Intelligence virtual analyst capability

Governing concepts and science and technology roadmap

Alexandre Bergeron-Guyard
Valérie Lavigne
Denis Poussart
Denis Gouin
Jean Roy

DRDC – Valcartier Research Centre

Defence Research and Development Canada

Scientific Report
DRDC-RDDC-2014-R156
December 2014

IMPORTANT INFORMATIVE STATEMENTS

TIF 05dz13

Abstract

This report presents the results of an investigation of Intelligent Software Assistant (ISA) technologies towards the development of an Intelligence Virtual Analyst Capability (iVAC). The idea behind the ISA is to synthesize the current state of Artificial Intelligence research to develop a personalized assistant that organizes information, learns processes, adapts to changing situations, and interactively supports individuals in their tasks in a seamless, intuitive fashion. The iVAC is meant as a computerized software assistant supporting the intelligence analysts in sensemaking tasks, while being ultimately capable of taking on autonomous analytical tasks in concert with other analysts (virtual or human).

To achieve this goal, literature surveys, technology studies, expert workshops and detailed analysis of iVAC sub-capabilities were performed. For every identified iVAC sub-capability, an assessment of the maturity of existing, relevant, applicable science and technology was performed. As a result, key research topics were identified. This document provides a description of the iVAC concept and of the research and development steps required to instantiate it.

Significance to defence and security

This effort lays the ground work and provides a way ahead for the development of an Intelligence Virtual Analyst Capability (iVAC). Following the concepts and research avenues proposed in this document will enable the development of an Intelligent Software Assistant (ISA) that will support Canadian Armed Forces analysts in their collection, processing, analysis and dissemination tasks, considerably reducing information and cognitive overload.

Résumé

Ce rapport présente les résultats d'une analyse des technologies en lien avec l'*Intelligent Software Assistant* (ISA) pour le développement d'un *Intelligence Virtual Analyst Capability* (iVAC). L'idée directrice du ISA est de combiner des approches d'intelligence artificielle pour développer un assistant personnel qui organise l'information, apprend les processus, s'adapte aux situations changeantes et supporte interactivement les individus dans l'exécution de leurs fonctions. L'iVAC est un assistant logiciel capable d'aider l'analyste du renseignement dans ses fonctions d'analyse et devrait ultimement arriver à effectuer des tâches de façon autonome en concert avec d'autres analystes qu'ils soient humain(s) ou virtuel(s).

Pour atteindre cet objectif, des revues de littérature, des études de technologies, des ateliers d'experts, ainsi qu'une analyse détaillée des fonctions du iVAC ont été effectués. Pour chaque fonction de l'iVAC identifiée, une évaluation de la maturité et de l'applicabilité des composantes technologique existantes a été effectuée. Des sujets de recherche clé ont été identifiés. Ce document donne une description des concepts sous-jacents à l'iVAC et identifie les composantes nécessaires à sa réalisation.

Importance pour la défense et la sécurité

Cet effort de recherche jette les bases et propose une direction pour le développement d'un *Intelligence Virtual Analyst Capability (iVAC).* Suivre les concepts et les avenues de recherche proposés dans ce document permettra le développement d'un *Intelligent Software Assistant (ISA)* qui aidera les analystes des forces armées canadiennes dans leurs tâches de collecte, traitement, analyse et dissémination en réduisant la surcharge cognitive et informationnelle.

Table of contents

Abstract …….. i
Significance to defence and security i
Résumé …….. ii
Importance pour la défense et la sécurité ii
Table of contents iii
List of figures vi
1 Introduction 1
2 Methodology 2
2.1 Literature survey 2
2.1.1 Initial literature survey 2
2.1.2 Human computer interaction study 3
2.1.3 State of the art in machine learning paradigms 3
2.2 Analysis of science and technology components 3
2.3 Decomposition of identified iVAC sub-capabilities 4
2.4 Mapping between S&T components and identified sub-capabilities 4
3 Intelligence virtual analysis capability and future intelligence analysis capability 5
3.1 Concept 5
3.1.1 Future intelligence analysis capability 6
3.1.2 The nature of a virtual analyst and its connection with FIAC 7
4 iVAC sub-capabilities 9
4.1 Description of sub-capabilities 9
4.1.1 Manage context 10
4.1.1.1 Manage a priori knowledge 10
4.1.1.2 Receive and analyse mandate 10
4.1.1.3 Personalize for each user 11
4.1.1.4 Manage system qualities 11
4.1.2 Acquire data, information and knowledge 11
4.1.2.1 Identify sources 12
4.1.2.2 Identify information needs 12
4.1.2.3 Evaluate information 12
4.1.2.4 Summarize and organize information 12
4.1.3 Monitor, schedule, manage and evaluate activities 12
4.1.4 Learn user and task models 13
4.1.5 Support complex intelligence tasks 13
4.1.6 Interact with humans 13
4.1.7 Interact with other systems (including other iVAC) 14
4.2 Capabilities focus points 14

4.2.1 Context ... 14
4.2.1.1 Context-awareness ... 16
4.2.2 Seamless natural interaction ... 16
4.2.3 Evaluating, gaining, and maintaining trust ... 17
4.2.4 Learning ... 18
5 iVAC S&T components ... 19
5.1 Human computer interactions ... 19
5.1.1 Embodied agents representation ... 19
5.1.2 Multimodal interaction ... 20
5.1.3 Display technologies ... 21
5.1.4 Information visualization/visual analytics ... 22
5.2 Artificial intelligence ... 23
5.2.1 Question answering ... 23
5.2.2 Learning by demonstration ... 24
6 iVAC capability mapping ... 25
6.1 Manage context ... 25
6.2 Acquire data, information and knowledge ... 27
6.3 Monitor, schedule, manage and evaluate activities ... 29
6.4 Learn user and task models ... 31
6.5 Support complex intelligence tasks ... 33
6.6 Interaction with humans ... 35
6.7 Interaction with other systems ... 38
7 Conclusion ... 39
References ... 41
Annex A ISA projects and technologies ... 45
A.1 DeepQA: Question answering capability ... 45
A.2 Siri ... 45
A.3 Wolfram Alpha ... 46
A.4 Disciple-LTA ... 46
A.5 CALO ... 46
A.5.1 CALO iLink ... 47
A.5.2 CALO Express: Personal desktop assistant ... 47
A.5.3 Adept: Task learning capability ... 47
A.5.4 WARP: Workflow Activity Recognition and Proactive assistance capability ... 47
A.5.5 Task Assistant ... 48
A.5.6 Meeting Assistant ... 48
A.5.7 FOAM: Form Autocompletion Manager ... 48
A.5.8 C2RSS: News aggregator that learns ... 48
A.5.9 Clustering suite ... 48
A.5.10 Classification suite ... 48
A.5.11 PrepPak: Relevant document recommender ... 48

A.5.12 Instrumentation and automation 49
A.5.13 Contact Management 49
A.5.14 Semantic Extraction Suite 49
A.5.15 Emma: Meeting and scheduling manager 49
A.5.16 Probabilistic Consistency Engine 49
A.5.17 IRIS: Semantic desktop framework 49
A.5.18 Tagomizer: Social bookmarking capability 49
A.5.19 Open Agent Architecture 50
A.5.20 SPARK: Agent framework 50
A.6 MALLET: General text mining suite 50
A.7 MinorThird: Text annotation and semantic extraction suite 50
A.8 Zypr: Cloud-based voice controlled mashup software development framework 50
A.9 Trapit: Newswire personalization capability 51
List of symbols/abbreviations/acronyms/initialisms 53

List of figures

Figure 1 iVAC concept evolution. 5

Figure 2 High level iVAC sub-capabilities. 10

Figure 3 Percentage of records that include each scientific topic [9]. 19

Figure 4 a) A humanoid ISA avatar with a text-based dialog [25]; b) A ISA avatar embedded in the display [22]. 20

Figure 5 Picture of the emotive EPOC BCI. 20

Figure 6 Google glasses. 21

Figure 7 Sharp PN-V601 60 Inch LCD video wall. 22

Figure 8 a) Visual analytics examples: River theme visualization [31]; b) Galaxy visualization [32]. 23

Figure 9 Context management capability mapping. 26

Figure 10 Acquire data, information and knowledge capability mapping. 28

Figure 11 Monitor, schedule, manage and evaluate activities capability mapping. . . . 30

Figure 12 Learn user and task models capability mapping. 32

Figure 13 Support complex intelligence tasks capability mapping. 34

Figure 14 Interact with humans capability mapping. 36

1 Introduction

The intelligence analysts of the Canadian Armed Forces have a mandate to collect, process and analyze information, and disseminate required intelligence. The main challenge facing the analysts is not a lack of data – in some ways they are drowning in data – but rather managing and making sense of the large amount of data being presented to them. This overload problem (at the information and cognition levels) has recently been addressed using a variety of tools that allow extracting, analyzing, and reasoning on information [1][2][3][4][5]. Still there remains a strong need to support the analysts, specifically in analyzing and making sense of the processed information in order to interpret its significance, and develop new knowledge.

In order to better address this problem, it is relevant to go beyond traditional knowledge exploitation and management approaches and make use of emerging cognitive support tools. Knowledge is a mix of experience, values, contextual information, expert insight and grounded intuition that provides an environment and framework for evaluating and incorporating new experiences and information. Once made symbolic, knowledge becomes information, which is data that are endowed with meaning and purpose [6]. Data are conceived of as symbols or signs, representing stimuli or signals, that are "of no use until...in a usable (that is, relevant) form" [7]. The notion of knowledge exploitation refers to the process of using existing, or available, knowledge in order to better understand a situation or solve a problem.

A very promising paradigm in Artificial Intelligence (AI) has emerged: the Intelligent Software Assistant (ISA). The idea behind this research is to use the ISA paradigm in the intelligence context and to synthesize the current state of AI research and to develop an Intelligence Virtual Analyst Capability (iVAC). An iVAC is a virtual analyst that organizes information, learns processes, adapts to changing situations, and interactively supports the analysts in their tasks in a seamless, intuitive fashion, eventually taking on autonomous tasks in concert with other analysts (virtual or human). An iVAC should be able to learn from its experience, by interacting with and being advised by its users. It should be able to explain what it is doing and why it is doing it. An iVAC should be aware of the context, such as traits and intent of its interactive “partners” and behave accordingly. An iVAC system should “be able to reflect on what goes wrong when an anomaly occurs, and anticipate such occurrences in the future. It should be able to reconfigure itself in response to contextual changes, and should be able to be configured, maintained, and operated by non-experts.”[8].

The goal of this research is to investigate the S&T elements behind the ISA concept to generate insight towards the development of an iVAC. This report presents a way ahead to develop an iVAC capability.

- Chapter 2 discusses the methodology employed.
- Chapter 3 provides detailed information on the iVAC concept.
- Chapter 4 provides details on the iVAC sub-capabilities.
- Chapter 5 discusses primordial iVAC S&T components.
- Chapter 6 explains the mapping between S&T components and iVAC sub-capabilities.

2 Methodology

In order to perform an investigation of the elements behind the ISA concept, a number of activities were conducted:

1. an initial survey of literature,

2. a thorough analysis of key science technology components,

3. a decomposition of identified iVAC sub-capabilities, and

4. a mapping between S&T components and identified sub-capabilities.

These activities are detailed in the following sections.

2.1 Literature survey

2.1.1 Initial literature survey

The general goal of this initial literature survey was to identify R&D trends and pathways that are of interest for the development of an iVAC. More specifically, four key questions were addressed:

1. Which key research projects or commercial efforts have focused on development of an ISA (in the past 10 years)? Which of these have been developed for the military intelligence domain?

2. Which tasks and/or functionalities were addressed by the software assistants (identified in question 1)? What are the emerging trends for functionality? Which of these could specifically be appropriate for an iVAC-type application?

3. Which key technologies (hardware, software, libraries, and platforms) and scientific topics support the ISA concept? What are the emerging research trends for these supporting technologies in the context of ISA?

4. Who are the key players (academic, government, military, and commercial organizations) in efforts to develop an intelligent software assistant? How are they connected?

Several searches were performed in scientific literature and patent databases to gather sets of bibliographic data for analysis of research trends. In total, 1,759 references to scientific papers, conference presentations, technical reports, theses and dissertations were analysed. In addition, 158 patents were gathered. Statistic literature analysis involving publication velocity, trend analysis, and patent analysis was conducted.

Over 90 known research and commercial projects were identified that focus on the development of an intelligent software assistant in the past 10 years. At least 14 projects were developed or used by military agencies, eight of which have been applied to the military intelligence domain.

The main high level functionalities addressed by the research projects include decision support, interaction management, learning and task management.

The scientific topics that support the intelligent software assistant concept can be grouped into five main areas including Human Computer Interaction (HCI), AI, Knowledge Management (KM), Natural Language Processing (NLP) and a last group including disparate yet interesting topics such as automation, adaptation, decision support, context-based military applications, planning and situational awareness.

Carnegie Mellon University, George Mason University, Stanford University, IBM, Microsoft, Avaya, SRI International, and Alcatel Lucent were identified as key players in the ISA research field.

Detailed results from this preliminary analysis can be found in [9] and will be discussed in Section 5.

2.1.2 Human computer interaction study

Aspects of the iVAC project related to HCI between the virtual assistant and the users were the focus of a particular research effort. Two requirements related to iVAC HCI i.e., to explore: the need to support a dialogue between the ISA and the analyst(s), and to optimize the presentation of the results, were identified as the subject of an initial study aiming at identifying enabling technologies.

Detailed results from this research can be found in [10] and are the object of discussion in Section 4.1.6.

2.1.3 State of the art in machine learning paradigms

A State of the Art (SOTA) in Machine Learning (ML) was produced. The SOTA covers the main learning categories usually encountered in the literature, namely: supervised, unsupervised and reinforcement learning. A description of the main algorithms for each category is presented with pointers for additional information. The SOTA also includes some sections to familiarize the reader with issues such as the theoretical foundations of this field of research, as well as some good practices for conducting empirical research.

Detailed results from this work can be found in [11] and are the object of discussion in Section 5.2.

2.2 Analysis of science and technology components

This analysis of Science and Technology (S&T) components was based on the results from the initial literature survey (described in 2.1.1). Initial focus was put on Intelligent Software Assistant

(ISA) related components, and several major ISA S&T components were identified and studied. Technologies such as DeepQA of IBM, the DARPA's CALO project, Apple's Siri, Trapit, Wolfram|Alpha, Disciple-LTA, and Pioneer's Zypr were investigated. As the study progressed, multiple related technologies were identified and analysed. For instance, multiple Siri alternatives, speech technologies, Question Answering (QA) technologies, technologies for authoring and managing and doing inference on ontologies, ML and NLP technologies, as well as technologies of distributed and high performance computing were later added to the list of ISA related technologies.

Detailed results from this work can be found in [10] and are the object of discussion in Section 5.

2.3 Decomposition of identified iVAC sub-capabilities

The identification of iVAC sub-capabilities was initially based on the literature and ISA S&T reviews performed previously. A careful review of the various technologies and approaches, and the capabilities they implemented provided with a starting point for the elicitation of desired, end-state, iVAC sub-capabilities. Working from this initial sub-capability gathering, a workshop was held to refine and augment the initial sub-capability set. This workshop brought together experts from the academia, DND scientific experts, as well as military Intelligence experts. A second workshop was held, focusing capabilities more specifically related to HCI.

The results from this work are documented in [10] and are discussed in further detail in Section 5.

2.4 Mapping between S&T components and identified sub-capabilities

A mapping between iVAC sub-capabilities and Intelligent Software Assistant (ISA) technologies and projects was performed. Each sub-capability was compared against a set of S&T components. Each evaluated S&T component was rated according to the level of applicability to the given sub-capability.

The results from this work are documented in [10] and are discussed in further detail in Section 6.

3 Intelligence virtual analysis capability and future intelligence analysis capability

3.1 Concept

An iVAC is a virtual analyst that organizes information, learns processes, adapts to changing situations, and interactively supports the analysts in their tasks in a seamless, intuitive fashion, eventually taking on autonomous tasks in concert with other analysts (virtual or human). An iVAC is a system composed of tools, user and task models, that allow the iterative understanding and supporting of analysts' various tasks and roles. An iVAC should be able to learn from its experience, by interacting with and being advised by its users. It should be able to explain what it is doing and why it is doing it.

At a high level, an iVAC should be able to connect to a variety of sources, including data, knowledge and human sources. Figure 1 gives a representation of the way an iVAC could perform and evolve over time.

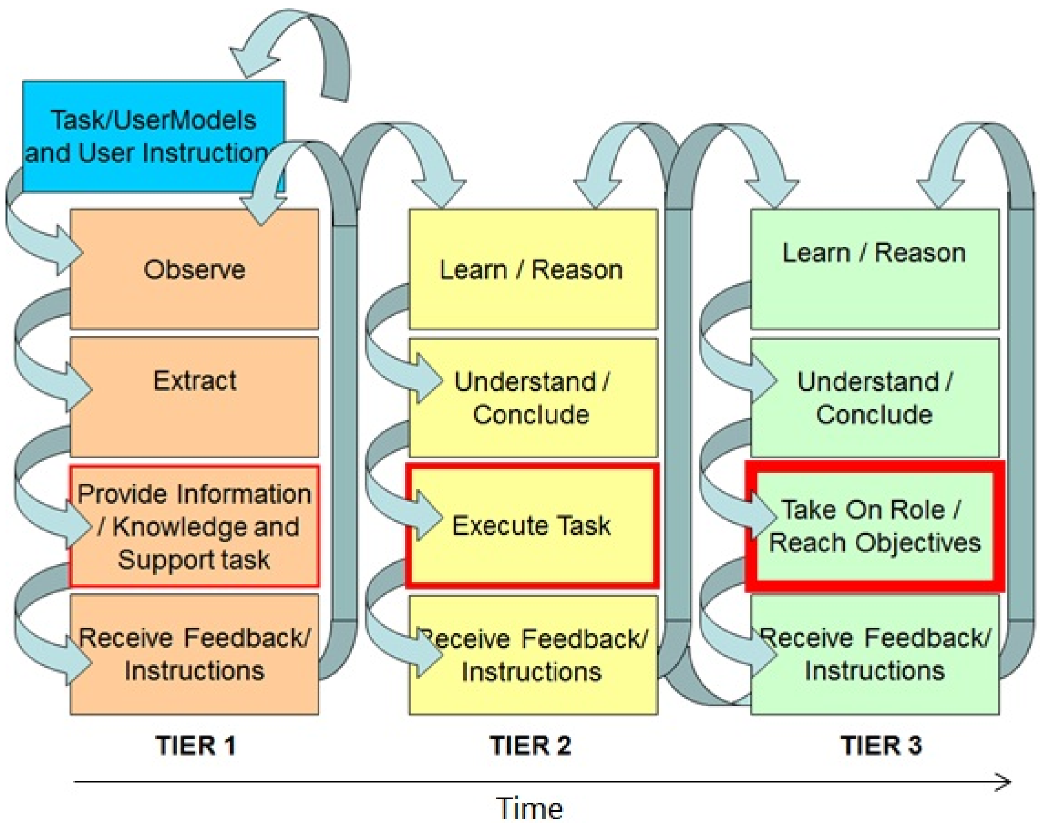

Figure 1: *iVAC concept evolution.*

An iVAC should start by acting in support to a human analyst conducting tasks of a more simple nature. At the start, the iVAC could be provided with task and user models as well as with a set of user instructions. By observing and interacting with the human analyst, the iVAC should begin to support the analyst in his task, and evolve (through feedback and instructions) towards a better understanding of the task being supported. After a certain number of iterations, the iVAC could evolve towards the actual execution of a task. In time, the iVAC should be able to conduct more complex tasks in a more autonomous way. Eventually, the iVAC may even be able to take on certain aspects of roles in order to help the human analyst reach his objectives.

The time required for an iVAC to progress from tier 1 to tier 3 is based on its performance which evaluated by the user, when providing feedback. If an iVAC were to perform extremely well from the beginning in a given task, and quickly receive positive user feedback, the time required to go from tier to tier would be short. On the other hand, if performing poorly, the iVAC would require more time to evolve from tier to tier. As shown on Figure 1, when moving to higher tiers, the iVAC takes on more responsibilities (from task support to taking on roles). The user being supported by the iVAC is likely to want to approve, or authorize, the iVAC taking on more autonomous roles. This ties in to the notion of trust, which the iVAC will have to gain through repetitive "good" performances, before an analyst agrees to delegate more of his responsibilities. In consequence, user feedback will take the form of various performance evaluations for different tasks, as well as trust evaluations, which will reflect the user's acceptance of the iVAC taking on more responsibilities. Evaluating, gaining, and maintaining trust is the topic of Section 4.2.3.

3.1.1 Future intelligence analysis capability

The Future Intelligence Analysis Capability (FIAC) vision seeks to develop advanced support toward the current and future challenges that analysts are facing in the ever increasing complexity of operating terrains.

Central issues in analysis center on two aspects:

- Data overloads: the volume of data that needs to be considered is increasing exponentially, mainly as a consequence of the widening diversity of relevant aspects that needs to be assessed, for instance in a Joint Interagency Multi-national and Public (JIMP) context. Technology has introduced a wealth of new sensors that generate huge amount of data, albeit often with a low signal/noise ratio. Cases in point include extended video recording where only a tiny fraction of the total footage may be significant, or social media information gathered from tweets, where again a small fraction of exchange may prove useful. Extracting, filtering, casting the major streams that originate from modern sensors (and sensor networks) in ways that facilitate further analysis of significant information, as well as assessing its reliability and context are major endeavors, with heavy burden on Human Analysts (HA), and
- Cognitive overloads: understanding a situation, i.e., "connecting its dots" such that high-level interpretation is possible – schemas that support a process of sense making – is again an enormous challenge. The challenge is further increased by the criticality of the conclusions or estimates/predictions that can be expressed, often within a minimal time available.

From a technical stand-point, the FIAC perspective includes the following tenets. With reference to the FIAC White Paper [8], there exists three main components that act in harmony while performing analysis:

- Neurons: the proposed capabilities are centered on the human analysts (HA) who remain in charge. They are responsible for all major decisions on the overall process and provide the key, highest level cognitive resource,
- Smart Bits: FIAC is built along a working environment where key computing resources are instantiated in an integrating software framework and platform such as a Service-Oriented Architecture (SOA). This SOA exploits a (potentially) large collection of tools – services – with each one corresponding to a specific processing task. These services are highly orthogonal, i.e., exhibit minimal functional overlap. They are implemented such as to be usable in a composed mode, with several services being "strung together" as a higher-level, meta-service. Each service is carefully validated before being put to general use, and is well documented, and
- Smart Pixels: advanced, smart interfaces – with a large component of visual element but also exploiting other sensory modalities as warranted, including haptics, gesture, sound, etc. – are seen as the bi-directional communication channel whereby the HA's and machine-based processing operate with utmost synergy. While the overall environment has a very high level of raw resources, it is through the refinement and usability of the smart pixels layer that it can be effectively exploited. In particular, it must not introduce spurious complexity which could easily arise just from the abundance of resources.

Through a careful integration of the key components above, FIAC aims at providing a unified, coherent and a cohesive information space that can be shared between HA's as they execute the various tasks of the intelligence cycle, from collation through dissemination. This integration relies on several aspects of the underlying infrastructure and design principles, from the ontology used in casting data and information in computer-based representations and common databases, up to the consistency of the metaphors and grammars associated with the sensory interfaces.

3.1.2 The nature of a virtual analyst and its connection with FIAC

The role of Virtual Analyst(s) is to mitigate some of the information and cognitive overload facing the HAs. This is to be achieved by delegating some of processing activities to machine-based resources, and doing so in ways that would decrease the spurious complexity that HAs are facing in managing their workflow and workload. Such mitigation is envisioned as a sophisticated integration of two main components:

- A robotic entity that can be requested to do specific tasks with a high level of autonomy and effectiveness, and
- An access metaphor providing bi-directional communication with the HA that provides a quantum step increase in usability.

The envisioned implementation of FIAC puts forward:

- the exploitation of a Service Oriented Architecture with clearly defined and standardized processing components, and
- the need to have outstanding usability to effectively support its "human-in-the-loop" paradigm.

These correspond, respectively, to the smart bits and smart pixels metaphors.

Clearly, the notion of virtual analyst (VA) is fully compliant with these key features. In fact, the VA consists of these two metaphors brought together to a further level of blending, where these resources are meshed in the metaphor of a fully operational assistant to whom high-level tasks may be delegated and executed with minimal supervision.

An alternate, complementary view is to consider the VA as a capability that provides for an automated sequencing of molecular services into processing with a distinctly higher-level flavor. This is in full accord with the FIAC premises that services must be defined, designed and implemented so as to allow for the composition of several of them into what might be referred to as a "macro-service". If such composition can be automatically realized, there would obviously be significant unloading of the workload of the HA, who could then concentrate more of the mental capabilities for the higher-level interpretation phases. Invoking this automated facility through an exceptionally well designed user-friendly interface – eventually that of a human avatar – would further contribute to a relief of the HA workload.

4 iVAC sub-capabilities

From a more general perspective, the iVAC refers to the notion of developing a capability in a more conceptual sense. However, such a capability is, in fact, a composition of a number of sub-capabilities. This section takes the high-level conceptual notion of the virtual analyst capability and breaks it down into its composing sub-capabilities.

The identification of iVAC sub-capabilities was based on literature and ISA S&T reviews, as well as on workshops which brought together experts from the academia, DND scientific experts, as well as military Intelligence experts. More details on this aspect can be found in Section 2.

Identifying all sub-capabilities of interest for a future iVAC system is an ambitious task, as the number of such sub-capabilities will vary with evolving users' needs and technological capabilities. Although extensive, the sub-capabilities presented in this section will therefore possibly evolve in the future.

The capabilities presented in this section describe what a full-fledged iVAC system should do, but how they could be achieved is not identified. The specific approaches relevant to implementing a given capability are the subject of Section 6.

4.1 Description of sub-capabilities

The capabilities of the iVAC system have been classified into seven broad categories, as shown in Figure 2, namely:

1. Manage context
2. Acquire data, information and knowledge
3. Monitor, schedule, manage and evaluate activities
4. Learn user and task models
5. Support complex Intelligence tasks
6. Interact with humans
7. Interact with other systems (including other iVAC)

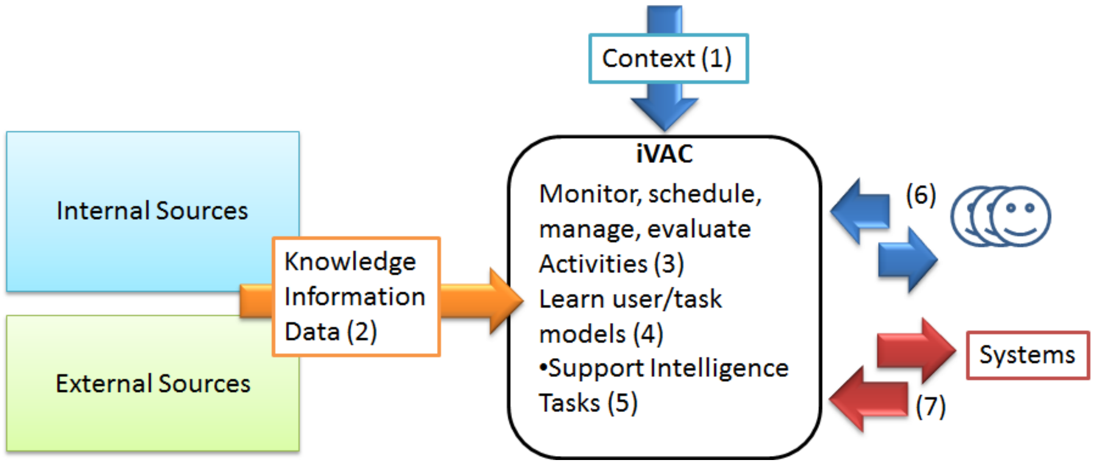

Figure 2: *High level iVAC sub-capabilities.*

4.1.1 Manage context

Context refers to the interrelated elements that make up the analyst's situation. Here the term situation refers both to the actual situation of the analyst, and to the situation that characterises the analyst's focus of interest (e.g., an analyst located in an office building in Ottawa could be analysing a situation in Afghanistan). Context awareness refers to the system's ability to react to the analyst's context. Context and context awareness are discussed in more details in Section 4.2.1.

It is obvious that, in order to provide sophisticated analysis capabilities, the iVAC must be able to manage the elements that make up the context. Such elements can include information about the current mission and objectives, as well as all other aspects that characterize the user's situation and personal profile.

4.1.1.1 Manage a priori knowledge

A priori knowledge refers to all background information and knowledge that surround the analyst's skill set. In other words, on top of knowing factual elements about a given situation, in order to perform actual relevant analysis, the analyst must know "what to do" with the facts. To be efficient, the analyst must be aware of the processes that can be applied to the information in order to extract relevant elements, he must be able to recognize patterns in the information in order to pin point aspects of the current situation that could be similar to past situations. A priori knowledge includes, but is not limited to, methods of analysis, doctrine, background knowledge, configuration parameters of algorithms used within iVAC and lessons learned. A priori knowledge also includes the information describing the user's context (detailed in Section 4.2.1).

4.1.1.2 Receive and analyse mandate

The ability to receive and understand direction from a commander is an essential component of the analyst's work. The system should be able to receive and process mission's objectives, goals, and directions from commanders. As a general goal of the system is to allow better production of intelligence products, potentially leading to actionable intelligence, this contextual information should be taken into account throughout the system when supporting the user.

4.1.1.3 Personalize for each user

The specificity of each analyst, in the way they execute tasks (interact, acquire, manage, and process information) is an intricate part of the context. The system should adapt to user's preferences. Over time, this means gaining a better grasp of the better ways to support and communicate with the analyst. The virtual analyst should provide interactive support to the analyst in a fashion that is better suited to his specificity as time goes by.

Following from 4.1.1.2, it is also important to understand and take into account the user's role in achieving mission's objectives. A virtual assistant should have a good grasp of general objectives, how to go about reaching them, and the specific contribution a user is expected to have towards reaching the objectives (what the user should achieve given his task, roles and responsibilities). Based on this contextual understanding the system should aim at combining the user's specific requirements in order to maximise the chances of reaching overall objectives.

4.1.1.4 Manage system qualities

A full-fledged iVAC system is bound to interact with a plethora of individuals and other computer systems. In line with this vision, it is likely that the efficacy with which the iVAC supports Intelligence analysis will rely on elements or resources external to itself. It is therefore required for the iVAC to track the quality of every asset that is involved in a given analysis process. Assets can refer to individuals, iVAC components, or external system components that are part of a given analysis process. Having an understanding of the quality of every asset that partakes in a given analysis will:

- allow for a better anticipation of performance in the execution of the analysis, and
- provide insight as to the internal quality of the analysis.

For instance, a given asset may be known to take very long to produce very reliable results, while another is more likely to provide instantaneous results with higher probabilities of error.

This capability quickly ties into the capability of properly fusing information. As more and more sources are discovered and used by the system, it is likely that redundant pieces of information will be encountered. It will be necessary for the iVAC to identify and deconflict these pieces of information in order to reduce potential overload to the user.

4.1.2 Acquire data, information and knowledge

As discussed earlier, an iVAC should be able to connect to a variety of sources, including data, knowledge and human sources. It should be able to collect, process and organize new information as it becomes available over the multiple sources available. The following sections provide additional details on these capabilities.

4.1.2.1 Identify sources

The system should identify new sources of information as they become available, with minimum user interactions. For example, the system could identify new sources by crawling the available networks, from emails received by the user indicating that a new web site or shared folder is now available and contains information of interest to the user. This also applies to knowledge and expertise sources. The iVAC should be able to peruse a network containing a catalog of exploitation service and recognize their potential use. The same goes for network of humans, which may direct to individuals or groups with particular expertise that may be of use in particular contexts.

4.1.2.2 Identify information needs

The system should be able to determine and identify information requirements, both from a user and a system's perspective. That is to say: what is the information the user needs to achieve his tasks and objective; and what information does the system need to effectively support the analyst in his task. These two parallel requirements are actually complimentary and will scale in importance as the iVAC progresses across the tiers of Figure 1. In other words, for a given task, as the iVAC becomes more autonomous, the information need may shift from the user to the system.

4.1.2.3 Evaluate information

The iVAC should be able to evaluate a piece of information in terms of credibility. This aspect also ties in with source discovery (4.1.2.1) as information sources should be evaluated for reliability. The iVAC should be able to assess the reliability of a source based on previous assessments of the source performed by system users. Assessment of the credibility of the information should be based on the information content as well as on what is already known (or present) in the system.

4.1.2.4 Summarize and organize information

The system should have the capability to summarize information that is found. Summaries should truly synthesize the content of information, such that users can quickly determine whether or not it is pertinent for a given activity. The capability of getting the general signification or the essence of a document could also be of use to rank information according to relevance for a given task.

4.1.3 Monitor, schedule, manage and evaluate activities

In order to gradually move from supporting a specific task to taking on a more complex role (3.1) the iVAC must be able to monitor, schedule, manage and evaluate activities. The monitoring aspect has a strong link with the system's learning phase, where it observes a task in order to learn how it is executed. Scheduling becomes relevant when tasks gain in complexity and need to be organized in a workflow. Managing comes into play when a more complex task implies the combination of numerous management or exploitation services, which require potentially complex orchestration. Evaluating activities touches the notion of performance evaluation, both from a user and a system perspective. In order to learn appropriate behaviours, the system will evaluate its own performance based on user feedback. This is first aspect of evaluation will help fine tune and improve the performance of the system. In order to provide the user with a measure of confidence with regards to information provided or task executed, the system also has to be able to evaluate its performance in various activities.

4.1.4 Learn user and task models

This capability is pivotal in order to allow the iVAC to evolve from simple, directed task support to taking on autonomous complex tasks or roles. This capability relates to both the user and the tasks, as a given task can be performed differently by various users.

From the user's perspective, the iVAC must be able to assess and learn the user's preferences on a large variety of aspects of interest as well as on presentation design preferences. The iVAC must be able to identify the information sources that are of particular interest for a user performing a specific task. The iVAC should also be able to characterise the information that is typically of use for a given task in order to be able to retrieve it from other possible sources. The key individuals, who are typical actors in a particular process or experts in a given field, must also be identified. From a task perspective, the iVAC must be able to learn from demonstration, which is to say by looking at examples of how a given task is usually performed. Learning is an important part of the iVAC construct and is the topic of Section 6.4.

4.1.5 Support complex intelligence tasks

This is obviously critical in order to achieve the full-fledged iVAC vision. As discussed in 3.1, the iVAC will evolve from simple task support to more complex task/role execution.

4.1.6 Interact with humans

It is essential, in the iVAC vision, to have a virtual assistant that interacts with users in an intuitive seamless fashion. This entails being able to effectively interact using various communication means: audio (voice), gestures, visual (video, avatar, text, graphical user interfaces), etc. There is a wide range of communication means that the iVAC must be able to understand and use. Beyond being able to use each means effectively, the iVAC must also be able to select which communication tool is the most appropriate based on the user and the situation.

The notion of interaction becomes even more complex by taking into account cultural aspects as people coming from various cultures may not react the same way to the same stimulus, or have the same requirements for presentation of information or context-related aspects (see Section 4.2.1).

Interaction with humans is very closely tied to the notion of seamless natural interactions and HCI, which are the specific topic of discussion of Sections 4.2.2 and 5.1.

4.1.7 Interact with other systems (including other iVAC)

An iVAC system will not operate as a standalone, disconnected application. To support the intelligence analyst efficiently, it has to interact with other systems (possibly other iVACs). This includes the capability of sharing data/information/knowledge with other systems as well as the ability to share information exploitation capabilities (newly developed set of rules for instance) across systems.

4.2 Capabilities focus points

A more detailed discussion on four topics that emerged as being central to the development of an iVAC is provided in this section: context definition and context awareness, seamless natural interaction, user trust, and the ability to learn.

4.2.1 Context

The Merriam Webster English Dictionary defines context as the interrelated conditions in which something exists or occurs [11]. The world of ubiquitous or pervasive computing has spawned many definitions of context. Schilit and Theimer [14] refer to context as location, identities of nearby people and objects, and changes to those objects. Dey [15] defines context as "any information that can be used to characterize the situation of an entity".

Context can be either internal or external. External (or physical) context refers to the aspects that can be directly assessed or measured by hardware sensors. Examples of external contextual elements include location, light, sound, movement, touch, temperature, air pressure, etc. Internal context (or logical) refers to elements that are specified by the user or captured by monitoring the user's interactions. Internal contextual elements include user's goals, tasks, work context, business processes, emotional state, etc.

In the context of iVAC, for the analyst, context refers to a combination of the following notions [16]:

Mission objective

The analyst may conduct work in the context of a long-term mission. This mission may include a certain number of high-level objectives that may influence the work of the analyst.

Role

The role refers to the specific function performed by the analyst in a particular process. An analyst performing a role will take on a specific set of tasks.

Task

In a given context, in a given role, the analyst will have a certain number of tasks to execute. A task may be divisible into sub-tasks. By executing a set of tasks, the analyst will fulfill a role.

Goal

Goals are the ends towards which effort is directed. In this sense, each task and sub-task could be associated to a specific goal.

Location (where?)

Location refers to the set of information that can be used to locate the analyst, and its current focus of interest. Location refers to the place where the analyst conducts its work. This includes the physical and geographical space. It also includes the location that is the focus of the current analysis. For instance an analyst could be in an analysis centre in Ottawa, working on events taking place in Afghanistan.

Identity (who?)

Identity refers to the set of information that can be used to uniquely identify the analyst. This includes basic personal information such as: full name, gender, place of birth, date of birth. It could also be extended to include system relative credentials or identifiers: social insurance number, driver's licence number, IP addresses, home address, and phone number. Finally it can also include biometric and genetic information such as face attributes, fingerprints, handwriting, DNA.

Time (when?)

Time refers to the date, time, and time zone where the analyst is located, as well as the date, time, and time zone of the analyst's focus of interest. The analyst's focus of interest could also cover periods of various length, and be situated in the past, present or future.

Activity (what?)

The activity refers to all the elements that describe the analyst's current endeavour. What activity is the user currently conducting? How does it related to a given goal, task, role, and objective?

Information space

What information is currently available to the user? This includes information present in computer systems, but could also be extended to reflect information that could be made available by other resources (human, non-electronic). Here the term information is used in a broad sense and also entails information about processes and information exploitation capabilities.

State of the user

This includes information that can reflect the user's physical and emotional state. State can also refers to the "location" of the user in a workflow or particular process. The state of the user can also reflect the stress level of the user, which could tie in to a reflection of the actual (task, cognitive or information) overload of the user.

Context also encapsulates a variety of other elements that, depending on the situation, can have a major impact on the way the analyst conducts activities. Among additional context-related elements are: user preferences, calendar, behavioural patterns, relationships, focus of attention, state of the physical environment, state of the computing system, history of a user.

4.2.1.1 Context-awareness

From an iVAC system perspective, context-awareness refers to the system's ability to adapt to the analyst's context [17]. A context-aware system is also aware of an analyst's state and surroundings, and adapts its behaviour [18]. A context-aware system should be able to adapt its operations to the current context without explicit user intervention. It should be able to understand what (why and how) the user is doing, and what he is likely to be doing next.

From a general perspective, a context-aware system should aim at [16]:

- Improving relevance,
- Minimizing disruption,
- Improving awareness,
- Reducing overload, and
- Selecting interactions.

Improving relevance involves deciding which information is relevant to the user's current (or future) situation. Minimizing disruption implies deciding when to notify the user of relevant information for its situation. By improving awareness, the system identifies information and mechanisms that can help the user better conduct his activities. An analyst faced with cognitive and information overload problems is going to miss relevant information or processes available to support him in his work. Identifying useful elements that have been missed is essential in order to improve the analyst's awareness. The keyword here is *useful* elements. The system should be able to provide those useful elements, and only those. Providing the analyst with useless information or processes would only contribute to the overload problem. Reducing overload means not immediately providing information deemed non-relevant in a given situation. By selecting interactions, the system determines the best way to communicate information to the user (HCI). Some situations may warrant notifying the user by email, while others (possibly more urgent) may require a notification by phone, for example.

The interface should adapt to the role and current tasks of the user, as well as his preferences. The iVAC should be proactive, observing what the user is doing, his mental, cognitive and emotional state and, should suggest sequences of events, or wait for an appropriate moment to step in.

4.2.2 Seamless natural interaction

It is essential, in the iVAC vision, to have a virtual assistant that interacts in an intuitive, seamless fashion. Not only in the way that it communicates and exchanges with the user, but also in the way that it can bring up topics or suggestions. The iVAC should understand natural language, recognize gestures (especially deictic gestures, e.g., pointing at something). The iVAC should be able to resolve ambiguities from different communication modes, like pointing at a display and saying “what is that?” The iVAC should enable a discourse between the analyst and the machine.

The iVAC should be aware of the user state (cognitive overload, stress, fatigue), physical context (physical properties of the room, devices available, mobility, spatial location of the analysts) and working situation (collaborating with other analysts, giving a briefing, etc.).

An iVAC should allow the users to interact with the system(s) using multimodal interaction (e.g., voice, pointing, writing, drawing, gesture, eye/gaze, neural/brain interfaces, emotion detection) [19]. Moreover, "full understanding requires identification of speakers and addressees, along with resolution of reference to other participants and objects, and integration of both verbal and non-verbal communication" [21]. In a military context, the communication with an iVAC will evolve around topics such as tasks, planning activities, standard operating procedures (SOPs), briefing material, documents.

In many cases, as the user will be using natural language interaction combined with movement, there will be some ambiguities. For example, in the PAL video [22] the user says: "These are my priorities… I'm attending this meeting… I need you to setup my briefing package", while pointing at the screen. The deictic references to 'my priorities and this meeting' cannot be resolved by speech alone, and in some cases the location being pointed to cannot be resolved to a high-enough precision by vision alone.

The iVAC should be able to interact with the users by conducting several tasks: present tools and information; answer questions; ask for clarification; remind the user of some tasks or procedures, and propose alternative possibilities. The iVAC may even give a briefing. This interaction is provided through voice output and/or information display. As illustrated in the PAL video [22], the ISA should be able to interact directly with the user's screen, such as: highlighting information elements already displayed; display information in a new window, organized as the user wants to see it; gather a set of documents; filter information based on user requests.

If the user moves to a large screen display, the user should be recognized using biometry and the interface should adapt to the distance of the user from the display in terms of font sizes, granularity and quantity of information presented.

4.2.3 Evaluating, gaining, and maintaining trust

The iVAC should, over time, behave and perform in such a way that it generates a high level of trust in the user. This is particularly important as in the intelligence domain, the problems addressed are often complex and involve critical consequences, the analysts are highly skilled people, and they are imputable in their decisions. As with any collaboration with another person, the level of trust in the iVAC will need to develop with time, supported by various strategies.

Some of the strategies could be the following: the iVAC should provide access to a traceability of the rationale behind any action taken or recommendation made (e.g., interpretation and argumentation, as well as the sources of information); the system should allow the user to associate confidence indicators for different types of tasks (e.g., the user may develop trust more rapidly in the weather prediction than in the recommendations for courses of action); appropriate training should be provided, in particularly, based on reinforcement learning; the iVAC should exploit emotion recognition (Section 4.2.2) and cognitive state recognition (Section 5.1.2), in order to determine if its actions and recommendations bring confusion to the user.

Reeves [23] states that “The presence of a character can increase and sustain trust”. A first factor is the mere presence, where people believe that social presence is useful and preferred, in particular in situations where errors are likely. People also believe that these characters can add specific information to a conversation or during the conduct of a task, which increases trust. It can be in the form of acknowledging a comment, gathering appropriate information, or providing credibility because of expertise in the related domain.

4.2.4 Learning

Learning is a central part of the iVAC as it is pivotal in order to achieve the iVAC vision detailed in Section 3.1. The system must evolve from observing the user while he performs a given task to toning on a specific role. This leads directly to the iVAC capability to learn both user and tasks models (Section 4.1). Learning is a very broad topic. Its use for iVAC is discussed in Section 5.2 and additional information on various learning approaches can be found in [24].

5 iVAC S&T components

This section provides an overview of the iVAC's most central S&T components. The main topics being explored are HCI and Artificial Intelligence (AI). A complete description of these topics can be found in [10].

The extensive literature survey described in Section 2.1 allowed for determining that AI and HCI represent over 50% of iVAC related research activities. Figure 3 shows the distribution of analysed scientific records across various research topics [9]. The total percentages reflected in the chart exceed 100% due to the overlap between terms. The figure shows that the remaining 50% is spread over numerous disciplines including KM, NLP and "other" disciplines (including automation, adaptation, decision support, context-based military applications, planning and situational awareness), often pointing back to AI or HCI.

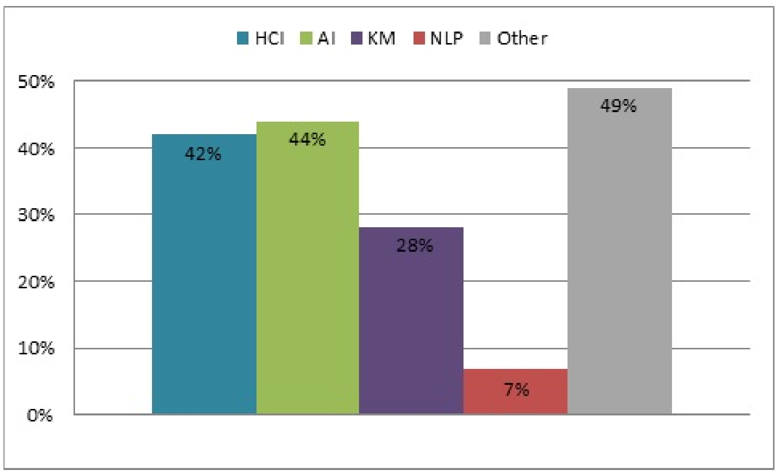

Figure 3: *Percentage of records that include each scientific topic [9].*

5.1 Human computer interactions

HCI handles the interactions between the virtual assistant and the users. More specifically, it supports the dialogue between the iVAC and the analyst(s), and optimizes the presentation of the results. This section presents some HCI elements that can help accomplish these objectives. Very detailed information on these and additional components can be found in [19][20].

5.1.1 Embodied agents representation

An Intelligent Software Assistant (ISA) can take different forms: a speech output, a simple icon, a two- or three-dimensional representation of a character (avatar), or simply appear through interactions such as highlighting or pop-ups. Figure 4 shows an automated online assistant, having a text-based dialog system and a humanoid avatar. The ISA can be personified, in particular from a gender perspective (male, female, neutral). The avatars can also exhibit other characteristics (age, ethnic group, profession – civil vs military) based on the users' social-cultural context and preferences. Facial and voice features (serious/smiling face, tone of the voice) could reflect the importance/urgency, or certainty of a message. Different avatars could be used to support different ISA tasks. For example, the avatar for the weather analyst might be different, in terms of gender, age and profession, from the one that recommends the course of action.

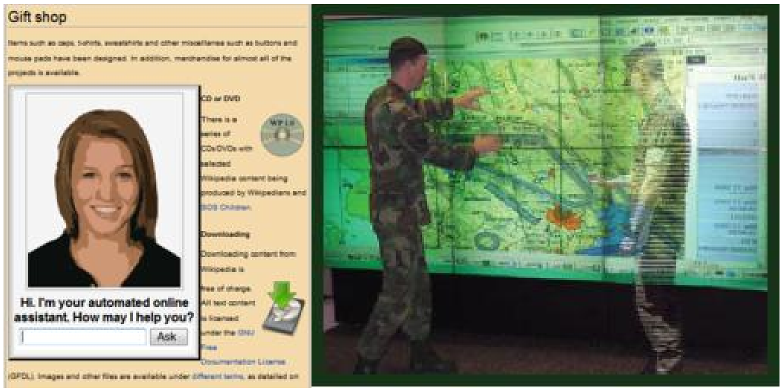

Figure 4: *a) A humanoid ISA avatar with a text-based dialog [25];*
b) An ISA avatar embedded in the display [22].

5.1.2 Multimodal interaction

Multimodal interaction provides the user with multiple modes of interfacing with a system. Combining multiple interaction modes can help resolving ambiguities in communications. Examples of interaction modes and technologies are: neural interfaces, gestures, multi-touch tables, etc.

A neural interface, or brain-computer interface (BCI), is sensing the brain waves in order to potentially detect certain states of the user. An example of such an interface is the Emotive EPOC neuroheadset [26] shown on Figure 5. It uses sensors to tune into electrical signals produced by the brain to detect user thoughts, feelings, and expressions.

Figure 5: *Picture of the emotive EPOC BCI.*

The use of a BCI is the object of a research effort that is part of the iVAC project. This research aims at sensing the user's internal context in order to streamline information based on factors such as relevance and attention. Detailed information on this activity can be found at [27].

5.1.3 Display technologies

Among other things, display technologies range from small to large, and from personal to shared devices. The interaction between the analyst and the iVAC needs to be tailored to the devices that are available to the iVAC for communicating with the user.

Interaction with an iVAC could involve multiple different display settings. First, in an individual setting, one can expect various combinations of display screens at the analyst's workstation. Head-mounted displays are also of interest, especially in mobile environments. Moreover, head-mounted and related display technologies enable a private view of the information and reduce the risk of sensitive information being seen by other people. An example of head-mounted display that has received a lot of attention recently are the Google glasses (see Figure 6). Although they say that it is still years away, the entry of Google in the augmented reality domain with it's 'project glass' has the potential to bring augmented reality head-mounted displays to a broader audience in the next few years.

Figure 6: *Google glasses.*

When the iVAC needs to communicate with a group of analysts, large group displays may be preferred. The technology for large group displays is continuously evolving, resulting in the availability of larger and more capable displays at lower costs. New interaction and collaboration technologies, such as the use of multi-modal interaction and more comprehensive support for co-located and distributed collaboration are also emerging. Figure 7 shows a mosaic of LCD panels. Gouin et al. [29] provide a number of human factor guidelines in terms of the design and use of LCDs. This included human engineering (e.g., organizational, team and cognitive considerations, legibility guidelines, layout considerations), interaction and collaboration (e.g., multi-modal interaction, alerts and notifications, local and distributed collaboration), organization of displayed information (e.g., content design heuristics) and content control.

Figure 7: *Sharp PN-V601 60 Inch LCD video wall.*

5.1.4 Information visualization/visual analytics

Beyond the display technologies used, a much more important concern is about what should be presented on these devices. The intelligence function is very complex. The multiple dimensions of the information and the quantity of linkages to understand may lead to a cognitive overload.

Recently, visual analytics has emerged as a multidisciplinary field of research that leverages information visualization and a number of other disciplines. “Visual analytics is the science of analytical reasoning facilitated by interactive visual interfaces” [30]. “Visual analytics focuses on the following areas:

– Analytical reasoning techniques that enable users to obtain deep insights that directly support assessment, planning, and decision making;

– Visual representations and interaction techniques that take advantage of the human eye’s broad bandwidth pathway into the mind to allow users to see, explore, and understand large amounts of information at once;

– Data representations and transformations that convert all types of conflicting and dynamic data in ways that support visualization and analysis; and

– Techniques to support production, presentation, and dissemination of the results of an analysis to communicate information in the appropriate context to a variety of audiences.” [30].

In other words, “the basic idea of visual analytics is to combine the strengths of automatic data analysis with the visual perception and analysis capabilities of the human user. It uses visualizations, user interaction and data analysis techniques to find insight from complex, conflicting and dynamic information. Visual analytics (Figure 8) is especially focused on situations where the huge amount of data and the complexity of the problem make automatic reasoning impossible without human interaction” [31].

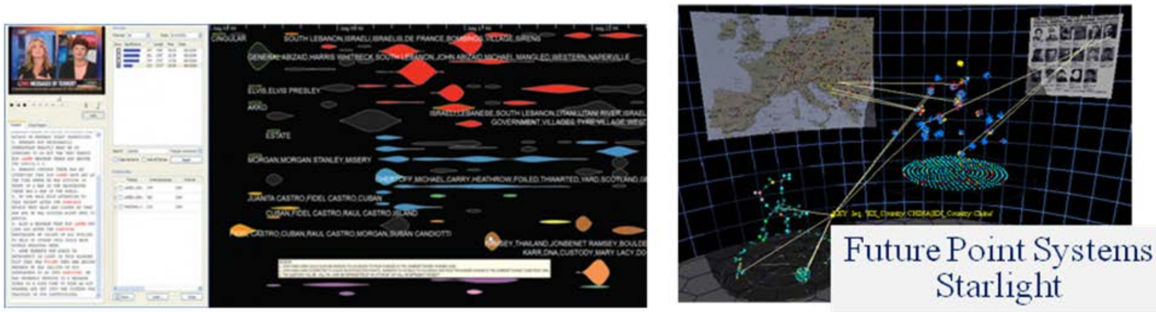

Figure 8: *a) Visual analytics examples: River theme visualization [31];*
b) Galaxy visualization [32].

5.2 Artificial intelligence

AI is a branch of computer science that aims at developing smart machines and software that emulate human thinking. This ambitious objective is tackled by research efforts in numerous AI disciplines. In the context of iVAC, the particular AI topics that are identified as especially relevant are: QA and Learning by Demonstration.

5.2.1 Question answering

QA is an emerging field in AI, in which a computer is specifically trained to provide answers to users' questions. QA was popularised in 2011 by the artificially intelligent computer system named Watson, which was developed by IBM's Deep QA project [32], and that was specifically developed to answer questions on the quiz show Jeopardy.

QA engines are an evolution of the typical search engines, where instead of having the computer retrieve documents which may contain a solution to a problem or the answer to a question, the system tries to extract a precise answer from retrieved documents.

QA is a combination of various fields that include:

- NLP, which is concerned with understanding the human's query and with producing an intelligible answer, and
- Information retrieval, which retrieves the collection of sources that are likely to contain relevant information.

In the context of iVAC, QA is of high interest as it directly supports the first tier of the vision explained in Section 3.1. QA is perfectly suited to provide support to the user in any given task, as long as there is an available corpus of documents that contain relevant information. QA is the object of a research effort that is part of the iVAC project. This effort aims at developing a framework for building adaptive intelligent virtual agents. In this framework, the QA system is implementing three main functions:

- Question analysis, which consists in understanding what is being asked by the user. At a minimum, the system must classify the topic of the question and extract important keywords featured in the question,
- Passage retrieval, which retrieves and filters out text fragments that might contain the answer sought by the user, and
- Answer generation, which presents the answers in a format acceptable to the user.

Detailed information about this activity can be found in [34].

5.2.2 Learning by demonstration

Learning by Demonstration (LBD) is a learning paradigm where an agent learns some new behavior through examples provided by a teacher (demonstrator). These examples indicate to the agent how a task should be accomplished. LBD is a form of supervised learning, which is to say it requires input from a trainer in order to function.

In a particular form of this approach, the agent observes the demonstration performed by the teacher and does not take part in it. This is called Imitation Learning. In Imitation Learning, observations are not directly applicable and the agent needs to convert them into useful examples. For instance, to imitate some human gestures, an agent would have to map the human moves into some internal representation. Once a mapping has been applied to the observations, techniques for LBD can be applied. This particular learning method is very well suited to iVAC to achieve the progression needed to evolve from task support to task and role execution.

6 iVAC capability mapping

Section 4 presented the sub-capabilities that need to be implemented in order to instantiate a full-fledged iVAC system. Section 5 has gone over S&T topics that are of particular interest for an iVAC. In this section, the mapping between iVAC capabilities and Intelligent Software Assistant (ISA) technologies and projects is described. This section is meant to provide the reader with an overview of the complete iVAC capability mapping which can be found in [12].

Information on which S&T component can be applied to a given capability is provided. Each iVAC sub-capability of Section 4 is addressed individually, and aspects of interest are highlighted. Detailed information on the ISA technologies mentioned in this section is provided in Annex A.

6.1 Manage context

Figure 9 provides an overview of the mapping between ISA technologies or research projects and context management sub-capabilities. The table intersects each sub-capability with a technology or research project. Each intersection is colored based on the applicability of a given technology/project to a capability.

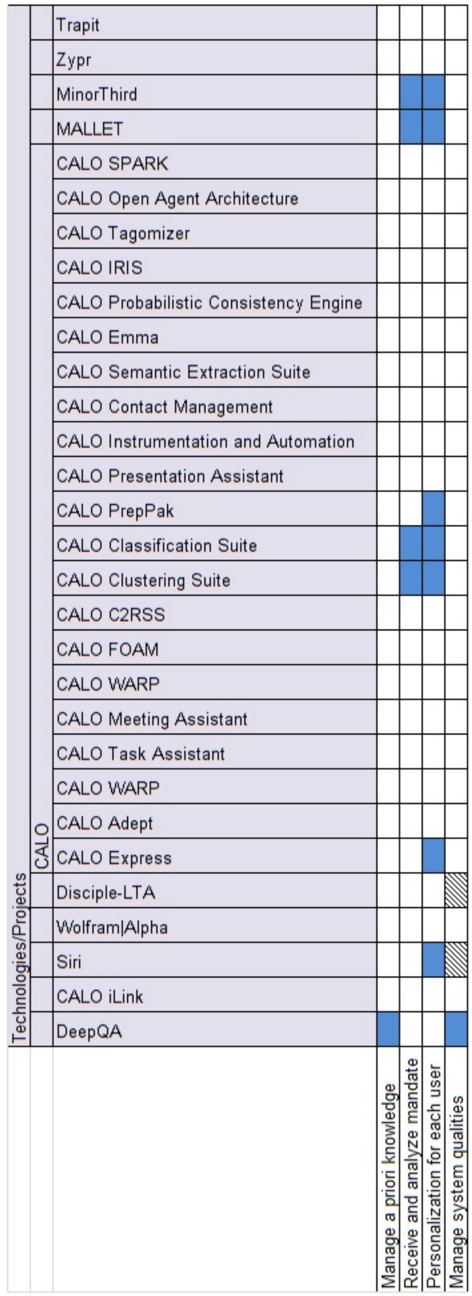

Legend	Value	Description
	Not applicable	It is considered that the technology does not apply to the given iVAC capability.
	Not enough information available	It is considered that the technology may be applicable, but the available information does not allow to estimate to which level.
	The technology can be adapted to meet the capability	The technology could be adapted to fulfil the needs of the capability, in whole or in part.

Figure 9: *Context management capability mapping.*

For the management of a priori knowledge, DeepQA [33] is the only project that contains capabilities allowing for the adaptation to a specific domain of knowledge, such as arts, medicine, politics, etc. However, it does not include the capability to represent more complex contextual information such as role, task, or state of the user.

For the analysis of the mandate, some of the CALO [35] modules implement one or multiple algorithms of ML: CALO Clustering Suite, CALO Classification Suite. It's assumed that, in practice, the mandate provided to the analyst would be clearly described. But its description might require some inference to fulfil some mission parameters left to the discretion of the analyst. Hence classification and regression techniques could be used to determine some of the parameters. MALLET [36] and MinorThird [37] implement multiple NLP algorithms and techniques, which could be necessary for doing an analysis of a mandate defined in a textual (i.e., unstructured) form.

CALO Express is using the PrepPak [38] module in order to give personalized suggestions for documents relevant to the user's current activity, such as writing an email message or scheduling a meeting. Taking into account user preferences is somewhat common in a variety of applications. However, tying preferences to the user's context as a whole is not common, even less so when one considers the user's inner context, which is harder to sense (see Section 4.2.1).

Siri is a voice activated personal assistant software for the Apple iPhone 4S, initially developed by Siri Inc., and then by Apple, Inc. Siri takes into account certain aspects of the user's external context (e.g., user location, time, user contact list) to answer simple requests (e.g., "When I get home email my wife.").

DeepQA has a built-in capability for quality management via a technique of hierarchical ML, where the performance of the system can be (semi-)automatically improved/tuned given new examples of question/answer pair.

At a glance, it is obvious that context awareness is an area that does not benefit from the support of a lot of technologies/projects. More specifically, the capability of a system to adapt to a user's context (both internal and external) as a whole is lacking.

6.2 Acquire data, information and knowledge

Figure 10 provides an overview of the mapping between ISA technologies or research projects and the acquisition of data, information and knowledge sub-capabilities.

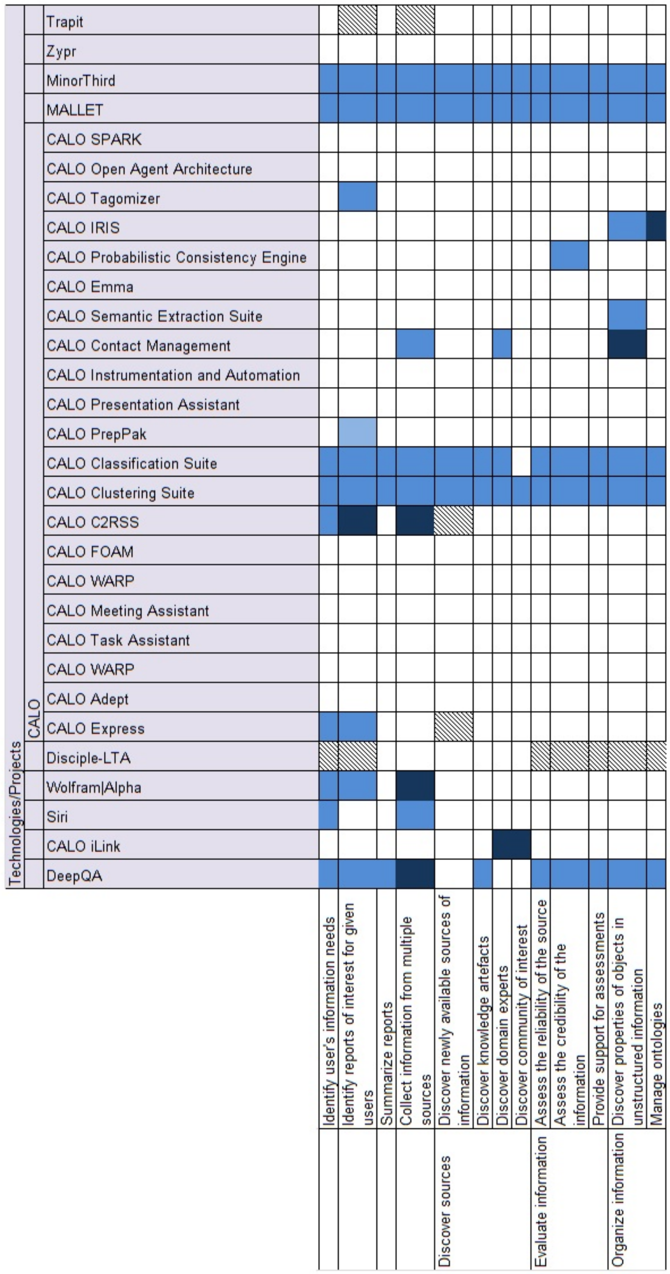

Legend	Value	Description
	Not applicable	It is considered that the technology does not apply to the given iVAC capability.
	Not enough information available	It is considered that the technology may be applicable, but the available information does not allow to estimate to which level.
	The technology can be adapted to meet the capability	The technology could be adapted to fulfil the needs of the capability, in whole or in part.
	The technology has the capability	It is considered that the technology addresses the capability in full or almost.

Figure 10: *Acquire data, information and knowledge capability mapping.*

This figure shows that out of 13 sub-capabilities, six are being directly addressed by a technology that meets the capability. The remaining seven sub-capabilities have technologies that can be adapted to meet their needs. Data, information and knowledge acquisition is a more mature research field, and as such it is normal that a larger number of solutions exist to support this capability. This is not to say that data, information and knowledge acquisition is of lesser importance or should not be considered as a potential research topic. However, one should be aware that there are a lot of potential leverage points in existing, ongoing research efforts and technologies.

6.3 Monitor, schedule, manage and evaluate activities

Figure 11 provides an overview of the mapping between ISA technologies or research projects and the monitor, schedule, manage, and evaluate activities sub-capabilities.

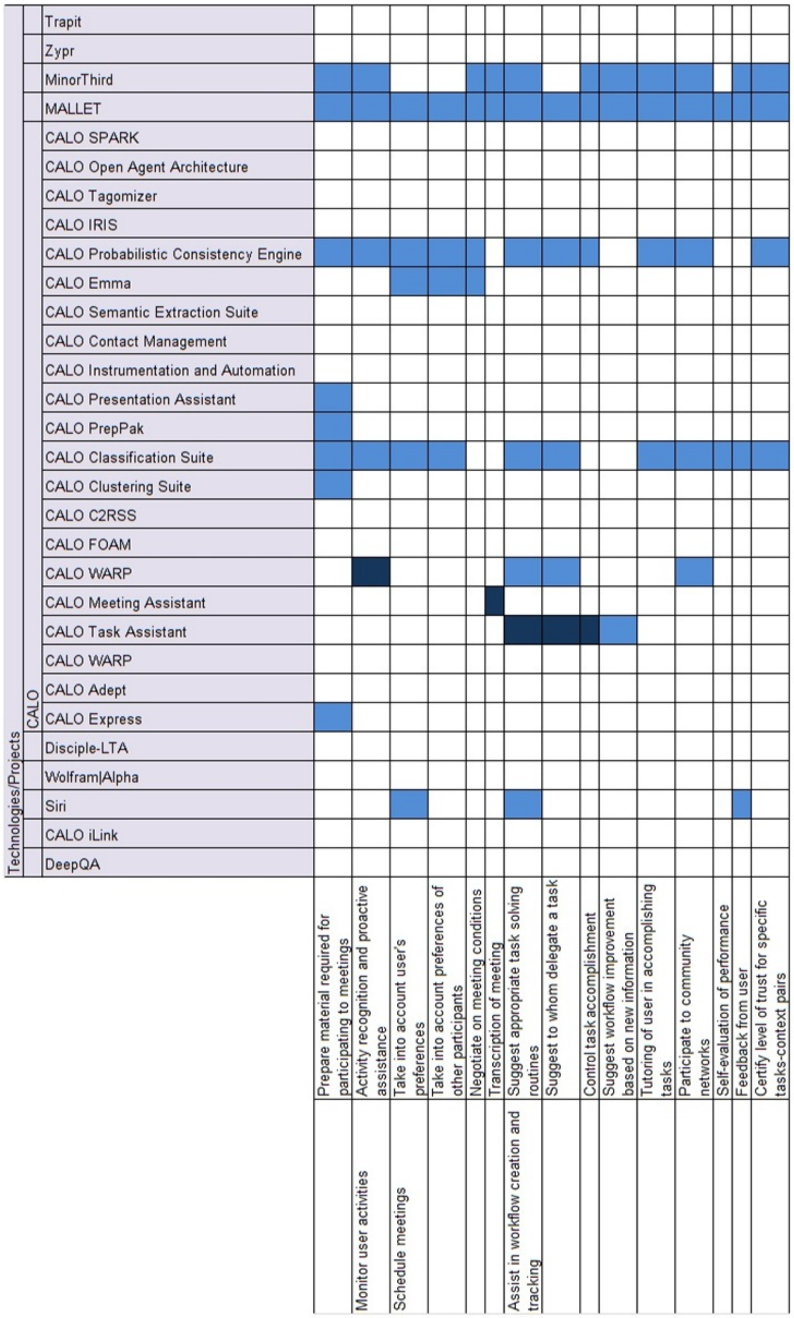

Legend	Value	Description
	Not applicable	It is considered that the technology does not apply to the given iVAC capability.
	Not enough information available	It is considered that the technology may be applicable, but the available information does not allow to estimate to which level.
	The technology can be adapted to meet the capability	The technology could be adapted to fulfil the needs of the capability, in whole or in part.
	The technology has the capability	It is considered that the technology addresses the capability in full or almost.

Figure 11: *Monitor, schedule, manage and evaluate activities capability mapping.*

A lot of the sub-capabilities identified in Section 4.1 have more to do with work organization than with actual analysis capabilities. However, the monitoring of user activities is important in order to provide the system with inputs allowing to support and learn user tasks and roles. As Figure 11 shows there are tools that support this sub-capability, which should be considered when planning the development of an iVAC. Self-evaluation and user feedback are also important. These capabilities will allow the system to improve its performance and evolve from simple task support to autonomously taking on tasks or roles. This aspect is not covered as well, and will likely require more effort.

6.4 Learn user and task models

Figure 12 provides an overview of the mapping between ISA technologies or research projects and the learn user and task models sub-capabilities.

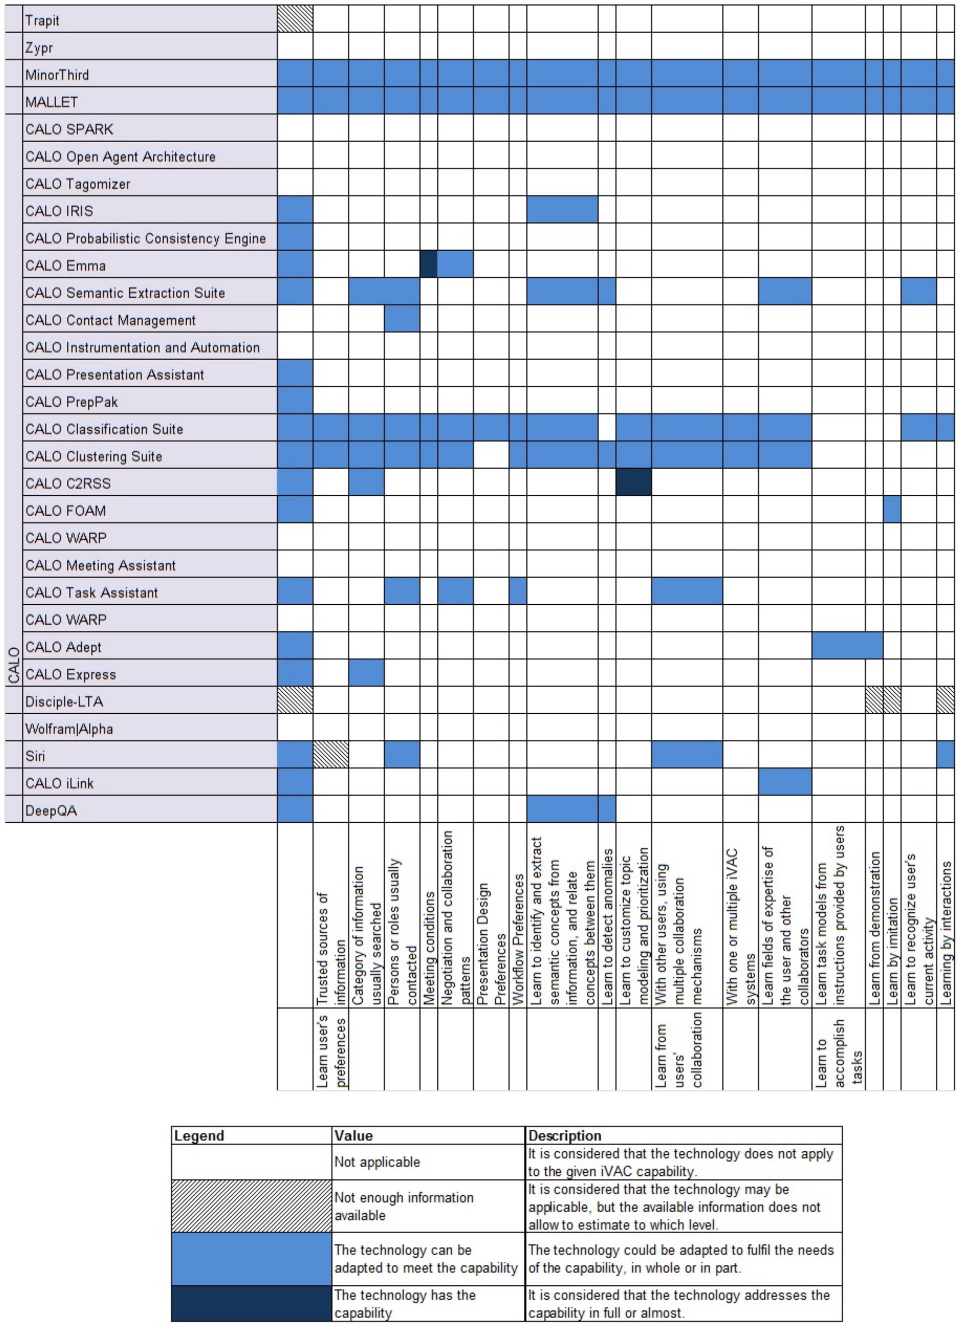

Figure 12: *Learn user and task models capability mapping.*

The capabilities associated with learning users' preferences have a variety of technologies or projects that could be adapted to address the need of these capabilities. The same goes for learning from users' collaboration. There are less existing solutions that support learning to accomplish tasks. This makes sense as these three sub-capabilities are of increasing complexity.

Indeed, it is obviously easier to learn a user's preference for a given task than it is to learn how to accomplish a complete task. This, again, directly relates to the challenges of achieving the complete iVAC vision (Section 3.1) which involves evolving from task support to autonomous task/role execution.

6.5 Support complex intelligence tasks

Figure 13 provides an overview of the mapping between ISA technologies or research projects and the support complex intelligence tasks sub-capabilities.

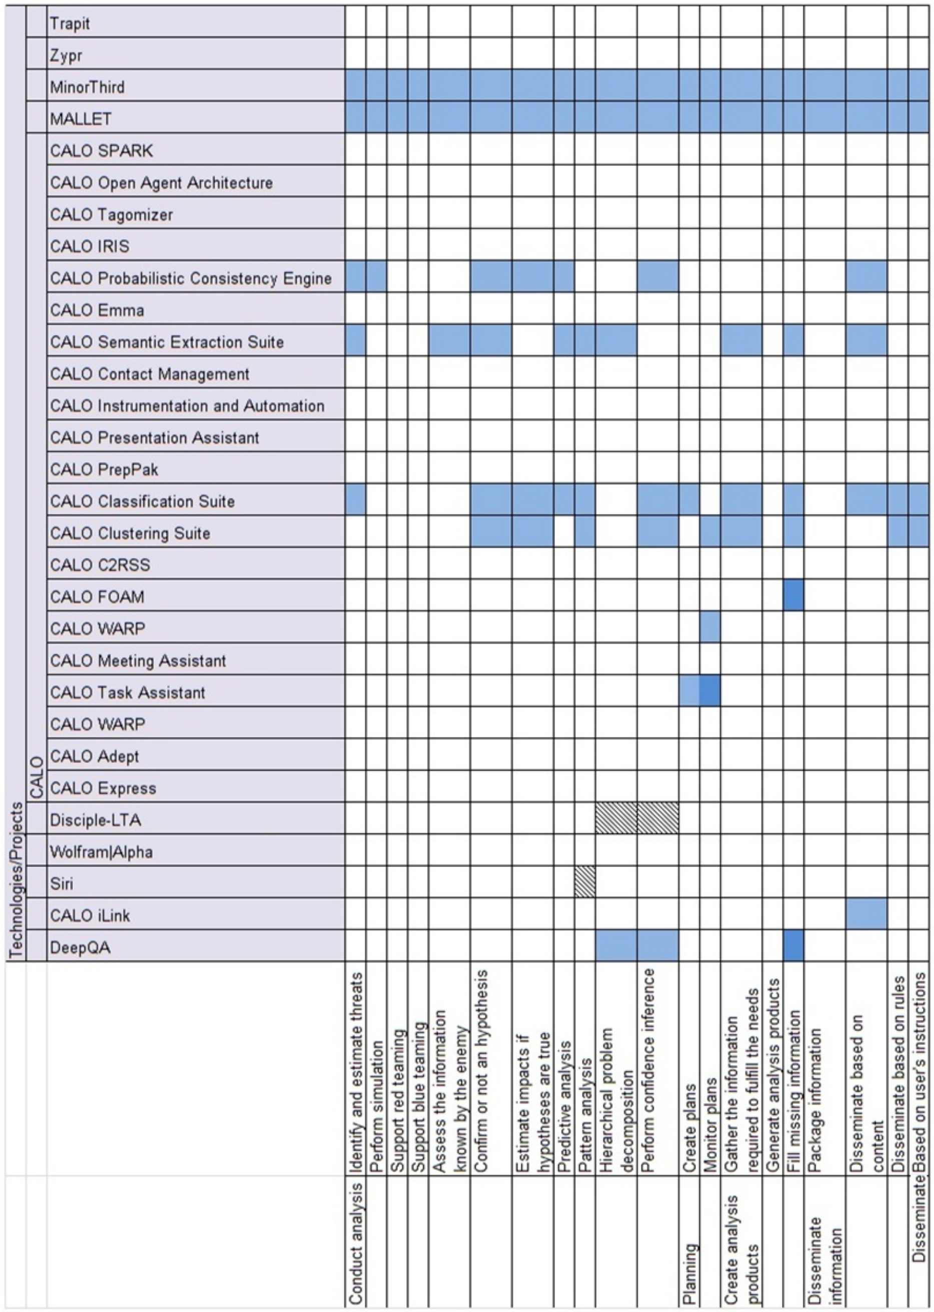

Figure 13: *Support complex intelligence tasks capability mapping.*

One can notice that none of the ISA technologies or projects of Figure 13 directly addresses the complex intelligence tasks. While some of the surveyed solutions could be adapted the meet the needs of intelligence analysis, none of them is directly meant to do so. On one hand, this is normal, as the basic focus of these technologies in not the intelligence domain, but rather the development of a more generic autonomous cognitive support tool. On the other hand, the development of intelligence analysis support tools is, and has been, a focus of research at

DRDC Valcartier for a number of years. Specific algorithms and technologies have been developed to support the Intelligence analysis problem. These tools range from Multiple Hypothesis Situation Analysis (MHSA) [39], automated reasoning to support the sense making process [2][4], text analytics [1], and knowledge representation [40]. These DRDC developed solutions should be considered as serious candidates for integration in an iVAC system. While the iVAC provides the seamless interface, adaptiveness, and learning aspects, the aforementioned DRDC solutions could represent the back bone of the intelligence analysis capability to support complex analysis tasks.

6.6 Interaction with humans

Seamless interactions with humans are at the core of the iVAC system (4.1.6). Figure 14 provides an overview of the mapping between HCI technologies and the interaction with humans capability. A complete, detailed mapping and description is available at [10]. It should be noted that for this particular section, HCI is considered in a more generic fashion, as the enumerated sections refer more to families of technologies than to an actual single technology.

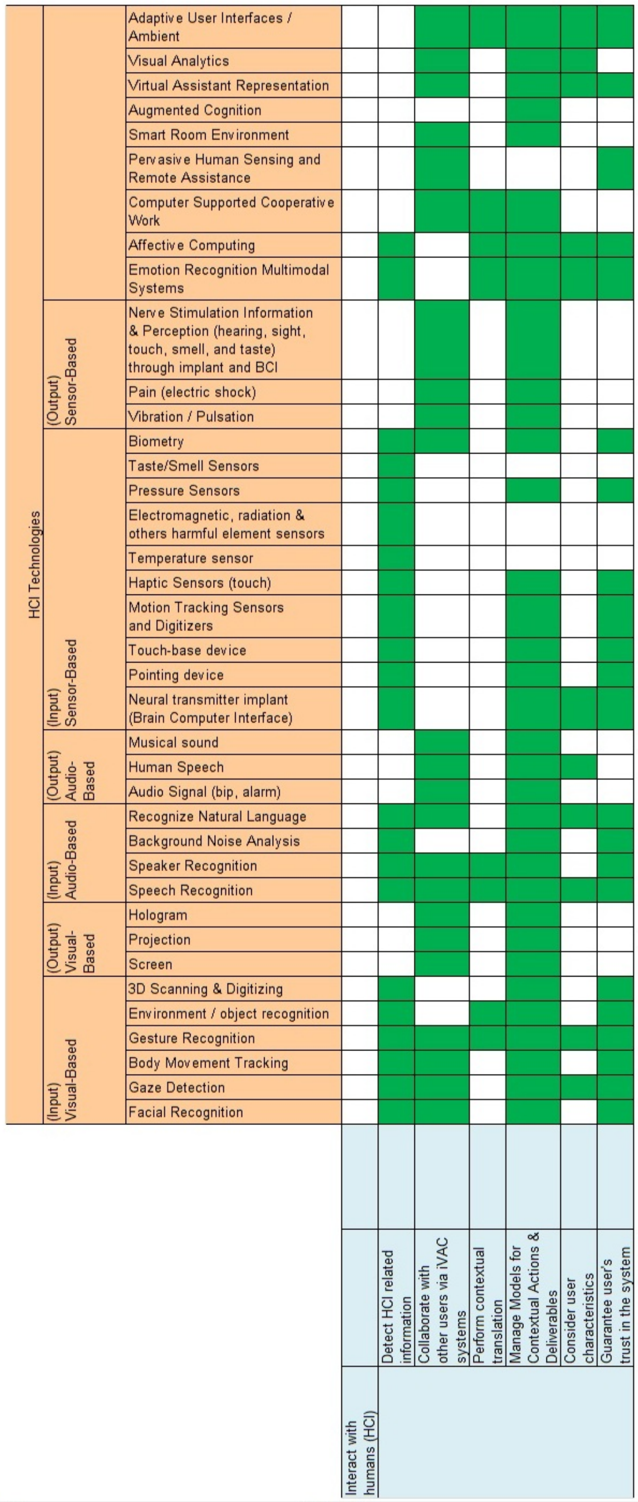

Legend	Value	Description
	Probably not applicable	It is considered that the technology may not apply to the given iVAC capability.
	Applicable to the capability	It is considered that the technology would be involved to address the capability.

Figure 14: *Interact with humans capability mapping.*

Detection of HCI related information is done by analysing information from multiple sources and detecting various aspects that will be influencing the interaction of the iVAC with the user. Any technology that can detect aspects related to the user and his environment will potentially contribute to addressing this iVAC capability. As the figure shows, there is a plethora of paradigms that fit the various iVAC HCI requirements in some fashion. However, it should be noted that none of these paradigms is directly suited to the iVAC context. Therefore, the challenges lie in the integration efforts that would be required between these various technologies for every iVAC HCI sub-capability.

Facial detection would be helpful in identifying the user, but also to identify the emotions that could be expressed by his face and detected by the system. Gaze detection could be used for detecting what the user is looking at. Body movement tracking would be useful to detect where the user is in the environment, but movements can also be used to derive emotions and intentions. Speech recognition and NLP are important to understanding what the user is saying as well as to infer cultural, emotional and intentions aspects. Multiple sensors can sense the environment, the user’s body, and even his brain.

Users are collaborating using multiple means. The system could allow them to collaborate through the iVAC and not only via "external" systems. Multiple HCI technologies may be involved within an iVAC system to enable collaboration with other users. Any technology used by the user when working alone could be involved when collaborating with other users. Technologies specifically designed for collaboration will definitely be applicable for addressing this capability, namely: computer supported cooperative work, smart room environments, visual analytics, adaptive user interfaces, etc. Similarly, technologies for detecting what the user is doing (gaze detection, gesture recognition, etc.) and transmitting images of him to others (holograms, projections, screen) will also be applicable as they will contribute by making collaboration environments more immersive. Sensor-based technologies will also apply, e.g., for making remote collaboration as seamless as possible and for breaking the distance barrier when collaborating users are working at remote sites.

Any interaction is performed in a specific environment or context. This context must be taken into account during these interactions to ensure that the context of each other is properly understood, as it impacts the way they behave. The importance of context awareness is the subject of Section 4.2.1.1 and this awareness is directly related to the notion of contextual translation. Contextual translation could be considered as an advanced capability of an iVAC and today’s technology may not be fully mature for supporting this capability. Nevertheless, advances in affective computing, emotion recognition and adaptive user interfaces will eventually be applicable to this capability as they allow detecting and using the users’ contextual information. Since context translation involves understanding and exploiting cultural aspects, technologies like speech, speaker, gesture and environment recognition will be applicable.

When interacting with the user, the system should consider user’s characteristics, like mental state, role and rank, security clearances, etc. The two main technologies applicable to this family of capabilities are the virtual assistant representation and adaptive user interfaces. For the virtual assistant, it is expected that the iVAC representation will change depending on user characteristics. Today’s technology is evolving towards that capability. For interfaces, it is expected that adaptive user interfaces must consider user’s characteristics (amongst other factor), when adapting an interface to the user.

Over time, the system should behave and perform in such a way that it generates a high level of trust from the user. As the user works with the system, he should trust the system more and more to support him for performing his tasks. For the system to reach that goal, it should have the capability to estimate the current level of user's trust. Then, it should adapt its behavior so as to generate trust from the user. This capability is at the heart of the relation between the user and the iVAC system, and more importantly at the heart of the success of an iVAC system deployment. This iVAC capability can be considered a futuristic capability: the system behaves in such way as to gain trust from the user. Based on the surveys and research performed, no technologies can achieve that today. Nevertheless, any HCI technologies that can potentially provide input to the system with respect to the user's trust would be applicable to that capability.

6.7 Interaction with other systems

At the centre of the iVAC vision lies the need to access a plethora of information sources in the hope of exploiting them using integrated and potentially heterogeneous analysis capabilities. Achieving this vision touches upon two major research challenges: the need to access and manage very large volumes of heterogeneous data sources, and the need to develop a flexible, scalable and efficient software integration architecture.

The term **big data** refers to a collection of data sets so large and complex that it becomes difficult to process using on-hand database management tools or traditional data processing applications [41]. This is the reality that modern systems are increasingly faced with, and the intelligence analysis world is certainly no exception. Intelligence analysts are provided with enormous amounts of data originating from a variety of source and sensors (texts, images, audio, video, etc…). In order for an iVAC to leverage this data, it has to be able to manage and access it efficiently. Big data is an active research topic, with initiatives from:

- the public sector: in March 2012, The White House announced a national "Big Data Initiative" that consisted of six Federal departments and agencies committing more than $200 million to big data research projects [42]; and
- the private sector, with software developers putting forward a variety of solutions: Hadoop (Apache Foundation) [43], MongoDB (MongoDB Inc.) [44], Splunk (Splunk INC) [45].

It is obvious that big data approaches will play a major role for the development of an iVAC system. As big data is a very active research topic, it is reasonable to assume that leveraging existing research would be a logical first step before investing research effort in this topic for iVAC.

7 Conclusion

This report gave an overview of the work performed in the context of the Intelligence Virtual Analysis Capability project. The objective of this document was not to provide the reader with every last detail of the work performed, but rather to present the concepts explored during the work, while providing pointers to additional documents containing detailed information and highlighting specific topics identified as having particular interest.

The methodology employed was discussed and referenced, from the literature reviews, state of the arts and workshops, to the analysis tasks. The iVAC concept was explained in relation with the Future Intelligence Analysis Capability (FIAC) vision. Sub-capabilities of the iVAC were identified and described. Some sub-capabilities were discussed in greater details as they are of great importance for the iVAC vision and were not the focus of many current research projects: context awareness, seamless natural interactions, evaluating, gaining, and maintaining trust and learning. iVAC sub-capabilities were finally enumerated and mapped to potentially relevant technologies or research projects.

DRDC Valcartier has developed a framework for the integration of heterogeneous analysis components [46], which forms a basis for the development of an iVAC. In addition, other research initiatives driven by DRDC Valcartier are focusing efforts on the development of QA functionalities [47], and context aware computing [48]. Future efforts will concentrate on integrating the functionalities developed under these initiatives into the iVAC framework.

Carrying out the iVAC vision is no small task. Although relevant science and technology components are maturing and becoming increasingly accessible, important integration and development efforts are required to build an iVAC. It is the belief of the authors of this report that this effort would be worth undertaking, and that this document provides a good starting point and a way ahead for this task.

This page intentionally left blank.

References

[1] Roy, J., Auger, A., “The Multi-Intelligence Tools Suite – Supporting Research and Development in Information and Knowledge Exploitation” in *Proceedings of 16th International Command and Control Research and Technology Symposium “Collective C2 in Multinational Civil-Military Operations”*, Québec City, Canada, June 21–23, 2011.

[2] Bergeron Guyard, A., Roy, J., “Toward Case-Based Reasoning for Maritime Anomaly Detection: A Positioning Paper” in *Proceedings of The Twelfth IASTED International Conference on Intelligent Systems and Control (ISC 2009)*, Cambridge, Massachusetts, November 2009.

[3] Bergeron Guyard, A., “Case-Based Reasoning for Maritime Anomaly Detection”, in *Proceedings of COGnitive systems with Interactive Sensors (COGIS 2010)*, Crawley, United Kingdom, November 2010.

[4] Roy, J., “Rule-Based Expert System for Maritime Anomaly Detection”, in *Proceedings of Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense VIX, SPIE Defense, Security, and Sensing 2010*, Orlando, FL, USA, 5 - 9 April 2010.

[5] Roy, J. and Bergeron Guyard, A., “Supporting Threat Analysis Through Description Logic Reasoning” in *Proceedings of , 2012 IEEE Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, New Orleans, LA, USA, March 6-8, 2012.

[6] Davenport, T.H., Prusack, L., “Working Knowledge: How Organizations Manage What They Know”. Boston: Harvard Business School Press, ISBN 0-585-05656-0, 1998.

[7] Rowley, J., Hartley, R., “Organizing Knowledge: An Introduction to Managing Access to Information”. Ashgate Publishing, Ltd., ISBN 978-0-7546-4431-6, 2006.

[8] Poussart, D., “Future Intelligence Analysis Capability – Towards a Cohesive R&D Program Definition”, *DRDC Valcartier Internal Draft Report, last revised in March 2013.*

[9] Wiseman, E., “CISTI, Intelligence Virtual Analyst Capability Literature Survey”, *RDDC Valcartier CR 2011-246*, 5 October 2011.

[10] Burkov, A., Bérubé, M., Lamontagne, L., Laviolette, F., “ISA technology and scientific literature Survey”, *RDDC Valcartier CR 2011-241*, 1 March 2012.

[11] Lamontagne, L., Laviolette, F., “State of the art study of automated/machine learning paradigms”, *RDDC Valcartier CR 2011-242*, 1 March 2012.

[12] Michaud, G., Burkov, A., Bérubé, M., Lamontagne, L., Laviolette, F., “Identification of iVAC capabilities”, *RDDC Valcartier CR 2011-243*, 1 March 2012.

[13] "Context." Merriam-Webster.com. Merriam-Webster, n.d. Web. 14 Mar. 2014. <http://www.merriam-webster.com/dictionary/context>.

[14] Schilit, B.N., Theimer, M.M. "Disseminating Active Map Information to Mobile Hosts", *IEEE Network 8 (5): 22–32*, 1994.

[15] Dey, Anind K., "Understanding and Using Context". Personal Ubiquitous Computing 5 (1): 4–7. 2001.

[16] Poellabauer, C., "Pervasive/Ubiquitous Computing – Context Awareness", *http://www.cse.nd.edu/~cpoellab/teaching/cse40816/Lecture5.pdf.*

[17] Schilit, B., Adams, N., Want, R., "Context-aware computing applications", *Mobile Computing Systems and Applications, 85-90,* 1994.

[18] Satyanarayanan, M., "Pervasive computing: Vision and challenges", *Personal Communications, IEEE 8(4), p. 10-17*, 2001.

[19] Gouin, D., Lavigne, V., "Assessment of Future Trends in Human-Computer Interaction for Military Intelligence", DRDC Valcartier TR 2012-425, June 2012.

[20] Gouin, D., Lavigne, V., Bergeron-Guyard, A., "Human-Computer Interaction with an Intelligence Virtual Analyst", *Proceedings of Knowledge Systems for Coalition Operations (KSCO 2012)*, Pensacola, FL, US, DRDC Valcartier SL 2012-148, February 2012.

[21] Purver, M., Niekrasz, J. and Peters, S., "Ontology-Based Multi-Party Meeting Understanding", in *Proceedings of CHI'05*, Portland, OR, April 2005.

[22] DARPA, "Personalized Assistant that Learns and Organizes", *video, http://www.youtube.com/watch?v=BF-KNFlOocQ&feature=youtu.be* , 2010.

[23] Reeves, B., "The Benefits of Interactive Online Characters", *Stanford University, http://noahx.com/resources/CSLI_Stanford_Study.pdf.*

[24] Lamontagne, L., Laviolette, F., "State of the art study of automated/machine learning paradigms", *RDDC Valcartier CR 2011-242*, 1 mars 2012.

[25] Wikipedia, "Automated Online Assistant", *http://en.wikipedia.org/wiki/File:Automated_online_assistant.*png.

[26] ENOTIV, "epoc", *http://emotiv.com/epoc/.*

[27] Ross, W., Morris, A., Ulieru, M., Bergeron Guyard, A., "RECON: An Adaptive Human-Machine System for Supporting Intelligence Analysis", in *Proceedings of IEEE International Conference on Systems, Man and Cybernetics (IEEE SMC 2013)*, Manchester, UK, 13-16 October 2013.

[28] Ackerman, E., "Google gets in your face", *IEEE Spectrum*, Vol 50, Issue 1, pp 26-29, Jan 2013.

[29] Gouin, D., Wark, S. and Vernik, R. (2009), Considerations and Guidelines for the Use of Large Group Displays in Command and Control Environments, TTCP C3I TP2 (Command Information Interfaces), December 2009.

[30] Thomas, J.J., Cook, K. A., “Illuminating the Path: The Research and Development Agenda for Visual Analytics”. *IEEE Press*, 2005.

[31] Järvinen, P., Puolamäki, K., Siltanen, P. and Ylikera la, M. (2009), Visual Analytics Final Report: Data Analytics, Visualization, Applications, Tools, VTT Working Papers 117, VTT Technical Research Centre of Finland, 45pp, 2009, http://www.vtt.fi/inf/pdf/workingpapers/2009/W117.pdf, retrieved January 2011.

[32] Ribarsky, B. (2009), Multimedia Visual Analytics, VAC Views, May 09, p. 10-11. Future Point (2011), Future Point Systems Inc., Starlight, http://www.futurepointsystems.com/, retrieved January 2011.

[33] IBM, “DeepQA Project”, https://www.research.ibm.com/deepqa/deepqa.shtml.

[34] Lamontagne, L., Bergeron Guyard, A., “A Framework for Building Adaptive Intelligent Virtual Assistants “, in *Proceedings of Artificial Intelligence and Applications 2014 (AIA2014),* Innsbruck, Austria, 17-19 February 2013.

[35] SRI International, “Cognitive Assistant that Learns and Organizes”, http://www.ai.sri.com/project/CALO.

[36] McCallum, A., K., “MALLET: A Machine Learning for Language Toolkit”, http://mallet.cs.umass.edu.

[37] Carnegy Mellon University, “MinorThird Documentation”, http://minorthird.sourceforge.net/old/doc/.

[38] SRI International, “PAL Framework Overview”, https://pal.sri.com/Plone/framework.

[39] Bergeron Guyard, A., Roy, J., “Multiple Hypothesis Situation Analysis: Governing Principles and Support System”, *RDDC Valcartier TR 2010-525*, December 2010.

[40] Bergeron Guyard, A., “Multi-Reasoner Framework to Support Situation Analysis and Context Awareness”, in *Proceedings of Artificial Intelligence and Applications 2013(AIA2013)*, Innsbruck, Autriche, 11-13 February 2013.

[41] White, T., “Hadoop: The Definitive Guide”. *O'Reilly Media. p. 3. ISBN 978-1-4493-3877-0*, 2012.

[42] The Computing Community Consortium Blog, “The White House, "Obama Administration Unveils "Big Data" Initiative: Announces $200 Million In New R&D Investments”, http://www.cccblog.org/2012/03/29/obama-administration-unveils-200m-big-data-rd-initiative/.

[43] Apache Foundation, “Hadoop”, http://hadoop.apache.org/.

[44] MongoDB, “MongoDB”, https://www.mongodb.org/.

[45] Splunk Enterprise, “Splunk”, http://www.splunk.com/.

[46] Bergeron-Guyard, A., “Integration Framework for Heterogeneous Analysis Components – Building a Context Aware Virtual Analyst”, DRDC-RDDC-2014-R138.

[47] Lamontagne, L., Bergeron Guyard, A., “A Framework for Building Adaptive Intelligent Virtual Assistants” , Proceedings of Artificial Intelligence and Applications 2014 (AIA2014), Innsbruck, Austria, RDDC Valcartier SL2014-670.

[48] Ross, W., Morris, A., Ulieru, M., Bergeron Guyard, A., “RECON: An Adaptive Human-Machine System for Supporting Intelligence Analysis”, Proceedings of IEEE International Conference on Systems, Man and Cybernetics (IEEE SMC 2013), Manchester, UK, RDDC Valcartier SL 2013-671.

Annex A ISA projects and technologies

This annex provides more information on ISA technologies and projects evaluated over the duration of the research. More detailed information and references can be found in [10].

A.1 DeepQA: Question answering capability

QA is an AI task lying at the intersection of the information retrieval (IR), NLP and ML domains of computer science. QA is the task of automatically answering a question posed in a natural language. In order to generate an answer to a question, a QA algorithm typically uses either a structured database or a collection of non-structured textual documents (called corpus) such as web pages, text files, Word documents, dictionaries, encyclopaedias, etc., as well as one or multiple algorithms from the NLP, IR, and ML domains in order to "understand" the question, find the relevant data in the corpus, identify and extract one or multiple candidate answers from the data, and finally select the correct answer from the list of candidate answers. The performance of a QA system is characterized by the number of correct answers it gives to a set of (initially unknown) questions using a certain corpus. TREC is one of the most frequently used frameworks to evaluate and compare the performance of QA systems.

DeepQA is a QA technology developed by IBM. The development of DeepQA started in 2005 with the primary focus on competing against humans in the Jeopardy! television quiz game. However, DeepQA is a general QA system capable of answering questions from a broad domain posed in a natural language. In order to distinguish between the general QA capability of DeepQA and its capability to play Jeopardy! against human contestants, the version of the DeepQA system tuned specifically to play Jeopardy! was named Watson, while DeepQA usually refers to the overall QA capability of IBM's system 1. Since February 2010, Watson is capable of beating human Jeopardy! contestants on a regular basis. The performance of DeepQA on open-domain questions (i.e., general questions not belonging to a particular knowledge domain) is similar to its performance on the questions from Jeopardy!.

A.2 Siri

Siri is a voice activated personal assistant software for the Apple iPhone 4S, initially developed by Siri Inc., and then by Apple, Inc. based on the results obtained in the CALO project. Siri is capable of answering questions posed in natural language using a mix of artificial intelligence technologies such as speech recognition, automated dialogues and natural language understanding, ML, evidential and probabilistic reasoning, ontology and knowledge representation, planning, reasoning, and service delegation. Siri is also capable of making recommendations and performing actions by delegating requests to a set of Web services. Apple, Inc. claims that the software adapts to its user's individual preferences over time and personalizes results. Siri understands natural language questions or requests and adapts itself to the user's accent. Currently (2012-02-15) Siri understands and speaks three languages: English (United States, United Kingdom, Australia), French (France), and German (Germany). When talking to Siri, the user is not required to change voice or tone, it is also not necessary to recall specific commands for each application or services controlled by Siri. For example, "What will the weather be like tomorrow?" or "Does it look like rain tomorrow?" or even "Will I need an umbrella tomorrow?" will be correctly understood by Siri and the weather application will be called with the necessary options (such as the date and location).

A.3 Wolfram Alpha

Wolfram|Alpha (WA) is a QA technology developed by Wolfram Research, LLC in 2009. Unlike DeepQA that heavily relies on NLP techniques to extract answers from a large corpus of structured and non-structured data (with a stronger emphasis on the non-structured data), WA searches for the answers in structured data (primarily reference tables) and only uses NLP techniques to "understand" the query. WA also generates answers using a variety of mathematics algorithms, primarily coming from Mathematica, a computational software program also developed by Wolfram Research. Similarly to Siri and DeepQA, WA understands queries in natural language; the queries do not need to follow a certain structure. For example, a user can ask "what is the wind speed in Birmingham", "is it windy in Birmingham?", or simply "Birmingham wind speed" and WA will convert them all into the same correct code (e.g., a Mathematica expression).

A.4 Disciple-LTA

Disciple-LTA (for learning, tutor, and assistant) is claimed to be "a personal cognitive assistant to an intelligence analyst that rapidly acquires expertise in intelligence analysis directly from the analyst, including the analyst's prior and tacit knowledge, biases and assumptions". A Disciple-LTA agent is capable of executing three main functions:

- Problem solving assistant: assist intelligence analysts in their daily analytic tasks,
- Learner: capture analyst's expertise, and
- Tutor: tutor new analysts.

As a problem solving assistant, Disciple-LTA is claimed to be capable of generating reasoning trees. In order to generate reasoning trees automatically, Disciple-LTA uses an object ontology (to represent knowledge) and a set of problem solving rules (to decompose a complex problem into a collection of simpler ones and to reconstruct a solution of a complex rule from solutions of simpler problems). The authors of the present document were not capable of finding additional technical detail on the used technologies. Neither the binaries nor the source code of Disciple-LTA have been found in public sources.

A.5 CALO

In 2003, the Defense Advanced Research Projects Agency (DARPA), as part of its Perceptive Assistant that Learns (PAL) program, has awarded SRI International (SRI) a contract to develop a personalized cognitive assistant. SRI has named its new project CALO for Cognitive Agent that Learns and Organizes. The goal of the project was to "create cognitive software systems, that is, systems that can reason, learn from experience, be told what to do, explain what they are doing, reflect on their experience, and respond robustly to surprise". All CALO sub-projects (or CALO modules) are presented next.

A.5.1 CALO iLink

iLink is an application that searches for answers to the questions the user is interested in. The search for an answer is done using the keywords of the question itself. The system searches for an answer to a given question in its own database of questions and answers (Q&A) as well as on the Web. The result of the search is a list of found questions and corresponding answers. The user then analyses the questions and answers returned by the search and decides whether one of them represents a satisfying answer. If not, the system identifies experts capable of answering the user's question and submits the question to them. To be capable of identifying experts whose expertise is relevant to the user's question, iLink regularly monitors the activity of experts in social networks and creates and updates a model of their expertise. In order to find a relevant expert, iLink first analyses the question and assigns a certain topic to it. Then it searches for experts that have expertise on this topic. When a certain expert gives an answer to the user's question, the user receives this answer and is asked to rate the question related to this answer. A favourable rating results in the answer being added to the Q&A repository and the answerer's expertise level for associated topics being increased. iLink uses topic discovery algorithms and classification algorithms to create the expertise model and to determine the best-matching answers for a question.

A.5.2 CALO Express: Personal desktop assistant

CALO Express (CE) is a "personal desktop assistant" software suite with learning capabilities allowing its user to effectively identify relevant information on the workstation. CE scans files located on the user's hard drive, user's email, their calendar, and contact information; this data is then used by the learning module of CE to provide search, presentation and meeting preparation capabilities.

A.5.3 Adept: Task learning capability

Adept is a task learning application that allows its user to teach the machine how to perform certain routine or time-consuming office tasks. In order to teach Adept how to execute a given task, the user simply has to execute the task himself, and, by so doing, to demonstrate to the machine how the task has to be executed. Adept is capable of observing the sequence of actions executed by the user, generalizing this sequence into an executable procedure, and executing this procedure in the future.

A.5.4 WARP: Workflow Activity Recognition and Proactive assistance capability

WARP (for Workflow Activity Recognition and Proactive Assistance) allows tracking the current activity of the user by observing the actions the user executes during his activities. WARP uses a collection of hand-crafted workflows and Logical Hidden Markov Models (HMM) to support the recognition and tracking of workflows in domains that may involve "hundreds to thousands of objects".

A.5.5 Task Assistant

Task Assistant (TA) is a software application allowing its users to create workflows, which are hierarchical task lists, i.e., sequences of tasks and subtasks that have to be accomplished in order to solve a certain problem.

A.5.6 Meeting Assistant

Meeting Assistant (MA) is a software application that allows for the distributed capture, annotation, automatic transcription, and semantic analysis of multiparty meetings.

A.5.7 FOAM: Form Autocompletion Manager

Form Online Analyzer and Manager (FOAM) is an application which helps its users to quickly fill complex electronic forms having many interrelated fields. When the user is filling a form, FOAM presents to the user in real time ranked suggestions for completing different fields of the form based on cross-field models learned from previous form-filling episodes.

A.5.8 C2RSS: News aggregator that learns

C2RSS is an application that gets as input a large set of generic and unclassified RSS feeds and generates a smaller set of personalized RSS feeds for its user. In this application, the user can create topics of interest and read articles from the input RSS feeds that belong to those topics.

A.5.9 Clustering Suite

Clustering Suite is a module of CALO that programmatically implements three clustering algorithms of ML: Latent Dirichlet Allocation, Lingo, and an algorithm referred to as Katz.

A.5.10 Classification Suite

Classification Suite is a module of CALO, which programmatically implements three classification learning algorithms: Transformed Weight-Normalized Complement Naïve Bayes, MaxEnt, and Decision Tree.

A.5.11 PrepPak: Relevant document recommender

PrepPak is an application that assists its users in different office tasks by suggesting in real time relevant documents while the user is working on the task. PrepPak is learning from its user to give better suggestions over time.

A.5.12 Instrumentation and automation

Instrumentation and automation play an essential role in CALO. These capabilities provide the means for CALO modules to "observe" user's actions in different desktop applications instrumentation and initiate activities within the desktop environment (automation).

A.5.13 Contact Management

The Contact Management (CM) application searches on the Web for properties (e.g., contacts and expertise) of a given person in an unstructured data.

A.5.14 Semantic Extraction Suite

The Semantic Extraction (SE) Suite is a collection of general-purpose ML techniques to recognize semantically meaningful entities in text. It is capable of generalizing from examples to recognize semantically meaningful entities such as names, addresses, geographic locations, email signatures, etc.

A.5.15 Emma: Meeting and scheduling manager

Emma is a meeting and scheduling management application that assists its user in the tasks of creating a meeting and time planning, and learns its user's preferences. Meeting requests can be submitted by the user through either a form-based or a restricted natural language interface.

A.5.16 Probabilistic Consistency Engine

The Probabilistic Consistency Engine (PCE) is a framework for probabilistic inference. Part of the CALO challenge resides in combining results from several learning algorithms. This requires taking into account both logical facts that are known to be true or uncertain and possibly conflicting predictions that different learning algorithms may produce. PCE is capable of handling this problem: the learned facts are weighted such that weights reflect degrees of uncertainty. PCE resolves conflicting predictions by constructing a model that attempts to satisfy all assertions or, when that is not possible, discounts the most uncertain facts.

A.5.17 IRIS: Semantic desktop framework

IRIS is a semantic desktop that enables users to create a "personal map" that semantically connects their desktop information objects (e.g., emails, tasks, contacts, files, calendar items). IRIS has ML capabilities that partially automate the process of map creation and update. It provides different map views (e.g., a dashboard view), and context.

A.5.18 Tagomizer: Social bookmarking capability

Tagomizer is a social bookmarking application. It allows a team of users to organize their archive of useful information entities (in particular, web pages) by saving (or in other words, by bookmarking) them in a common knowledge base of the team, and by associating informative labels (aka tags) with those entities.

A.5.19 Open Agent Architecture

Open Agent Architecture (OAA) is a framework for integrating a community of heterogeneous software agents in a distributed environment. Communication and cooperation between agents are mediated by one or more facilitators, which are responsible for matching requests from users or software agents with descriptions of the capabilities of other software agents.

A.5.20 SPARK: Agent framework

SPARK (for SRI Procedural Agent Realization Kit) is an agent framework that enables building software "agent systems that can scale to real world applications, yet retain the clean semantic underpinning of more formal agent frameworks". SPARK provides a general-purpose agent technology for a range of applications that require reactive task execution.

A.6 MALLET: General text mining suite

MALLET is a collection of ML algorithms specifically adapted to textual data and implemented in Java programming language. MALLET is an open-source project developed at the University of Massachusetts, partially with DARPA funding. It can be used for solving problems of NLP, document classification, clustering, topic modelling, information extraction, and several other problems related to text mining.

A.7 MinorThird: Text annotation and semantic extraction suite

MinorThird is a collection of methods for learning how to annotate (i.e., determine classes for words or subsequences of words) and categorize (i.e., determine the topic for the document) texts and to extract entities (i.e., identify and extract from the input text meaningful elements such as names, addresses, etc.).

A.8 Zypr: Cloud-based voice controlled mashup software development framework

Zypr is an online platform for making mashup software applications with voice control. The application programming interface (API) of Zypr gives third-party software developers a platform allowing incorporating voice-activated commands in their software applications for navigation, social media, maps, and calendars. The Zypr API was designed to work across mobile devices, web applications, consumer electronics devices, and automotive information and entertainment systems.

A.9 Trapit: Newswire personalization capability

Trapit is a personal assistant software that uses AI techniques to learn how to understand interests of the user, crawls the web, and personalizes the stream of content to the user based on the learned model of the user's interests.

This page intentionally left blank.

List of symbols/abbreviations/acronyms/initialisms

AI	Artificial Intelligence
DND	Department of National Defence
DRDC	Defence Research and Development Canada
DSTKIM	Director Science and Technology Knowledge and Information Management
HA	Human analysts
HCI	Human Computer Interaction
ISA	Intelligent Software Assistant
iVAC	Intelligence Virtual Analyst Capability
JIMP	Joint Interagency Multi-national and Public
KM	Knowledge Management
ML	Machine Learning
NLP	Natural Language Processing
QA	Question Answering
R&D	Research & Development
SOA	Service Oriented Architecture
SOP	Standard Operating Procedures
S&T	Science and Technology
VA	Virtual analyst

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security markings for the title, abstract and indexing annotation must be entered when the document is Classified or Designated)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g., Centre sponsoring a contractor's report, or tasking agency, are entered in Section 8.) DRDC – Valcartier Research Centre Defence Research and Development Canada 2459 route de la Bravoure Québec (Québec) G3J 1X5 Canada	2a. SECURITY MARKING (Overall security marking of the document including special supplemental markings if applicable.) UNCLASSIFIED 2b. CONTROLLED GOODS (NON-CONTROLLED GOODS) DMC A REVIEW: GCEC DECEMBER 2012

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Intelligence virtual analyst capability : Governing concepts and science and technology roadmap

4. AUTHORS (last name, followed by initials – ranks, titles, etc., not to be used)

Bergeron-Guyard A.; Lavigne V.; Poussart D.; Gouin D; Roy J.

5. DATE OF PUBLICATION (Month and year of publication of document.) December 2014	6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.) 67	6b. NO. OF REFS (Total cited in document.) 48

7. DESCRIPTIVE NOTES (The category of the document, e.g., technical report, technical note or memorandum. If appropriate, enter the type of report, e.g., interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Scientific Report

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

DRDC – Valcartier Research Centre
Defence Research and Development Canada
2459 route de la Bravoure
Québec (Québec) G3J 1X5
Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.) TIF05dz13	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.) DRDC-RDDC-2014-R156	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

This report presents the results of an investigation of Intelligent Software Assistant (ISA) technologies towards the development of an Intelligence Virtual Analyst Capability (iVAC). The idea behind the ISA is to synthesize the current state of Artificial Intelligence research to develop a personalized assistant that organizes information, learns processes, adapts to changing situations, and interactively supports individuals in their tasks in a seamless, intuitive fashion. The iVAC is meant as a computerized software assistant supporting the intelligence analysts in sensemaking tasks, while being ultimately capable of taking on autonomous analytical tasks in concert with other analysts (virtual or human).

To achieve this goal, literature surveys, technology studies, expert workshops and detailed analysis of iVAC sub-capabilities were performed. For every identified iVAC sub-capability, an assessment of the maturity of existing, relevant, applicable science and technology was performed. As a result, key research topics were identified. This document provides a description of the iVAC concept and of the research and development steps required to instantiate it.

Ce rapport présente les résultats d’une analyse des technologies en lien avec l’*Intelligent Software Assistant* (ISA) pour le développement d’un *Intelligence Virtual Analyst Capability* (iVAC). L’idée directrice du ISA est de combiner des approches d’intelligence artificielle pour développer un assistant personnel qui organise l’information, apprend les processus, s’adapte aux situations changeantes et supporte interactivement les individus dans l’exécution de leurs fonctions. L’iVAC est un assistant logiciel capable d’aider l’analyste du renseignement dans ses fonctions d’analyse et devrait ultimement arriver à effectuer des tâches de façon autonome en concert avec d’autres analystes qu’ils soient humain(s) ou virtuel(s).

Pour atteindre cet objectif, des revues de littérature, des études de technologies, des ateliers d’experts, ainsi qu’une analyse détaillée des fonctions du iVAC ont été effectués. Pour chaque fonction de l’iVAC identifiée, une évaluation de la maturité et de l’applicabilité des composantes technologique existantes a été effectuée. Des sujets de recherche clé ont été identifiés. Ce document donne une description des concepts sous-jacents à l’iVAC et identifie les composantes nécessaires à sa réalisation.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g., Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Intelligent Software Assistant; Artificial Intelligence; Intelligence Analysis