

Cognitive Computing for Coalition Situational Understanding

Alun Preece*, Federico Cerutti*, Dave Braines*[†], Supriyo Chakraborty[‡], Mani Srivastava[§]

*Crime and Security Research Institute, Cardiff University, Cardiff, UK; contact email: PreeceAD@cardiff.ac.uk

[†]IBM Emerging Technology Services, Hursley, UK

[‡]IBM TJ Watson Research Center, Hawthorne NY, USA

[§]Department of Electrical Engineering, University of California, Los Angeles, USA

Abstract—The term **cognitive computing (CC)** refers to computer systems that harness multiple techniques from artificial intelligence (AI) and signal processing (SP). Situational understanding (SU) involves creating and reasoning about models of an environment and events. Coalition operations are defined by multiple partners seeking to achieve a common purpose. This paper characterises the SU problem in a coalition operations context — coalition situational understanding (CSU) — in terms of a set of problem attributes. The paper argues that CSU problems require CC system solutions involving a hybrid of AI and SP approaches. The paper outlines some of the architectural choices for CC CSU systems.

I. INTRODUCTION

Cognitive Computing (CC) has emerged in recent years as a term for computational platforms that, in general, draw on the domains of artificial intelligence (AI) and signal processing (SP) [1]. While there is no generally agreed definition of the term, CC is characterised by a systems-oriented view in which multiple AI and/or SP technologies are integrated to address problems generally involving sizeable volumes of data. The archetypal CC system was IBM’s original Watson, that integrated AI techniques including natural language processing, human-computer interaction, and knowledge representation and reasoning to achieve high performance in open-domain question answering [2]. Since the original Watson project, the concept of CC has broadened somewhat, such that we can identify a wider range of AI and SP capabilities occurring in cognitive systems. Some of these are outlined in Section II.

A coalition is commonly defined as a temporary alliance of partners formed to achieve some common purpose or combined action. Examples include multiple agencies participating in a response to a major emergency, a political coalition of multiple parties governing together, or a coalition of military forces engaged in a shared mission. There is often a time-constrained need to form a coalition, to agree common objectives, and to establish policies for command, control, and sharing of information and resources. This paper specifically considers the problem of coalition situational understanding (CSU), where the purpose and actions of a coalition are contingent on gaining and maintaining models of an environment and events. The paper has two goals: (1) to characterise the CSU problem in terms of a set of attributes, and to show that these characteristics necessitate a CC systems approach, i.e., that CSU is fundamentally a CC problem; and (2) to begin

to explore some of the systems architecture choices that need to be made when addressing CSU problem instances using appropriate CC technologies.

The paper is organised as follows: Section II provides a brief overview of key CC technologies; Section IV defines the characteristics of the CSU problem, and identifies corresponding CC approaches in each case; Section V discusses systems architecture issues in addressing CSU problem instances with CC technologies; Section VI makes some initial recommendations and identifies areas for immediate future work.

II. COGNITIVE COMPUTING APPROACHES

As noted in the introduction, there is no commonly accepted definition of CC; instead, the term tends to reflect systems that harness a hybrid set of AI and SP approaches. Moreover, there is no generally agreed architecture for CC systems. There is a sizeable literature on cognitive architectures that attempt to operationalise theories of human cognition in computer systems — of which some of the best known are ACT-R [3] and SOAR [4] — but as yet there is no consensus that cognitive architectures offer the most appropriate means of engineering CC systems. Section V will return to this issue.

This section offers a brief survey of AI and SP approaches commonly considered as being in the scope of CC.¹

Human-computer collaboration (HCC) CC systems commonly perform tasks where human interaction is important. This includes providing human-friendly means of inputting queries/requirements and refining them, as well as providing ways for human users to comprehend, explore, and potentially challenge the output from the CC system. Often, the nature of the tasks performed by a CC system is such that the interaction between humans and the system is in the form of a collaborative activity [5].

Knowledge representation and reasoning (KRR) A CC system often requires the ability to represent explicit models of a problem domain, and to perform inference — in classic AI, a knowledge representation is considered a surrogate for the modelled problem domain, allowing the CC system to reason about the world [6]. A key feature of KRR approaches is that they rely on ontological commitments to map terms in the model to entities in the problem domain.

¹These approaches are listed alphabetically; no relative importance is implied.

Machine learning (ML) CC systems generally require abilities to generate models from data, and dynamically adapt to changes in their problem domain. Significant advances have been made in recent years in ML, particularly in the area of so-called deep learning involving neural networks with multiple hidden layers [7]. Deep learning tends to perform well in problem domains featuring large volumes of training data, and is therefore a good fit for many CC applications.

Multi-agent systems (MAS) MAS technologies address the harnessing of multiple AI components in distributed systems, focusing on issues of inter-agent communication and coordination protocols, as well as algorithms for collective behaviour [8]. MAS approaches are relevant to CC in that they offer a distributed systems architecture for integrating hybrid AI/SP techniques in a principled manner.

Natural language processing (NLP) A key element of IBM’s original Watson [2], NLP is an essential component of CC systems that must process textual data and/or interact with human users in natural language.

Vision and speech processing (VSP) Processing of imagery (still and video) and speech audio data are two specific areas of AI/SP that characterise CC systems. Intelligent handling of speech input has led to the current growth in virtual assistants such as Amazon’s Alexa² and Apple’s Siri³, while computer vision remains one of the key challenges in AI.

The next section introduces the problem of situational understanding and situates it in a coalition operations context.

III. COALITION SITUATIONAL UNDERSTANDING

Situational understanding (SU) is commonly defined as the ‘product of applying analysis and judgment to the unit’s situation awareness to determine the relationships of the factors present and form logical conclusions concerning threats to the force or mission accomplishment, opportunities for mission accomplishment, and gaps in information’ [9]. SU is asserted⁴ as corresponding to *Level 2 situational awareness (SA)* in Endsley’s model — see Figure 1.

UK military doctrine [10] defines *understanding* in the following terms (Figure 2):

Comprehension (Insight) = Situational Awareness and Analysis

Understanding (Foresight) = Comprehension and Judgement

Here, understanding includes *foresight*, i.e., an ability to infer (predict) potential future states, and in this sense corresponds more closely to *Level 3 situational awareness* in Endsley’s terms, while also being compatible with the common definition above that SU involves being able to draw conclusions concerning threats [9]. However, foresight necessarily includes an ability to process and reason about information temporally, spanning approaches in KRR, ML, and VSP.

²<http://web.archive.org/web/20170208020449/https://developer.amazon.com/alexa>

³<https://web.archive.org/web/20160628202009/http://www.apple.com/ios/siri/>

⁴http://web.archive.org/web/20170307202806/https://en.wikipedia.org/wiki/Situation_awareness

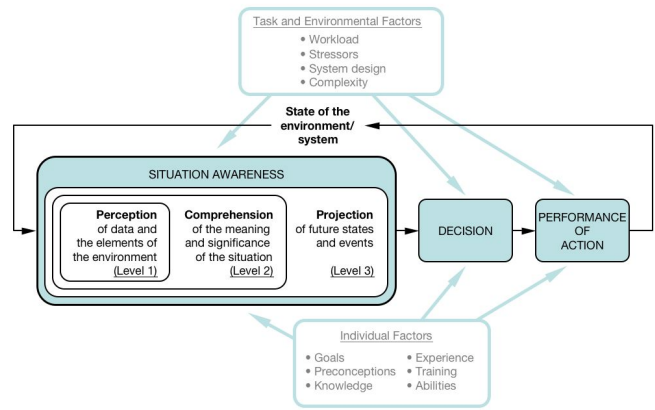


Fig. 1. Endsley’s model of SA (a synthesis⁴)

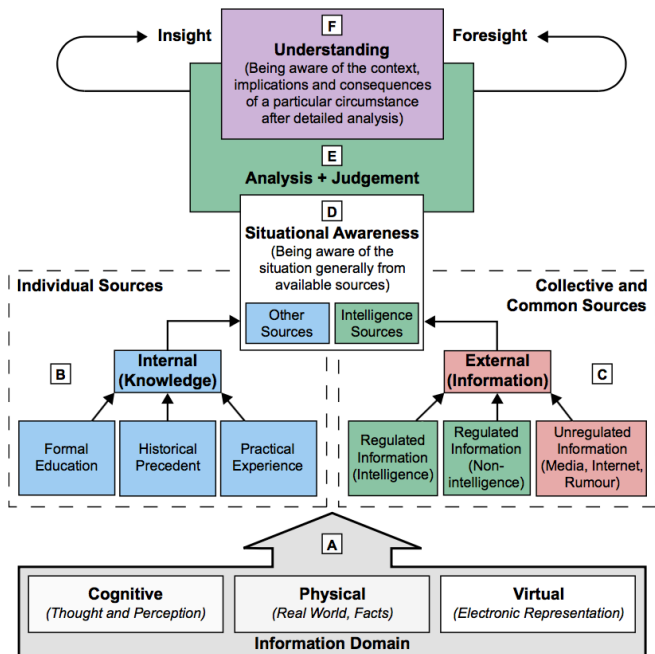


Fig. 2. Sources of understanding [10]

These views of SU are intrinsically linked to information fusion in that they involve the collection and processing of data from multiple environmental sources as input to deriving SA and ultimately SU; the data sources feed into the left of Figure 1 and the bottom of Figure 2. Moreover, the sources will commonly span multiple modalities (for example, imagery, acoustic and textual data [11]) requiring NLP and VSP in addition to ML.

The importance of KRR, and the human user, in the information fusion process is reflected in the *user fusion model* [12] (Figure 3), an extension of the standard JDL Model⁵. The key point here is that the user’s requirements for SU place constraints on all levels of the JDL Model, for example,

⁵More fully, the Joint Directors of Laboratories Data Fusion Information Group (JDL/DFIG) model.

prioritising and valuing of particular kinds of data, objects, contexts and intents [12]. Supporting these user refinement needs in an information fusion system for SU requires a means for the user to interact with elements of the system at both relatively high and low levels, requiring HCC, KRR and NLP.

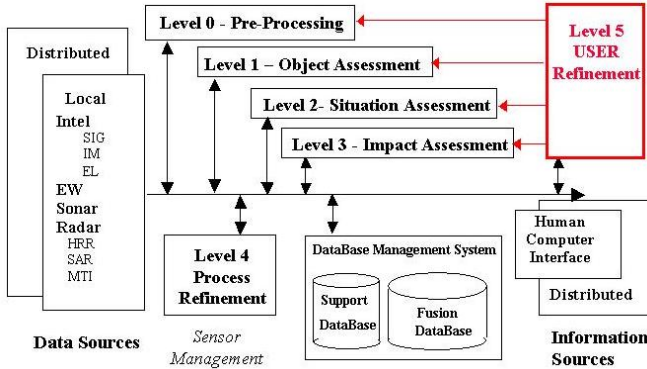


Fig. 3. User fusion model (from [12])

The coalition operations information fusion environment is characterised by a high degree of dynamicity, needs for effective information and asset sharing, and constraints dictated by organisational and mission policies [13]. These features of the coalition environment require techniques from KRR and MAS in the context of an extended formulation of the SU problem, characterised in the next section as *coalition situational understanding* (CSU).

IV. MAPPING CSU PROBLEM ATTRIBUTES TO CC APPROACHES

We identify a set of discrete attributes for the CSU problem as follows.

1) *Level of Understanding*: In terms of the JDL Model (Figure 3), a CSU problem may address relatively high or relatively low levels of understanding, in terms of the kinds of semantic entities and relationships considered. For example, at the relatively low levels a CSU problem may be concerned with only the detection, identification and localization of objects such as vehicles or buildings (JDL Model Levels 1 and 2). At higher levels, a CSU problem would be concerned with determining threats, intent, or anomalies (JDL Model Level 3).

We denote this attribute of the CSU problem as \mathcal{U} ; a CSU problem with a low level of understanding is denoted as $\mathcal{U}^{\mathcal{L}}$ and one with a high level of understanding as $\mathcal{U}^{\mathcal{H}}$. A CSU problem in which both levels of understanding are required is denoted as $\mathcal{U}^{\mathcal{HL}}$.

While it is reasonable to expect a $\mathcal{U}^{\mathcal{L}}$ CSU problem to rely largely on ML, NLP and VSP technologies, by definition (understanding requiring both insight and foresight), $\mathcal{U}^{\mathcal{H}}$ CSU problems require a degree of world-modelling and support for reasoning, hence needing KRR.

2) *Temporal*: As discussed in Section III, *understanding* is associated with *foresight*, requiring an ability to process

information temporally. We denote this attribute of the CSU problem as \mathcal{T} .

We can distinguish between CSU problems that involve relatively short vs relatively long time-scales, denoted by $\mathcal{T}^{\mathcal{L}}$ and $\mathcal{T}^{\mathcal{S}}$ respectively. The latter would be characterised by foresight involving events in the very near future, requiring fine-grained consideration of time (hours, minutes, seconds, or less), while the latter would involve foresight of events in the more distant future, allowing a coarser granularity of time (hours, days, weeks, etc). A problem requiring consideration of time at both scales would be denoted $\mathcal{T}^{\mathcal{LS}}$.

Temporally-oriented ML and VSP techniques abound in the literature, including several deep learning methods such as RNNs and LSTM networks; however, current performance tends to be best on short-term sequences [7]. Activity recognition on the kinds of longer-term sequences important for SU is an active deep learning research area at present. There has also been a large amount of research in KRR temporal reasoning techniques with approaches derived from the event calculus (EC) being particularly well suited to SU reasoning [14]. Modern approaches combining EC with probabilistic logic programming [15] seem particularly promising to address both $\mathcal{T}^{\mathcal{L}}$ and $\mathcal{U}^{\mathcal{H}}$ attributes.

3) *Multimodal Data*: Where a CSU problem involves the processing and fusion of multimodal data we denote this attribute as \mathcal{M} . As described in Section III, the data may be a product of physical sensors, so-called *hard* data, or originate from humans in the form of (usually textual) *soft* data.

Where a CSU problem involves multimodal hard data alone — for example, imagery and acoustic data — we denote it as $\mathcal{M}^{\mathcal{H}}$. Where a CSU problem involves multimodal soft data alone — for example, combining mainstream media with social media reports — we denote this as $\mathcal{M}^{\mathcal{S}}$. Where a CSU problem involves fusion of hard and soft data we denote this as $\mathcal{M}^{\mathcal{HS}}$.

Handling multimodal data is one of the fundamental drivers of the information fusion field, as discussed in Section III. The dominant approach here is to apply SP techniques at JDL Model Levels 0–3, including ML. Where data includes imagery, VSP is an essential component; similarly NLP for text data. $\mathcal{M}^{\mathcal{HS}}$ problems will require both. However, it is increasingly recognised in the fusion domain [16] that KRR plays a key role, at least at the higher JDL levels (including Level 5, Figure 3); nevertheless, we consider this more as an aspect of the 5th attribute (Human-in-the-Loop) below rather than a feature of handling multimodal data *per se*.

4) *Distributed*: As described in Section III, the CSU problem is distributed in nature, partly due to the structure of a coalition. Usually, each coalition partner will have data sources and processing resources, together with constraints on how those data and resources can be shared with other partners. Moreover, these distributed resources (data and processes) will be distributed between the centre and the edge of the network. We denote the distributed attribute of the CSU problem as \mathcal{D} and the aspect of distribution that arises from the structure of the coalition as $\mathcal{D}^{\mathcal{C}}$.

CSU attribute			Primary CC approaches
\mathcal{U}	Level of understanding	High Low High-to-Low	$\mathcal{U}^{\mathcal{H}} \mathcal{U}^{\mathcal{H}}$ $\mathcal{U}^{\mathcal{L}} \mathcal{U}^{\mathcal{L}}$ $\mathcal{U}^{\mathcal{H}\mathcal{L}}$ KRR ML, NLP, VSP KRR, ML, NLP, VSP
\mathcal{T}	Temporal	Long Short Long-to-Short	$\mathcal{T}^{\mathcal{L}}$ $\mathcal{T}^{\mathcal{S}}$ $\mathcal{T}^{\mathcal{L}\mathcal{S}}$ KRR, ML, VSP KRR, ML, VSP KRR, ML, VSP
\mathcal{M}	Multimodal Data	Hard Soft Hard & Soft	$\mathcal{M}^{\mathcal{H}}$ $\mathcal{M}^{\mathcal{S}}$ $\mathcal{M}^{\mathcal{L}\mathcal{S}}$ ML, VSP ML, NLP ML, NLP, VSP
\mathcal{D}	Distributed	Coalition Heterogeneous Coalition & Heterogeneous	$\mathcal{D}^{\mathcal{C}}$ $\mathcal{D}^{\mathcal{H}}$ $\mathcal{D}^{\mathcal{C}\mathcal{H}}$ MAS, ML KRR, MAS KRR, MAS, ML
\mathcal{H}	Human-in-the-Loop	Interpretable Tellable Interpretable & Tellable	$\mathcal{H}^{\mathcal{I}}$ $\mathcal{H}^{\mathcal{T}}$ $\mathcal{H}^{\mathcal{I}\mathcal{T}}$ HCC, KRR, ML HCC, KRR, NLP HCC, KRR, ML, NLP

TABLE I
SUMMARY OF CSU PROBLEM ATTRIBUTES

A second aspect of the CSU problem that makes it distributed in nature arises when different approaches are required to process the data. For example, hard and soft data will generally require different analytic techniques where the CSU problem is characterised above as $\mathcal{M}^{\mathcal{H}\mathcal{S}}$. In other cases, it may be advantageous or required to process data of the same modality using a variety of techniques. In either case, when a CSU problem involves heterogeneous processing methods, we denote this as $\mathcal{D}^{\mathcal{H}}$. In the case where a CSU problem is distributed both due to coalition structure and heterogeneity, we denote this as $\mathcal{D}^{\mathcal{C}\mathcal{H}}$.

The CC approach that directly addresses the \mathcal{D} attribute is MAS, allowing heterogeneous AI/SP components to be dynamically composed across a coalition network, as exemplified (for the $\mathcal{D}^{\mathcal{C}\mathcal{H}}\mathcal{M}^{\mathcal{H}\mathcal{S}}$ problem case) in [17]. KRR also plays a role here, providing ontological commitments to support robust communication and interoperability in a heterogeneous system, $\mathcal{D}^{\mathcal{H}}$. There is a sizeable body of research in distributed ML, including a growing body of work in distributed deep learning; notably, recent work addresses ML over decentralised data, highly appropriate to the CSU $\mathcal{D}^{\mathcal{C}}$ case [18].

5) *Human-in-the-Loop*: The final attribute represents the extent to which the human is in the loop of a CSU problem. As discussed in Section III and illustrated in Figure 3, the human user typically has a number of interaction points with an information fusion system, generally in terms of setting requirements and preferences at the various JDL levels, denoted by \mathcal{H} .

Specifically, where the human provides input to the CSU problem in terms that change the representation and reasoning of the corresponding CSU system — for example, by providing key information currently unknown to the CSU system — we denote this as $\mathcal{H}^{\mathcal{T}}$ and refer to the CSU system as *tellable* by humans.

In cases where the CSU system cannot be a ‘black box’ but must be in some senses transparent to users — for example, able to generate explanations for its output, we denote this as $\mathcal{H}^{\mathcal{I}}$ and refer to the CSU system as *interpretable* by humans.

Finally, where a CSU problem requires both interpretability and provision of key human input, we denote this as $\mathcal{H}^{\mathcal{I}\mathcal{T}}$.

The primary CC approach here is obviously HCC. Communication is also facilitated via KRR and NLP (particularly for processing user NL input). Interpretability of ML approaches is an active area of current research [19]; a particular challenge exists in the CSU problem case $\mathcal{H}^{\mathcal{I}}\mathcal{M}^{\mathcal{H}\mathcal{S}}$ where interpretability is required in the presence of multimodal (hard and soft) data. While techniques are emerging for interpretability of ML over imagery (chiefly heatmapping, e.g., [20]) and text (e.g., [21]), it is an open question as to how to integrate these. Potentially KRR has a role here too, similar to its role in addressing higher-level SU ($\mathcal{U}^{\mathcal{H}}$).

The $\mathcal{H}^{\mathcal{T}}$ attribute is also afforded by KRR: these approaches in AI are seen as having the desirable property that both specific facts and general rules or heuristics can be concisely added to a model without the need (as in ML) for a sizeable number of training examples [22].⁶

Table I summarizes the attributes described above and shows the primary relevant CC approaches required to handle each attribute in a CSU system. The next section discusses architecture choices for a CC CSU system.

V. TOWARDS A SYSTEMS ARCHITECTURE FOR CSU

We begin this discussion by outlining a layered conceptual architecture for a CSU system, illustrated in Figure 4. The lowest layer consists of a collection of data sources (physical sensors and human-generated content), accessible across the coalition, providing hard and soft data. The three upper layers roughly correspond to Levels 0–3 of the JDL Model. For each layer, the figure shows the primary CC techniques employed, though others may be exploited also. The information representation layer uses incoming data streams to learn concepts and model entities together with their relationships at multiple levels of semantic granularity, addressing primarily attributes $\mathcal{M}^{\mathcal{H}\mathcal{S}}\mathcal{T}^{\mathcal{S}}\mathcal{U}^{\mathcal{L}}$. The history of past observations is encoded in

⁶ $\mathcal{H}^{\mathcal{T}}$ is referred to in [22] as ‘teachability’.

these representations, explicitly or implicitly. CC techniques used in this layer are likely to be drawn predominantly from ML, NLP, and VSP. MAS will also play a role in communication of sensor-produced data and coordination among distributed processing services, i.e., in view of the CSU \mathcal{D}^{CH} attribute.

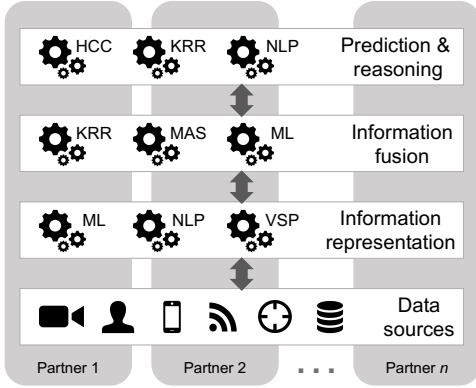


Fig. 4. SU layered model distributed virtually across a coalition

The information fusion layer employs algorithms and techniques developed to perform fusion over concepts and entities derived from the information representation layer. This layer estimates the current state of the world, providing the *insight* (situational awareness) as discussed in Section III and depicted in Figure 2, addressing primarily attributes $\mathcal{T}^{LS}\mathcal{U}^{\mathcal{H}\mathcal{L}}$. While ML, NLP, and VSP approaches play a role at this level also, KRR and MAS have a significant role here.

The prediction and reasoning layer then uses the estimated current state, together with the state space of the models to predict the future state. KRR plays a key role in reasoning about the future state (providing the *foresight* necessary for situational understanding) while HCC, NLP, and interpretable ML approaches address the interaction between this layer and human users, i.e., attributes $\mathcal{H}^{\mathcal{I}\mathcal{T}}\mathcal{U}^{\mathcal{H}}$.

The figure depicts a virtual view of the coalition: all four layers are distributed across the coalition, with MAS and distributed ML techniques addressing attributes \mathcal{D}^{CH} .

In accordance with the User Fusion model (Figure 3), the upper layers in Figure 4 need to be open to humans to provide expert knowledge for reasoning; these layers also need to be open to the human user in terms of being able to generate explanations ($\mathcal{H}^{\mathcal{I}}$) of the insight and foresight generated by the system. There is a bi-directional exchange of information occurring between the different layers: in the upward (feedforward) direction, the inferences at the lower layer act as input for the next higher layer; in the downward (feedback) direction, information is used to adjust the model and algorithm parameters and possibly task the sensors differently. Creating better systems to support CSU necessitates the development of mature models and algorithms that can over a period of time reduce the human intervention and attain greater autonomy, but without entirely replacing human involvement and oversight.

The conceptual distributed architecture shown in Figure 4 most obviously maps to a physical architecture wherein the various models are managed by services and interoperate by message exchange (MAS-style). This approach is exemplified in our previous [17] and current [23], [24] work, managing both the \mathcal{H}^{CH} attribute and also the hybrid nature of a CC system approach.

A key issue in the integration of hybrid approaches is how to combine KRR and SP/ML approaches. Figure 5 shows these as separate subsystems with points of connection. For example, classes forming the output of an ML classifier form part of a KRR ontology, allowing classified instances to be passed upward (feedforward direction in Figure 4) to become part of a represented model and used in reasoning processes. Generally, the output of the classifier will have some degree of uncertainty associated with it, that will also need to be propagated upwards.

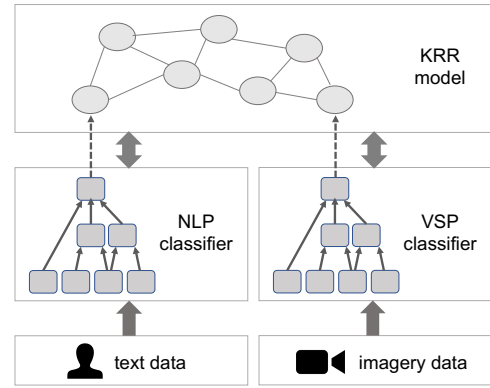


Fig. 5. Interconnection between KRR and SP/ML approaches in a CSU system

The KRR subsystem in Figure 5 also feeds downwards to the SP/ML subsystem in a number of ways. The ontological commitments defined in the KRR model provide a ‘frame of reference’ for the SP/ML subsystem. For example, they define the important entities and relationships of the SU problem, as discussed in Section III (in relation to Figure 3, the user fusion model). Going further, they provide a basis for selection and composition of information processing services. For example, if a problem requires the localisation of particular kinds of semantic objects — say, vehicles — then a CSU system could attempt to look up appropriate pre-trained ML services on the coalition network to classify instances of the target classes using MAS-style matchmaking approaches [8].⁷ Research has also shown that KRR models can be used as prior knowledge to improve the performance of image processing for scene understanding [25].

Integrating the CC approaches as separate-but-interconnected subsystems as in Figure 5 has several advantages:

⁷This idea is illustrated for a particular ‘smart city’ SU problem in [24] and extended to cover the tasking of physical sensors to collect appropriate data in [17].

- it is naturally distributed (necessary for $\mathcal{D}^{\mathcal{C}\mathcal{H}}$);
- it affords flexibility in combining a rich variety of ML and KRR techniques;
- it is currently achievable for the $\mathcal{U}^{\mathcal{H}}$ case given that the current state-of-the-art in ML is weak at learning causal models of the world that support understanding [26].

It should be noted, however, that improving ML techniques to address the third point above is seen as a key future goal in machine learning [27] and deep learning [7] research. Progress in this area — in particular on approaches that seek to utilise vector-space representations in place of rule-based manipulation of symbolic expressions — would offer a more uniform means of integrating ML and KRR, though with three important caveats for CSU:

- for $\mathcal{H}^{\mathcal{I}}$, that the interpretability problem for such a unified approach would also need to be solved;
- for $\mathcal{H}^{\mathcal{T}}$, that the unified learning framework should be capable of learning from small amounts of data/human input (a similar point is made in [26]); and
- for $\mathcal{D}^{\mathcal{C}\mathcal{H}}$, that there should be a means of distributing the elements of the system across a coalition (a problem that has so far been little addressed, though [28] considers steps towards it).

VI. CONCLUSION AND FUTURE WORK

We have defined the CSU problem and shown that it is by nature a CC problem, requiring a hybrid set of AI/SP techniques in a distributed context. We have begun to consider an appropriate systems architecture for addressing CSU problem instances and have noted that, in the near term, a 4-layer model including dual KRR and ML subsystems appears to be a natural approach. However, anticipated advances in ML suggests a more unified approach might become feasible in the future, though for CSU open problems of supporting interpretability, tellability, and distributed SU processes across a coalition need be addressed.

ACKNOWLEDGEMENT

This research was sponsored by the U.S. Army Research Laboratory and the UK Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the UK Ministry of Defence or the UK Government. The U.S. and UK Governments are authorised to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] J. Kelly and S. Hamm, *Smart Machines: IBM's Watson and the Era of Cognitive Computing*. Columbia Business School Publishing, 2013.
- [2] D. Ferrucci, E. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. A. Kalyanpur, A. Lally, J. W. Murdock, E. Nyberg, J. Prager, N. Schlaefer, and C. Welty, "Building watson: An overview of the deepqa project," *AI Magazine*, vol. 31, no. 1, pp. 59–79, 2010.
- [3] J. R. Anderson, *How Can the Human Mind Occur in the Physical Universe*. Oxford, 2007.
- [4] J. E. Laird, *The Soar Cognitive Architecture*. MIT Press, 2012.
- [5] L. Terveen, "Overview of human-computer collaboration," *Knowledge-Based Systems*, vol. 8(2), pp. 67–81, 1995.
- [6] R. Davis, H. Shrobe, and P. Szolovits, "What is a knowledge representation?" *AI Magazine*, vol. 14, no. 1, pp. 17–33, 1993.
- [7] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 5 2015.
- [8] M. Wooldridge, *An Introduction to MultiAgent Systems*. Wiley, 2009.
- [9] B. C. Dostal, "Enhancing situational understanding through the employment of unmanned aerial vehicles," *Army Transformation Taking Shape... Interim Brigade Combat Team Newsletter*, no. 01–18, 2007.
- [10] "Understanding: Joint Doctrine Publication 04 (JDP 04)," Ministry of Defence, UK, 2010.
- [11] D. Lahat, T. Adali, and C. Jutten, "Multimodal data fusion: An overview of methods, challenges, and prospects?" *Proceedings of the IEEE*, vol. 103, no. 9, 2015.
- [12] E. Blasch, "Level 5 (user refinement) issues supporting information fusion management," in *9th International Conference on Information Fusion*, 2006.
- [13] T. Pham, G. Cirincione, D. Verma, and G. Pearson, "Intelligence, surveillance, and reconnaissance fusion for coalition operations," in *Proc 11th International Conference on Information Fusion*, 2008.
- [14] R. Kowalski and M. Sergot, "A logic-based calculus of events," *New Generation Computing*, vol. 4, no. 1, pp. 67–96, 1986.
- [15] A. Skarlatidis, A. Artikis, J. Filippou, and G. Paliouras, "A probabilistic logic programming event calculus," *Theory and Practice of Logic Programming*, vol. 15, no. 2, pp. 213–245, 2015.
- [16] D. L. Hall and J. M. Jordan, *Human-Centered Information Fusion*. Artech House, 2010.
- [17] A. Preece, D. Braines, D. Pizzocaro, and C. Parizas, "Human-machine conversations to support multi-agency missions," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 18, no. 1, pp. 75–84, 2014.
- [18] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017*, 2017.
- [19] Z. C. Lipton, "The mythos of model interpretability," in *2016 ICML Workshop on Human Interpretability in Machine Learning (WHI 2016)*, 2017.
- [20] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep taylor decomposition," *Pattern Recognition*, vol. 65, pp. 211–222, 2017.
- [21] T. Lei, R. Barzilay, and T. Jaakkola, "Rationalizing neural predictions," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016, pp. 107–117.
- [22] R. Guha, "Towards a model theory for distributed representations," in *Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches: Papers from the 2015 AAAI Spring Symposium*, 2015.
- [23] D. Braines, A. Thomas, L. Kaplan, M. Sensoy, M. Ivanovska, A. Preece, and F. Cerutti, "Human-in-the-loop situational understanding via subjective bayesian networks," in *Fifth International Workshop on Graph Structures for Knowledge Representation and Reasoning (GKR 2017)*, 2017.
- [24] A. Nottle, S. Quintana-Amate, D. Harborne, M. Alzantot, D. Braines, R. Tomsett, L. Kaplan, M. Srivastava, S. Chakraborty, and A. Preece, "Distributed opportunistic sensing and fusion for traffic congestion detection," in *First International Workshop on Distributed Analytics InfraStructure and Algorithms for Multi-Organization Federations*, 2017.
- [25] K. Chen, Y. Lai, Y.-X. Wu, R. R. Martin, and S.-M. Hu, "Automatic semantic modeling of indoor scenes from low-quality rgb-d data using contextual information," *ACM Transactions on Graphics*, vol. 33, no. 6, pp. 133–149, 2014.
- [26] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," *CoRR*, vol. abs/1604.00289, 2016.
- [27] L. Bottou, "From machine learning to machine reasoning," *Machine Learning*, vol. 94, no. 2, pp. 133–149, 2014.
- [28] D. Verma, G. Bent, and I. Taylor, "Towards a distributed federated brain architecture using cognitive iot devices," in *The Ninth International Conference on Advanced Cognitive Technologies and Applications*, 2017.