HC3 at pixel (1246,381)?	object_definition	T	H	T	0.44	19.72
query-45	Is there a head <head1> in the FOV of HC3 at pixel (638,239)?	object_definition	T	H	T	0.44	40.69
query-46	Is there a hand <hand1> in the FOV of HC3 at pixel (737,522)?	object_definition	T	H	T	0.13	225.86
query-47	Is there a lower-body <lower-body1> in the FOV of HC3 at pixel (1232,540)?	object_definition	T	H	T	0.44	24.42
query-48	Is <hand1> part-of <person10>?	part_of	F	H	F	0.79	15.95
query-49	Is <hand1> part-of <person9>?	part_of	T	H	T	0.44	13.17
query-50	Is <lower-body1> part-of <person9>?	part_of	F	H	F	0.79	12.52
query-51	Is <head1> part-of <person9>?	part_of	T	H	T	0.44	11.83
query-52	Is <head1> below <hand1>?	relationships	F	H	F	0.79	12.56
query-53	Is <hand1> touching <person10>?	relationships	F	H	T	0.44	10.66
query-54	Is <person9> talking?	attributes	T	H	Other		5.62
query-55	Is <person11> talking?	attributes	F	H	Other		14.76

query-56	Is <lower-body1> part-of <person11>?	part_of	T	H	T	0.44	9.80
query-57	Is <person9> occluding <person10> in the FOV of HC3?	spatial	T	H	Other		2.69
query-58	Is <person9> farther from <person10> than <person11>?	spatial	F	H	T	0.95	512.56
query-59	Is <person11> stationary?	tracking	F	H	F	0.79	17.81
query-60	Is there a person <person12> in the FOV of IP5 at pixel (309,280)?	object_definition	T	H	T	0.69	8.73
query-61	Is there a luggage <luggage2> in the FOV of IP5 at pixel (279,326)?	object_definition	T	H	T	0.69	7.71
query-62	Is there an arm <arm1> in the FOV of IP5 at pixel (332,304)?	object_definition	T	H	UnknownObject		238.27
query-63	Is there a head <head2> in the FOV of IP5 at pixel (288,250)?	object_definition	T	H	T	0.11	1947.92
query-64	Are <person12> and <person11> the same object?	relationships	F	H	F	0.79	19.23
query-65	Are <person12> and <person7> the same object?	relationships	T	H	F	0.79	11.84
query-66	Is <person12> standing?	attributes	F	H	T	0.69	5.88
query-67	Is <person12> talking?	attributes	F	H	F	0.79	13.77
query-68	Is <person12> sitting?	attributes	T	H	F	0.79	15.35
query-69	Is <head2> part-of <person12>?	part_of	T	H	F	0.79	13.40
query-70	Is <arm1> part-of <person12>?	part_of	T	H	Skipped		0.01
query-71	Is <person12> unloading <luggage2>?	relationships	F	H	F	0.79	15.22
query-72	Are there at least 2 people in the geodetic polygon?	classification	F	H	T	0.95	5.49
soc-Schiciano-2014-02-22-Testing							

storyline-classification							
					ucla		
	Query	Category	Assessor				Time
query-1	Is there at least one table in the loc-aud-entrance?	classification	T	H	T	0.35	125.35
query-2	Are there more than 10 people in the FOV of obs-HC3 at time-registration-start?	classification	F	H	T	0.35	8.28
query-3	Are there at least 5 items of luggage in the FOV of obs-HC2 at time 19:28:59?	classification	T	H	T	0.38	17.67
query-4	Are there at least 7 people in loc-auditorium-left?	classification	T	H	T	0.44	30.38
query-5	Are there at least 2 people standing in loc-auditorium-left?	attributes	F	H	F	0.91	10.71
query-6	Are there at least 2 hats in loc-aud-entrance?	classification	T	H	F	0.91	7.88
query-7	Are there any vehicles in loc-aud-entrance?	classification	F	H	F	0.91	11.24
query-8	Are there any chairs in loc-aud-entrance?	classification	T	H	T	0.38	7.79

storyline-part-of-relationships							
					ucla		
	Query	Category	Assessor				Time
query-1	Identify person obj-instructor1	object_definition	T	H	T	0.95	9.45
query-2	Identify arm.	object_definition	T	H	T	0.95	6.64
query-3	Is obj-arm part of obj-instructor1?	part_of	T	H	T	0.95	6.27
query-4	Identify obj-head	object_definition	T	H	T	0.69	1367.62
query-5	Is obj-head part of obj-instructor1?	part_of	F	H	F	1.00	23.37
query-6	Identify obj-student1.	object_definition	T	H	T	0.26	31.01
query-7	Identify obj-lower-body.	object_definition	T	H	T	0.26	15.81
query-8	Is obj-lower-body part of obj-student1?	part_of	F	H	T	0.26	16.62
query-9	Identify obj-hand at	object_definition	T	H	T	0.26	16.68

	(1411, 924).						
query-10	Is obj-hand part of obj-student1?	part_of	T	H	T	0.26	9.22
query-11	Is there at least one room in obs-IP2?	classification	T	H	F	0.91	8.57
query-12	Identify Wall.	object_definition	T	H	T	0.69	71.10
query-13	Is obj-wall part of a room?	part_of	T	H	F	0.91	8.93
query-14	Are there at least 3 doors that are part of obj-wall?	part_of	F	H	F	0.91	15.14
query-15	Identify obj-student2 at (1509, 712).	object_definition	T	H	T	0.22	18.39
query-16	Identify obj-student3 at (1310,618).	object_definition	T	H	T	0.22	26.85
query-17	Identify obj-head2 at (1502, 614)	object_definition	T	H	T	0.22	124.34
query-18	Idenity obj-arm2 at (1417, 777).	object_definition	T	H	T	0.22	144.05
query-19	Is obj-arm2 part of obj-student3?	part_of	F	H	F	1.00	11.10
query-20	Is obj-head2 part of obj-student2?	part_of	T	H	F	1.00	10.71
query-21	Is obj-head2 part of obj-student3?	part_of	F	H	F	1.00	15.26
query-22	Are obj-head2 and obj-arm2 part of the same person?	part_of	T	H	F	0.91	10.17
query-23	Is obj-student2 a female?	classification	T	H	T	0.22	12.44
query-24	Is obj-student1 a female?	classification	F	H	F	1.00	10.10
query-25	Is obj-student3 a female?	classification	T	H	T	0.22	8.79
query-26	Identify obj-student4 in obs-IP5.	object_definition	T	H	T	0.44	13.05
query-27	Identify obj-head3 in obs-IP5.	object_definition	T	H	T	0.44	13.32
query-28	Identify obj-instructor2 in obs-IP5.	object_definition	T	H	T	0.44	14.47
query-29	Identify obj-head4 in obs-IP5.	object_definition	T	H	T	0.44	13.66
query-30	Is obj-head4 part of obj-student4?	part_of	F	H	F	1.00	10.14

query-31	Is obj-head4 part of obj-instructor2?	part_of	T	H	T	0.44	6.90
query-32	Is obj-head3 part of obj-student4?	part_of	T	H	T	0.44	6.92
query-33	Is obj-head3 part of obj-instructor2?	part_of	F	H	F	1.00	9.80
query-34	Identify obj-person1 in obs-IP2.	object_definition	T	H	T	0.95	7.65
query-35	Identify obj-person2 in obs-IP2.	object_definition	T	H	T	0.95	7.52
query-36	Identify obj-arm3 in obs-IP2.	object_definition	T	H	T	0.95	10.36
query-37	Identify obj-arm4 in obs-IP2.	object_definition	T	H	T	0.95	104.73
query-38	Is obj-arm4 part of obj-person2?	part_of	F	H	T	0.95	6.91
query-39	Is obj-arm3 part of obj-person2?	part_of	T	H	T	0.95	5.66
query-40	Is obj-arm4 part of obj-person1?	part_of	T	H	F	1.00	7.80
query-41	Identify obj-person3 in obs-IP2.	object_definition	T	H	T	0.69	9.04
query-42	Identify obj-person4 in obs-IP2.	object_definition	T	H	T	0.69	30.86
query-43	Identify obj-hand2 in obs-IP2.	object_definition	T	H	UnknownObject		485.30
query-44	Is obj-hand2 part of obj-person3?	part_of	F	H	Skipped		0.01
query-45	Is obj-hand2 part of obj-person4?	part_of	T	H	Skipped		0.01
query-46	Identify obj-lower-body2 in obs-HC3.	object_definition	T	H	T	0.26	949.01
query-47	Identify obj-person5 in obs-HC3.	object_definition	T	H	T	0.22	37.77
query-48	Identify obj-lower-body3 in obs-HC3.	object_definition	T	H	T	0.26	48.21
query-49	Is obj-lower-body3 part of obj-person5?	part_of	F	H	F	1.00	12.81
query-50	Is obj-lower-body2 part of obj-person5?	part_of	T	H	F	1.00	10.66
storyline-registration							
					ucla		
	Query	**Category**	**Assessor**				**Time**

query-1	Identify person obj-reg1.	object_definition	T	H	T	0.24	17.00
query-2	Is obj-reg1 stationary?	tracking	T	H	T	0.35	8.95
query-3	Is obj-reg1 moving?	tracking	T	H	T	0.35	9.20
query-4	Is obj-reg1 standing?	attributes	T	H	T	0.35	14.47
query-5	Is obj-reg1 walking?	attributes	T	H	T	0.35	10.25
query-6	Is obj-reg1 pointing?	attributes	F	H	Other		7.33
query-7	Is obj-reg1 sitting?	attributes	T	H	F	0.91	10.55
query-8	Is obj-reg1 talking?	attributes	T	H	Other		5.41
query-9	Is obj-reg1 running?	attributes	F	H	T	0.35	10.50
query-10	Is obj-reg1 crawling?	attributes	F	H	F	0.91	7.58
query-11	Identify person obj-student1.	object_definition	T	H	T	0.38	77.30
query-12	Is obj-student1 stationary?	tracking	T	H	T	0.40	9.02
query-13	Is obj-student1 moving?	tracking	T	H	T	0.40	7.23
query-14	Is obj-student1 standing?	attributes	T	H	T	0.40	10.57
query-15	Is obj-student1 walking?	attributes	T	H	T	0.40	8.65
query-16	Is obj-student1 pointing?	attributes	F	H	Other		4.35
query-17	Is obj-student1 sitting?	attributes	F	H	F	0.91	14.55
query-18	Is obj-student1 talking?	attributes	T	H	Other		5.89
query-19	Is obj-student1 running?	attributes	F	H	F	0.91	7.92
query-20	Is obj-student1 crawling?	attributes	F	H	F	0.91	6.84
query-21	Identify bag obj-luggage.	object_definition	T	H	T	0.38	14.09
query-22	Is obj-student1 touching obj-luggage?	relationships	T	H	T	0.40	7.26
query-23	Is obj-student1 putting-down obj-luggage?	relationships	F	H	F	0.91	8.28
query-24	Is obj-student1 throwing obj-luggage?	relationships	F	H	F	0.91	35.38
query-25	Is obj-student1 carrying obj-luggage?	relationships	T	H	F	0.91	11.97
query-26	Is obj-student1 picking-up obj-	relationships	F	H	F	0.91	19.82

	luggage?						
query-27	Is obj-student1 dropping obj-luggage?	relationships	F	H	Other		18.19
query-28	Is obj-student1 catching obj-luggage?	relationships	F	H	Other		7.38
query-29	Identify person obj-student2.	object_definition	T	H	T	0.38	18.47
query-30	Is obj-student1 same-object obj-student2?	relationships	F	H	T	0.40	7.77
query-31	Is obj-student1 same-motion obj-student2?	tracking	T	H	F	0.91	8.83
query-32	Is obj-student1 together obj-student2?	relationships	T	H	F	0.91	12.62
query-33	Is obj-student1 below obj-student2?	relationships	F	H	F	0.91	9.82
query-34	Is obj-student1 touching obj-student2?	relationships	F	H	F	0.91	8.20
query-35	Is obj-student1 following obj-student2?	tracking	F	M	F	0.91	9.63
query-36	Is obj-student1 opposite-motion obj-student2?	tracking	F	H	F	0.91	8.74
query-37	Is obj-student1 passing obj-student2?	tracking	F	H	F	0.91	10.57
query-38	Identify person obj-student3.	object_definition	T	H	T	0.29	16.66
query-39	Identify backpack, obj-backpack.	object_definition	T	H	UnknownObject		416.48
query-40	Is obj-student3 touching obj-backpack?	relationships	T	H	Skipped		0.01
query-41	Is obj-student3 on obj-backpack?	relationships	F	H	Skipped		0.04
query-42	Is obj-student3 putting-down obj-backpack?	relationships	T	H	Skipped		0.01
query-43	Is obj-student3 throwing obj-backpack?	relationships	F	H	Skipped		0.01
query-44	Is obj-student3 dropping obj-backpack?	relationships	F	H	Skipped		0.04
query-45	Is obj-student3 picking-up obj-	relationships	T	H	Skipped		0.01

	backpack?						
query-46	Is obj-student3 carrying obj-backpack?	relationships	T	H	Skipped		0.05
query-47	Identify person obj-student4.	object_definition	T	H	T	0.38	17.10
query-48	Is obj-student4 moving?	tracking	T	H	T	0.38	7.46
query-49	Is obj-student4 stationary?	tracking	T	H	T	0.38	9.71
query-50	Is obj-student4 standing?	attributes	T	H	T	0.38	7.93
query-51	Is obj-student4 walking?	attributes	T	H	T	0.38	11.45
query-52	Is obj-student4 pointing?	attributes	T	H	Other		5.64
query-53	Is obj-student4 sitting?	attributes	F	H	F	0.91	9.02
query-54	Is obj-student4 starting?	tracking	T	H	T	0.38	7.09
query-55	Is obj-student4 talking?	attributes	T	H	Other		10.35
query-56	Is obj-student4 running?	attributes	F	H	F	0.91	8.15
query-57	Is obj-student4 crawling?	attributes	F	H	F	0.91	7.68
query-58	Is obj-student4 stopping?	tracking	T	H	T	0.38	6.53
query-59	Is obj-student3 same-object obj-student4?	relationships	F	H	F	0.91	8.65
query-60	Is obj-student3 together obj-student4?	relationships	T	H	F	0.91	8.89
query-61	Is obj-student3 clear-line-of-sight obj-student4?	spatial	T	H	F	0.91	10.12
query-62	Is obj-student3 facing obj-student4?	spatial	F	H	F	0.91	13.96
query-63	Is obj-student3 touching obj-student4?	relationships	F	H	F	0.91	12.96
query-64	Is obj-student3 facing-opposite obj-student4?	spatial	F	H	F	0.91	7.46
query-65	Identify table as obj-table.	object_definition	T	H	T	0.29	18.08
query-66	Identify water bottle as obj-water-bottle.	object_definition	T	H	T	0.29	42.12

query-67	Identify person as obj-student5.	object_definition	T	H	T	0.29	97.02
query-68	Is obj-water-bottle below obj-table?	relationships	F	H	F	0.91	8.68
query-69	Is obj-water-bottle touching obj-table?	relationships	T	H	T	0.38	12.51
query-70	Is obj-water-bottle on obj-table?	relationships	T	H	F	0.91	7.79
query-71	Is obj-student5 putting-down obj-water-bottle?	relationships	F	H	F	0.91	8.21
query-72	Is obj-student5 throwing obj-water-bottle?	relationships	F	H	F	0.91	8.98
query-73	Is obj-student5 dropping obj-water-bottle?	relationships	F	H	Other		14.35
query-74	Is obj-student5 touching obj-water-bottle?	relationships	T	H	F	0.91	9.38
query-75	Is obj-student5 picking-up obj-water-bottle?	relationships	T	H	F	0.91	8.44
query-76	Is obj-student5 carrying obj-water-bottle?	relationships	T	H	F	0.91	14.46
storyline-presentation							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Identify auditorium(room) as obj-auditorium.	object_definition	T	H	T	0.95	11.55
query-2	Are there at least 5 people who enter the auditorium during time-enter?	relationships	T	H	Other		1749.09
query-3	Are there at least 5 people who exit the auditorium during time-enter?	relationships	F	H	Other		7.07
query-4	Identify person as obj-student3.	object_definition	T	H	T	0.49	41.87
query-5	Is obj-student3 outside obj-auditorium?	relationships	T	H	F	0.91	9.31
query-6	Is obj-student3	tracking	F	H	F	0.91	8.98

	stationary?						
query-7	Is obj-student3 moving?	tracking	T	H	F	0.91	8.15
query-8	Is obj-student3 walking?	attributes	T	H	T	0.49	5.24
query-9	Is obj-student3 pointing?	attributes	F	H	Other		3.99
query-10	Is obj-student3 running?	attributes	F	H	F	0.91	5.96
query-11	Is obj-student3 stopping?	tracking	F	H	F	0.91	23.32
query-12	Identify person as obj-student1.	object_definition	T	H	T	0.44	70.35
query-13	Is obj-student1 inside obj-auditorium?	relationships	T	H	T	0.44	6.59
query-14	Is obj-student1 moving?	tracking	T	H	F	0.91	9.49
query-15	Is obj-student1 stationary?	tracking	T	H	F	0.91	8.41
query-16	Is obj-student1 walking?	attributes	T	H	T	0.44	7.26
query-17	Is obj-student1 turning-left?	tracking	T	H	F	0.91	7.50
query-18	Is obj-student1 pointing?	attributes	F	H	Other		49.08
query-19	Is obj-student1 sitting?	attributes	T	H	F	0.91	8.35
query-20	Is obj-student1 starting?	tracking	F	H	F	0.91	7.62
query-21	Is obj-student1 turning-right?	tracking	F	H	F	0.91	26.41
query-22	Is obj-student1 running?	attributes	F	H	F	0.91	7.36
query-23	Is obj-student1 crawling?	attributes	F	H	F	0.91	7.43
query-24	Identify bag as obj-bag.	object_definition	T	H	UnknownObject		396.67
query-25	Is obj-student1 touching obj-bag?	relationships	T	H	Skipped		0.01
query-26	Is obj-student1 carrying obj-bag?	relationships	T	H	Skipped		0.01
query-27	Is obj-student1 putting-down obj-bag?	relationships	T	H	Skipped		0.01
query-28	Is obj-student1 throwing obj-bag?	relationships	F	H	Skipped		0.01

query-29	Is obj-student1 picking-up obj-bag?	relationships	F	H	Skipped		0.01
query-30	Is obj-student1 dropping obj-bag?	relationships	F	H	Skipped		0.01
query-31	Is obj-student1 catching obj-bag?	relationships	F	H	Skipped		0.01
query-32	Identify person as obj-presenter1.	object_definition	T	H	T	0.95	919.70
query-33	Identify person as obj-presenter2.	object_definition	T	H	T	0.69	7.79
query-34	Identify person as obj-student2.	object_definition	T	H	T	0.69	8.26
query-35	Does obj-student2 have a clear line of sight to obj-presenter1?	spatial	T	H	T	0.69	5.41
query-36	Does obj-student2 have a clear line of sight to obj-presenter2?	spatial	T	H	T	0.69	5.29
query-37	Is obj-presenter1 closer to obj-student2 than obj-presenter2?	spatial	F	H	T	0.69	7.04
query-38	Is obj-presenter1 same-object obj-student2?	relationships	F	H	F	0.91	12.66
query-39	Is obj-presenter1 same-motion obj-student2?	tracking	T	H	Other		46.90
query-40	Is obj-presenter1 facing-opposite obj-student2?	spatial	T	H	F	0.91	9.93
query-41	Is obj-presenter1 facing obj-student2?	spatial	T	H	F	0.91	9.64
query-42	Is obj-presenter1 touching obj-student2?	relationships	F	H	F	0.91	7.52
query-43	Is obj-presenter1 following obj-student2?	tracking	T	H	Other		74.56
query-44	Is obj-presenter1 passing obj-student2?	tracking	F	H	Other		11.64
storyline-presentation2							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Identify person as obj-	object_definition	T	H	T	0.95	27.66

	student1.						
query-2	Is obj-student1 standing?	attributes	F	H	F	0.91	8.21
query-3	Is obj-student1 walking?	attributes	F	H	F	0.91	18.07
query-4	Is obj-student1 pointing?	attributes	F	H	Other		5.33
query-5	Is obj-student1 sitting?	attributes	T	H	T	0.44	6.47
query-6	Is obj-student1 stationary?	tracking	T	H	T	0.44	6.84
query-7	Identify person as obj-presenter2.	object_definition	T	H	T	0.44	11.40
query-8	Identify hat as obj-hat.	object_definition	T	H	UnknownObject		233.72
query-9	Is obj-presenter2 donning obj-hat?	relationships	F	H	Skipped		0.01
query-10	Is obj-presenter2 doffing obj-hat?	relationships	F	H	Skipped		0.01
query-11	Is obj-presenter2 wearing obj-hat?	relationships	F	H	Skipped		0.01
query-12	Is obj-presenter2 putting-down obj-hat?	relationships	F	H	Skipped		0.01
query-13	Is obj-presenter2 throwing obj-hat?	relationships	F	H	Skipped		0.01
query-14	Is obj-presenter2 touching obj-hat?	relationships	T	H	Skipped		0.01
query-15	Is obj-presenter2 picking-up obj-hat?	relationships	F	H	Skipped		0.01
query-16	Identify person as obj-student2.	object_definition	T	H	T	0.95	1325.21
query-17	Is there a clear line of sight from student2 to presenter2?	spatial	T	H	F	0.91	9.27
query-18	Is there a clear line of sight from obs-MC1 to obj-student1?	spatial	F	H	F	0.91	10.74
query-19	Is there a clear line of sight from obs-MC1 to obj-hat?	spatial	F	H	Skipped		0.01
query-20	From the point of view of MC1, is presenter2 occluding obj-hat?	spatial	T	H	Skipped		0.01
query-21	Is obj-presenter2 standing?	attributes	T	H	F	0.91	9.34
query-22	Is obj-presenter2	tracking	T	H	F	0.91	7.29

	moving?						
query-23	Is obj-presenter2 walking?	attributes	T	H	F	0.91	10.99
query-24	Is obj-presenter2 running?	attributes	F	H	F	0.91	8.88
query-25	Is obj-presenter2 sitting?	attributes	T	H	F	0.91	7.79
query-26	Is obj-student1 closer to obj-presenter2 than obj-student2?	spatial	T	H	F	0.91	9.90
query-27	Is obj-student1 farther from obj-student2 than obj-presenter2?	spatial	F	H	F	0.91	11.77
query-28	Are there at least 4 people in the location loc-front-auditorium?	classification	T	H	T	0.95	6.11
query-29	Identify person as obj-presenter1.	object_definition	T	H	UnknownObject		364.94
query-30	Identify person as obj-student3.	object_definition	T	H	T	0.95	8.44
query-31	Identify person as obj-student4.	object_definition	T	H	T	0.27	801.39
query-32	Is obj-presenter1 touching obj-student3?	relationships	F	H	Skipped		0.01
query-33	Is obj-presenter1 together with obj-student3?	relationships	T	H	Skipped		0.02
query-34	Is obj-presenter1 facing-opposite obj-student3?	spatial	T	H	Skipped		0.01
query-35	Is obj-student4 touching obj-student3?	relationships	F	H	F	0.91	201.04
query-36	Is obj-student4 together with obj-student3?	relationships	F	H	F	0.91	7.36
query-37	Is obj-student4 facing-opposite obj-student3?	spatial	F	H	F	0.91	6.34
query-38	Is obj-student3 stationary?	tracking	T	H	T	0.95	6.71
query-39	Is obj-student3 standing?	attributes	T	H	T	0.95	3.30
query-40	Is obj-student3 talking?	attributes	T	H	F	0.91	7.02
query-41	Is obj-student3 pointing?	attributes	F	H	Other		1068.71

query-42	Is obj-student3 sitting?	attributes	F	H	F	0.91	73.26
query-43	Does obj-presenter1 have a clear line of sight to obj-student3?	spatial	T	H	Skipped		0.01
query-44	Does obj-presenter1 have a clear line of sight to obj-student4?	spatial	F	H	Skipped		0.01
query-45	Is obj-presenter1 standing?	attributes	T	H	Skipped		0.01
query-46	Is obj-presenter1 talking?	attributes	F	H	Skipped		0.01
query-47	Is obj-presenter1 walking?	attributes	F	H	Skipped		0.01
query-48	Is obj-presenter1 sitting?	attributes	F	H	Skipped		0.01
query-49	Is obj-student4 standing?	attributes	F	H	T	0.38	6.92
query-50	Is obj-student4 talking?	attributes	F	H	Other		298.78
query-51	Is obj-student4 walking?	attributes	T	H	F	0.93	14.94
query-52	Is obj-student4 sitting?	attributes	F	H	F	0.93	8.39
storyline-panic							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Identify interior door as obj-door1.	object_definition	T	H	T	0.22	12.34
query-2	Identify second interior door as obj-door2.	object_definition	T	H	UnknownObject		141.88
query-3	Is obj-door1 open?	attributes	F	H	F	0.91	11.74
query-4	Is obj-door2 open?	attributes	T	H	Skipped		0.01
query-5	Identify person as obj-person1.	object_definition	T	H	T	0.22	70.86
query-6	Identify person as obj-person2.	object_definition	T	H	T	0.22	25.67
query-7	Is obj-person1 same-object obj-person2?	relationships	F	H	F	0.91	9.12
query-8	Is obj-person1 same-motion obj-person2?	tracking	T	H	Other		13.43
query-9	Is obj-person1 together obj-person2?	relationships	T	H	F	0.91	10.81
query-10	Is obj-person1 facing-	spatial	F	H	F	0.91	5.21

	opposite obj-person2?						
query-11	Is obj-person1 touching obj-person2?	relationships	F	H	F	0.91	7.46
query-12	Is obj-person1 following obj-person2?	tracking	F	H	Other		11.77
query-13	Is obj-person1 opposite-motion obj-person2?	tracking	F	H	Other		9.13
query-14	Is obj-person1 passing obj-person2?	tracking	F	H	Other		9.51
query-15	Identify person as obj-person3.	object_definition	T	H	T	0.31	68.88
query-16	Identify luggage as obj-bag.	object_definition	T	H	UnknownObject		316.39
query-17	Is obj-person3 touching obj-bag?	relationships	T	H	Skipped		0.01
query-18	Is obj-person3 carrying obj-bag?	relationships	T	H	Skipped		0.01
query-19	Is obj-person3 dropping obj-bag?	relationships	F	H	Skipped		0.01
query-20	Is obj-person3 putting-down obj-bag?	relationships	F	H	Skipped		0.01
query-21	Is obj-person3 picking-up obj-bag?	relationships	F	H	Skipped		0.01
storyline-scene-locations							
					ucla		
	Query	**Category**	**Assessor**				**Time**
query-1	Is there at least one person in the cartesian polygon loc-lobby at time-1?	classification	F	H	T	0.40	8.24
query-2	Is there at least one person in the cartesian polygon loc-hallway at time-1?	classification	T	H	T	0.95	23.17
query-3	Is there at least one person in the cartesian polygon loc-entry at time-1?	classification	F	H	T	0.49	12.37

IV. Publications funded under this program

- Jake Bouvrie, Mauro Maggioni, "Multiscale Markov Decision Problems: Compression, Solution, and Transfer Learning", submitted (arXiv 1212.1143 [cs.AI]).
- William K. Allard, Guangliang Chen, Mauro Maggioni, "Multiscale Geometric Methods for Data Sets II: Geometric Multi-Resolution Analysis", Appl. Comp. Harm. Anal., Vol. 32(3), May 2012, 435-462.
- J. Bouvrie and M. Maggioni, "Geometric Multiscale Reduction for Autonomous and Controlled Nonlinear Systems", IEEE Conference on Decision and Control (CDC), 2012.
- M. Iwen, M. Maggioni, "Approximation of Points on Low-Dimensional Manifolds Via Random Linear Projections", Information & Inference, Feb. 2013.
- G. Chen and M. Iwen and S. Chin and M. Maggioni, "A Fast Multiscale Framework for Data in High Dimensions: Measure Estimation, Anomaly Detection, and Compressive Measurements", Visual Communications and Image Processing (VCIP), Nov. 2012 IEEE.
- E. Hall and R. Willett. "Foreground and background reconstruction in Poisson video, ICIP 2013.
- A. K. Oh, Z. T. Harmany, and R. Willett. "Logarithmic total variation regularization for cross-validation in photon-limited imaging, ICIP 2013.
- E. Hall and R. Willett. "Dynamical Models and Tracking Regret in Online Convex Programming, ICML 2013, arXiv:1301.1254, 2013.
- Y. Xie, J. Huang, and R. Willett. "Changepoint detection for high-dimensional time series with missing data. arXiv:1208:5062, 2012.
- Banerjee A, Murray J & Dunson D (2013). Bayesian learning of joint distributions of objects. *International Conference on Artificial Intelligence & Statistics*
- Bhattacharya A, Pati D & Dunson D (2013). Adaptive dimension reduction with a Gaussian process prior. *Annals of Statistics*, to appear.
- Bhattacharya A, Pati D & Dunson D (2013). Posterior convergence rates in non-linear latent variable models. *Electronic Journal of Statistics*, to appear.
- Bhattacharya A, Pati D, Pillai N & Dunson, D (2013). Bayesian shrinkage. *arXiv:1212.6088*.
- Guhaniyogi R & Dunson D (2013). Bayesian compressed regression. *arXiv:1303.0642*.
- Pati D, Bhattacharya A, Pillai N & Dunson D (2013). Posterior contraction in sparse Bayesian factor models for massive covariance matrices. *Annals of Statistics*, under revision.
- Strawn N, Armagan A, Saab R, Carin L & Dunson D (2013). Finite sample posterior concentration in high-dimensional regression. *arXiv:1207.4854*.

- Yang Y & Dunson D (2013). Bayesian manifold regression. *arXiv:1305.0617*.
- Zhou J, Bhattacharya A, Herring H & Dunson D (2013). Bayesian factorizations of big sparse tensors. *arXiv:1306.1598*.
- L. Ren, Y. Wang, D. Dunson and L. Carin, "The kernel beta process, Neural Information Processing Systems (NIPS), 2011.
- Z. Xing, M. Zhou, A. Castrodad, G. Sapiro and L. Carin, "Dictionary learning for noisy and incomplete hyperspectral images, SIAM Journal on Imaging Sciences, 2011.
- E. Wang, E. Salazar, D. Dunson and L. Carin, "Spatio-Temporal Modeling of Legislation and Votes, to appear in Bayesian Analysis
- M. Ding, L. He, D. Dunson and L. Carin, "Nonparametric Bayesian Segmentation of Multivariate Inhomogeneous Space-Time Poisson Process, to appear in Bayesian Analysis
- M. Zhou, L. Hannah, D. Dunson and L. Carin, "Beta-negative binomial process and Poisson factor analysis, Artificial Intelligence and Statistics (AISTATS), 2012
- B. Chen, G. Polatkan, G. Sapiro, D. Blei, D. Dunson and L. Carin, "Deep learning with hierarchical convolutional factor analysis, to appear in IEEE Trans. Pattern Analysis Machine Intelligence
- Y. Wang and L. Carin, "Levy Measure Decompositions for the Beta and Gamma Processes, Proc. Int. Conf. Machine Learning (ICML), 2012
- E. Salazar and L. Carin, "Inferring Latent Structure From Mixed Real and Categorical Relational Data, Proc. Int. Conf. Machine Learning (ICML), 2012
- M. Chen, W. Carson, M. Rodrigues, R. Calderbank and L. Carin, "Communications Inspired Linear Discriminant Analysis, Proc. Int. Conf. Machine Learning (ICML), 2012
- S. Han, X. Liao and L. Carin, "Cross-Domain Multitask Learning with Latent Probit Models, Proc. Int. Conf. Machine Learning (ICML), 2012
- M. Zhou and L. Carin, "Lognormal and Gamma Mixed Negative Binomial Regression, Proc. Int. Conf. Machine Learning (ICML), 2012
- J. Silva and L. Carin, "Active Learning for Online Bayesian Matrix Factorization, Proc. SIGKDD Conf. Knowledge Discovery and Data Mining, 2012.
- X. Chen, M. Zhou and L. Carin, "The Contextual Focused Topic Model, Proc. SIGKDD Conf. Knowledge Discovery and Data Mining, 2012.
- L. Li, X. Zhang, M. Zhou and L. Carin, "Nested Dictionary Learning for Hierarchical Organization of Imagery and Text, Proc. Uncertainty in Artificial Intelligence (UAI), 2012
- M. Zhou and L. Carin, "Augment-and-Conquer Negative Binomial Processes, Proc. Neural and

Information Processing Systems (NIPS), 2012

- X. Zhang and L. Carin, “Joint Modeling of a Matrix with Associated Text, Proc. Neural and Information Processing Systems (NIPS), 2012
- J.M. Duarte-Carvajalino, G. Yu, L. Carin and G. Sapiro, “Task-driven adaptive statistical compressive sensing of Gaussian mixture models, to appear in IEEE Trans. Signal Processing

V. Patents

None

VI. PhD Students Graduating

Eric Wang, Mingyuan Zhou, John Paisley, Minhua Lu

REFERENCES

[1] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *JMLR*, 2003.

[2] M. Zhou, L. Hannah, D. Dunson, and L. Carin, "Beta-negative Binomial process and Poisson factor analysis," *AISTATS*, 2012.

[3] S. Williamson, C. Wang, K. Heller, and D. M. Blei, "The IBP compound Dirichlet process and its application to focused topic modeling," *ICML*, 2010.

[4] D. M. Blei, T. Griffiths, M. I. Jordan, and J. B. Tenenbaum, "Hierarchical topic models and the nested Chinese restaurant process," *NIPS*, 2004.

[5] J. Paisley, C. Wang, D. M. Blei, and M. I. Jordan, "Nested hierarchical Dirichlet processes," *PAMI*, 2015.

[6] R. Salakhutdinov and G. E. Hinton, "Replicated softmax: an undirected topic model," *NIPS*, 2009.

[7] G. E. Hinton, "Training products of experts by minimizing contrastive divergence," *Neural computation*, 2002.

[8] H. Larochelle and S. Lauly, "A neural autoregressive topic model," *NIPS*, 2012.

[9] G. E. Hinton, S. Osindero, and Y. W. Teh, "A fast learning algorithm for deep Belief nets," *Neural computation*, 2006.

[10] R. Salakhutdinov and G. E. Hinton, "Deep Boltzmann machines," *AISTATS*, 2009.

[11] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," *ICLR*, 2014.

[12] A. Mnih and K. Gregor, "Neural variational inference and learning in belief networks," *ICML*, 2014.

[13] J. R. Danilo, M. Shakir, and W. Daan, "Stochastic backpropagation and approximate inference in deep generative models," *ICML*, 2014.

[14] Z. Gan, R. Henao, D. Carlson, and L. Carin, "Learning deep sigmoid belief networks with data augmentation," *AISTATS*, 2015.

[15] G. E. Hinton and R. Salakhutdinov, "Discovering binary codes for documents by learning deep generative models," *Topics in Cognitive Science*, 2011.

[16] N. Srivastava, R. Salakhutdinov, and G. E. Hinton, "Modeling documents with deep Boltzmann machines," *UAI*, 2013.

[17] R. M. Neal, "Connectionist learning of belief networks," *Artificial intelligence*, 1992.

[18] R. Guhaniyogi, S. Qamar, and D. B. Dunson, "Bayesian conditional density filtering," *arXiv:1401.3632*, 2014.

[19] N. Ding, Y. Fang, R. Babbush, C. Chen, R. D. Skeel, and H. Neven, "Bayesian sampling using stochastic gradient thermostats," *NIPS*, 2014.

[20] M. Zhou and L. Carin, "Negative Binomial process count and mixture modeling," *PAMI*, 2015.

[21] Q. Ho, J. Cipar, H. Cui, J. K. Kim, S. Lee, P. B. Gibbons, G. A. Gibbons, G. R. Ganger, and E. P. Xing, "More effective distributed ml via a stale synchronous parallel parameter server," *NIPS*, 2013.

[22] M. Li, D. Andersen, A. Smola, and K. Yu, "Communication efficient distributed machine learning with the parameter server," *NIPS*, 2014.

[23] N. G. Polson, J. G. Scott, and J. Windle, "Bayesian inference for logistic models using pólya-Gamma latent variables," *JASA*, 2013.

[24] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient Langevin dynamics," *ICML*, 2011.

[25] T. Chen, E. Fox, and C. Guestrin, "Stochastic gradient hamiltonian monte carlo," *ICML*, 2014.

[26] M. E. Tuckerman, *Statistical Mechanics: Theory and Molecular Simulation*. Oxford University Press, 2010.

[27] S. Patterson and Y. W. Teh, "Stochastic gradient Riemannian Langevin dynamics on the probability simplex," *NIPS*, 2013.

[28] A. Yuille, "The convergence of contrastive divergences," *NIPS*, 2005.

[29] D. M. Blei and J. D. Lafferty, "A correlated topic model of science," *The Annals of Applied Statistics*, 2007.

[30] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Hierarchical dirichlet processes," *JASA*, 2006.

[31] L. Maaloe, M. Arngren, and O. Winther, "Deep belief nets for topic modeling," *arXiv:1501.04325*, 2015.

[32] R. Ranganath, L. Tang, L. Charlin, and D. M. Blei, "Deep exponential families," *arXiv:1411.2581*, 2014.

[33] M. D. Hoffman, D. M. Blei, and F. Bach, "Online learning for latent Dirichlet allocation," *NIPS*, 2010.

[34] D. Mimno, M. D. Hoffman, and D. M. Blei, "Sparse stochastic inference for latent dirichlet allocation," *ICML*, 2012.

[35] C. Wang and D. M. Blei, "Truncation-free stochastic variational inference for bayesian nonparametric models," *NIPS*, 2012.

[36] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, "Stochastic variational inference," *JMLR*, 2013.

[37] M. Girolami and B. Calderhead, "Riemann manifold langevin and hamiltonian monte carlo methods," *J. R. Statist. Soc. B*, 2011.

[38] S. Byrne and M. Girolami, "Geodesic monte carlo on embedded manifolds," *Scandinavian J. Statist*, 2013.