Conversational Sensemaking

Alun Preece[a], Will Webberley[a], Dave Braines[b]

[a]School of Computer Science and Informatics, Cardiff University, Cardiff, UK
[b]Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, UK

ABSTRACT

Recent advances in natural language question-answering systems and context-aware mobile apps create opportunities for improved sensemaking in a tactical setting. Users equipped with mobile devices act as both sensors (able to acquire information) and effectors (able to act in situ), operating alone or in collectives. The currently-dominant technical approaches follow either a pull model (e.g. Apple's Siri or IBM's Watson which respond to users' natural language queries) or a push model (e.g. Google's Now which sends notifications to a user based on their context). There is growing recognition that users need more flexible styles of conversational interaction, where they are able to freely ask or tell, be asked or told, seek explanations and clarifications. Ideally such conversations should involve a mix of human and machine agents, able to collaborate in collective sensemaking activities with as few barriers as possible. Desirable capabilities include adding new knowledge, collaboratively building models, invoking specific services, and drawing inferences. As a step towards this goal, we collect evidence from a number of recent pilot studies including natural experiments (e.g. situation awareness in the context of organised protests) and synthetic experiments (e.g. human and machine agents collaborating in information seeking and spot reporting). We identify some principles and areas of future research for "conversational sensemaking".

Keywords: sensemaking; situation awareness; crowdsourcing; natural language; controlled natural language

1. INTRODUCTION

Pirelli and Card's seminal 2005 paper[1] defined the sensemaking process as a set of interconnected loops, shown in Figure 1. The sensemaking process[*] is divided into two parts: a *foraging loop* in which data is gathered from the external environment and assembled into a body of evidence, and a *sensemaking loop* in which schematised evidence is connected to hypotheses and cases are built to inform decision making. The loops denote feedback in the respective parts of the process; further feedback loops exist between each pair of successive steps in the process. The progression of the process from left to right and bottom to top represents increasing effort on the part of the analyst, and increasing structure in the information artefacts created.

Our recent work has focussed on supporting intelligence, surveillance and reconnaissance tasks by increasing the degree of automation in setting up "data to decision" (D2D) pipelines,[2] with an overall aim of improving operational agility and empowering human actors at the network edge.[3] We observed that there is a tendency to view D2D pipelines as uni-directional, flowing from data sources such as sensing assets to human decision makers, and to downplay the importance of the bi-directional aspects of the pipelines, which are where some of the greatest opportunities exist for improving agility and empowering edge users. For example, we developed approaches to automating the assignment of sensing assets to mission tasks in a highly dynamic way, by following the backward chain from "decision (requirement) to data (sources)".[4] By leading a user to focus upon *what* information they need rather than *how* to obtain that information in terms of specific sensing assets, we can open up the search space of available assets, which helps make the most efficient use of available resources. Opening the search space as widely as possible is particularly important in a coalition setting where assets are shared among multiple partners. This approach of separating *what* the user wants from *how* to provide it reduces the requirements on users to know about the effectiveness or availability of particular kinds of assets. This is important because studies have shown that knowledge of what assets can best serve ISR tasks is not generally possessed in great depth by edge users.[5]

Send correspondence to A.Preece, E-mail: PreeceAD@cardiff.ac.uk, Tel: +44 (0) 29 2087 4653

[*]Pirolli and Card use "sense-making" but here we prefer "sensemaking".

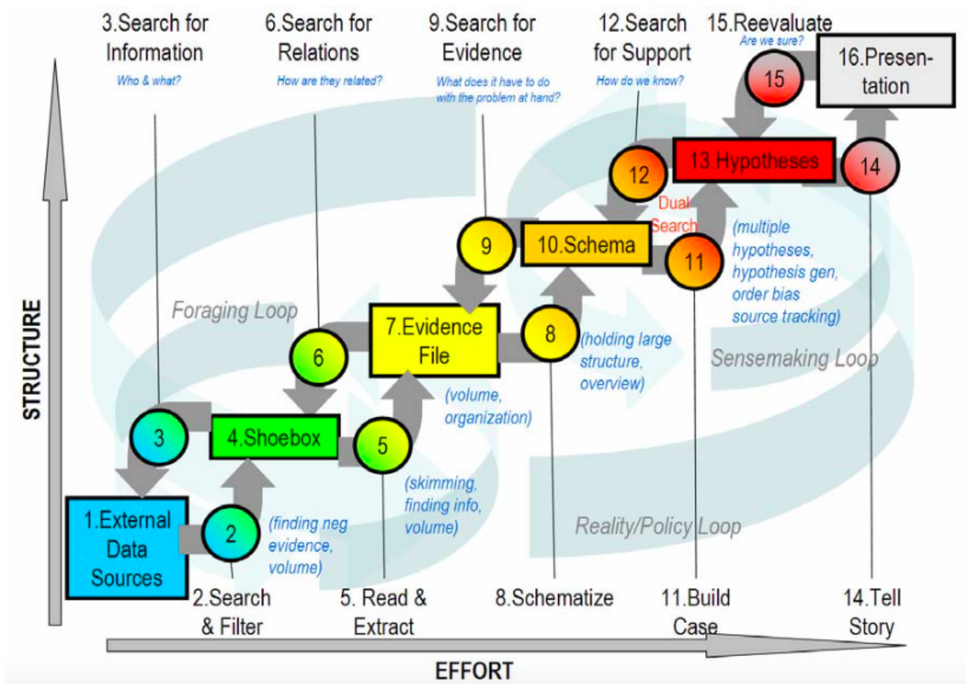

Figure 1. Pirolli and Card's model of the sensemaking loop for intelligence analysis[1]

We hypothesise that, a result of this increased automation and agility, the user becomes a more active participant in the process, able to ask the system for information as well as receive it. Moreover, we observed that users, especially those at the network edge, are often themselves sources of information — they are human sensors.[6] These observations led us to examine the role of *conversational* modes of interaction in D2D processes, looking at a number of illustrative use cases including spot reporting, crowdsourcing, asset tasking, and fusion of soft and hard information. Inspired by recent developments in natural language "intelligent assistants" such as Apple's Siri[†] and Microsoft's Cortana[‡], we endeavoured to explore areas where human-machine conversational interactions might further improve operational agility and reduce training overheads for users at or near the tactical edge. To this end, we proposed a protocol[7] to support free-flowing conversations among human and machine agents that supports both natural language (NL) input and output, as well as the use of controlled natural language (CNL) as a uniform information representation that is both machine-processable and human-consumable.[8] The protocol allows conversational sequences to be initiated by either human or machine agents, for either party to use NL or CNL, for any agent to issue queries (*ask*) or provide information (*tell*) to another agent, and for any agent to seek an explanation (*why*) for some piece of information.

Viewing a conversation in this way, as a sequence of messages of various kinds (ask, tell, why), is consistent both with linguistic "speech act" theories[9] and approaches to managing communication between software agents.[10] However, an alternative view of conversation as an element of cognition sees it as a process of co-constructing shared informational artefacts.[11] These views are compatible and we will argue in this paper that they reflect the different roles conversational processes can play in the sensemaking process as framed by Pirolli and Card. We observe that the backward transitions between each successive pair of steps — detailed along the top of Figure 1 — take the form of queries (*ask*), while the forward transitions — detailed along the bottom of the figure — impart information (*tell*). This is consistent with the speech act view of conversational actions. Viewed at a more macro scale, the process of adding structure that comes with increasing effort, moving upwards and rightwards in the figure, corresponds to the view of conversation as co-construction of shared knowledge. As the conversation progresses, the potential for increased structure, increased awareness of the different conversational partners' "views of the world" (conceptual models), and increased alignment of the agents to the context of the

[†]https://www.apple.com/ios/siri/
[‡]http://www.windowsphone.com/cortana

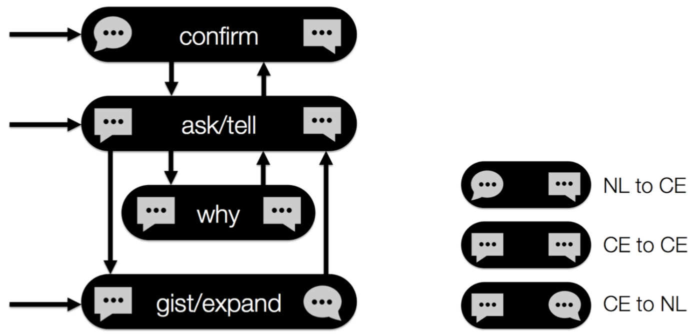

Figure 2. Model for human-machine conversational interactions

current task can occur.

The remainder of this paper seeks to develop this relationship between conversational interaction and the sensemaking loop, with a practical aim of beginning to identify some areas where our approach to supporting human-machine conversations may assist in the aspects of sensemaking that are situated closest to the tactical edge. We start in Section 2 with a brief review of our conversational protocol and some examples of its use. Section 3 then examines the role of conversational interactions in Pirolli and Card's foraging loop, and Section 4 does the same for the sensemaking loop. Both sections draw on experiments we have performed to date, in both natural and synthetic settings. Finally, Section 5 provides some further discussion and points towards future work.

2. HUMAN-MACHINE CONVERSATIONAL MODEL

The main features of our conversation model[7] are summarised in Figure 2, which is adapted from a previously-published version.[2] The model differentiates between natural language (NL) and controlled natural language (CNL) content. NL content is represented as text but its origin can be speech, directly-typed input, or material gathered from external sources including documents, SMS, or social media. It can also be generated by machine agents. CNL content uses ITA Controlled English (CE) to define its syntax and semantics.[12] Some details and examples of CE are given later in this section.

To be processable by a machine agent, NL content must first be interpreted into CNL, and we refer to this conversational sequence as a *confirm* interaction. When the NL input originates from a human user, this interaction typically involves showing the user a piece of CE generated via natural language processing (NLP) of their input, and asking them explicitly to confirm or edit it. When the interpretation has low ambiguity or does not originate from a user — for example, if it comes from an external social media source — the explicit confirmation step may be skipped. In such cases, the system will still retain the original NL in case it needs to be revisited later.

While NL can be ambiguous to the point where it is not always possible to determine whether a sentence is a query or states a fact[§], no ambiguity remains once the sentence is interpreted as CE. From this point agents can engage in query-response exchanges which we refer to in the model as *ask/tell* interactions. Such an exchange can be as simple as a single *tell* — it does not need to involve an *ask*.

A key feature of the approach is to allow any agent to obtain an explanation, justification or provenance for a piece of information. For example, if told a fact, an agent may seek an explanation of how that fact was obtained or inferred. This is the purpose of the *why* interaction. CE has a specific syntax for rationale which begins with the keyword `because`.

As noted above, a machine agent may generate NL purely for human convenience purposes — for example, to make some output more easily readable when the human user is engaged in tasks that require digestible

[§]See the famous Abbott and Costello "Who's on first?" comedy routine, or the title of the British television programme "Doctor Who".

information — or when the device form factor is not suitable for the consumption of longer pieces of text. (CE is intended to be human-readable but it is often rather verbose and can be especially difficult to comprehend on mobile devices where smaller screens favour shorter messages.) We refer to such generated NL as "gist" and the protocol requires any agent that issues gist to be able to provide a full CE expansion of this if required by the receiver: this is shown as the *gist/expand* interaction.

In our approach, ITA Controlled English (CE) is the sole formal knowledge and information representation: the machine agents process this formalism directly, rather than translating it to some other form such as the Resource Description Framework (RDF).[13] CE offers approximately the same expressivity in terms of representing models (ontologies) and facts as the W3C's Web Ontology Language (OWL),[14] and includes a rule language with similar inferential power to the Semantic Web Rule Language.[15] For illustration, an example CE model definition is shown below.

```
conceptualise a ~ protest ~ P that
  is an event.

conceptualise an ~ event ~ E that
  has the time ST as ~ start time ~ and
  has the time ET as ~ end time ~ and
  ~ involves ~ the agent A and
  ~ is located at ~ the place P.
```

A `conceptualise` sentence defines a new concept in a CE model (ontology). New terms in the model are introduced between the tilde (~) symbols. The example defines the concept `protest` as being a child of the parent concept `event`, and having properties `start time` and `end time`. We can also define relationships between concepts, for example, the relationship `involves` relates an `event` to an `agent` (which can be a human, machine agent, or organisation) and the relationship `is located at` relates an `event` to a `place`.

CE facts (instance data) are defined via a similar syntax. The example below shows the CE that creates an instance of the concept `protest`.

```
there is a protest named 'Central Square protest' that
  has the time 4-9-2014-12:00 as start time and
  involves the group 'Blue Group' and
  is located at the place 'Central Square'.
```

This instance is named `Central Square protest`. It has `4-9-2014-12:00` as the value of its `time` property, and an `involves` relationship with an instance of the concept `group` (as defined above). The `group` instance is named `Blue Group`. The `protest` instance `Central Square protest` also has an `is located at` relationship with an instance of the concept `place`, named `Central Square`.

Our current natural language processing (NLP) approach to interpreting NL as CE in confirm interactions uses a bag-of-words technique to map elements of NL sentences to CE models.[2] This allows us to generate CE for relatively simple NL input, as discussed in the next section. For the generation of gist NL from CE we use a simple template-based approach that preserves the mapping from the original CE to the gist form, to allow expansion back into full CE if requested by the recipient.

3. CONVERSATIONAL FORAGING

The ability to interpret NL statements into CE to support subsequent machine analysis and fusion offers considerable potential in terms of supporting foraging activities. One of our original ISR use cases was to enable direct submission of in situ reports via confirm interactions in the conversational protocol.[2] The original approach was to make this facility available to users through a dedicated "text chat" app on a mobile device, called Moira (Mobile Intelligence Reporting Agent). In more recent work we have experimented with allowing a user to post

Figure 3. Tweeted spot report as natural language input to the Moira agent

brief reports via social media, specifically Twitter. (The rationale for this is that Twitter tends to be the dominant social medium for eyewitness reporting of "breaking events".[16]) The Moira agent can be configured either to "follow" specific accounts — public or private — or to process data collected in bulk via Twitter's APIs.[17]

An example of a natural language report posted on Twitter is shown in Figure 3. This report was posted from a private account used by members of our research team in experiments using Twitter for in situ reporting.[17] The equivalent CE form of this report was shown in the previous section as the example instance of a `protest`. To generate the CE, our Moira agent needs to assign identifiers such as `Central Square protest` to entities identified via the bag-of-words NLP approach. Some elements such as the name of the group involved and the location of the event are extracted from the text; other elements such as the start time can be inferred from metadata such as the timestamp of the tweet.

When the dedicated Moira app is used, we tend to give the user an opportunity to manually confirm the generated CE, or to edit it if they require. For input via Twitter it doesn't make sense to do this as the user will typically be a member of the public who would find such an interaction very confusing (even if it were possible to fit a typical CE sentence into the 140 characters of a tweet). To explore the extent to which our Moira agent is in principle able to extract usable information from crowdsourced data we conducted an experiment in which 20 users were asked to describe (in NL) images they were shown of common street scenes involving emergency services such as police, firefighters, and medics.[2] The participants had no training in writing CE and were encouraged to use simple, normal English. They submitted their NL input to a version of the Moira agent that had been rapidly constructed over a 2 week period to perform bag-of-words NLP on sentences relating to the domain of the images. The main result was that an average of 2.3 usable CE elements (recognised concepts or relationships) were extracted from each input by the untrained users. We believe this result shows that our approach is promising for information foraging via crowdsourcing or open social media. Further experiments are ongoing.

In addition to supporting crowdsourcing and in situ reporting, we have explored the use of rapid fact acquisition to build CE knowledge bases of background information such as people, organisations, and places, useful as a means of connecting or contextualising other information. This background knowledge is also valuable in NLP, performing an important role in named entity recognition as illustrated by the (fictitious) examples "Central Square" and "Blue Group" in Figure 3. During a series of studies in 2014 focussed on monitoring community reactions to disruptive events in the South Wales region of the UK, we constructed models and fact sets for notable local places, organisations, and public individuals including politicians and journalists[17]¶. This allowed our Moira agent to perform question answering via the conversational interface on this body of rapidly-sourced background data, as illustrated in Figure 4.

Essentially, we view all these activities as illustrative of how the conversational approach can support the foraging loop in sensemaking. We regard part of our CE knowledge base as the "shoebox" shown in Figure 1. This contains both NL and CE input from external sources, including crowdsourced data, tweets, and rapidly-acquired background information. It is perhaps different from the traditional sensemaking "shoebox" in that use of CE for interpreted data and metadata gives the information a degree of structure from the moment it arrives in the system. However, it is important to emphasise that this is typically "low level" structure and interpretation — we are not attempting to "make sense" of the data at this point in the ways the later stages of the process do. Put another way, our CE shoebox covers basic facts like *who*, *what*, *when* and *where*, but does not yet attempt to capture anything about *how* or *why*. We should also note that this point that the contents of the shoebox are

¶For this we used public sources such as Tweetminister, http://tweetminster.co.uk, Muckrack, http://muckrack.com, and Wikipedia/Dbpedia.

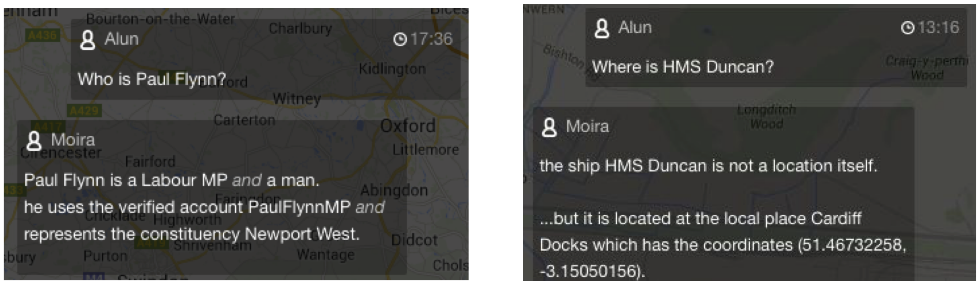

Figure 4. An example of question answering from foraged data using Moira

not limited to soft data such as that acquired from humans or documents. Our earlier work examined the use of CE to support linking and fusion of hard and soft data, including processed data from physical sensors.[18]

The progression from shoebox to evidence file involves some degree of processing of the data in the former. Some of this involves linking, summarisation and inference. One of the advantages of casting acquired information into a CE knowledge base is that we automatically build a graph of linked data, around common entities such as people, organisations, and places, and the relationships defined in the CE models between these and other entities. Note that the CE models should not be considered static: they evolve as new relationships and concepts are found. This can be done semi-automatically using techniques like automatic term recognition to propose new concept names;[19] it can also be done manually by querying and filtering the data.

As an example of the latter, part of our focus in the 2014 South Wales region study was to explore the value of social media as a sensor to detect public protest events. The details of this work are reported elsewhere[17] but a key result was that Twitter proved a valuable source of data for real-time situational awareness in this context. The example elements of our CE model shown in Section 2 (`event` and `protest`) are drawn from this work. However, it quickly became apparent to us that an important distinction needed to be drawn between "expected" and "unexpected" protests: the former are ones that the police and local authorities will be aware of and will have planned how to manage them, while the latter are ones that may be more disruptive in that they have an element of surprise. We therefore extended our model to create two sub-classes of `protest` — `expected protest` and `unexpected protest` — and wrote CE rules to infer that a detected `protest` was an `expected protest` if it had particular properties such as occurring at a particular place or time. This allowed us to enrich the shoebox with useful additional context, and facilitate the moving of some of the acquired data to the evidence file to support higher-level sensemaking.

4. CONVERSATIONAL SENSEMAKING

In Pirolli and Card's model (Figure 1), the first step of the sensemaking loop as it follows from the end of the foraging loop involves schematising the evidence, adding further structure and interpretation. Our approach in using CE introduces a degree of schematisation "for free" at the lower levels of the process, and at the end of the previous section we discussed how the approach allows progressive enrichment of the model, to better contextualise evidential information. We therefore do not see a hard boundary between the two loops, more a gradual refinement and enhancement of the knowledge base. Pirolli and Card observed a tight coupling between the schematising and hypothesising stages of the model; indeed, the latter frequently feeds back to the earlier stages of the process, as hypotheses trigger requests for additional information. This concurs with our previous observations on the bi-directionality of the data to decision chain, discussed in Section 1.

In the previous discussion of foraging we avoided mention of the hypotheses that framed our work. For example, the South Wales studies were partly framed by social science motivations to understand the impacts of various events that were happening between the spring and autumn of 2014, including: (i) community reactions to media reports of three local youths who had reportedly travelled to Syria in early 2014 to join the Islamic State organisation[||] and (ii) the region's hosting of the international NATO Summit in September 2014, involving major local disruption due to increased security, and threats of significant protest by a broad spectrum of groups[**].

[||]http://www.bbc.co.uk/news/uk-wales-south-east-wales-28116575
[**]http://www.bbc.co.uk/news/uk-wales-28953445

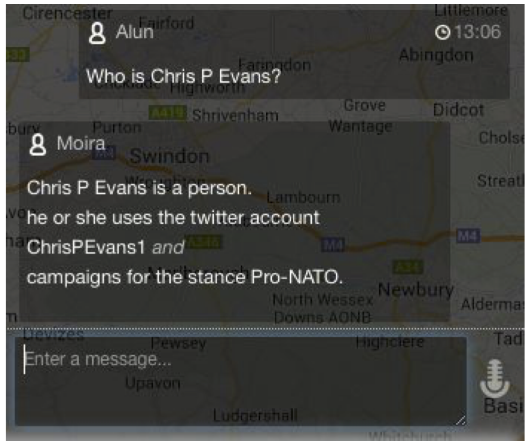

Figure 5. Use of the `stance` relationship in a conversation with Moira[17]

The requirement to make sense of these unfolding situations led us to frame much of our foraging activity around event detection from social media — especially protests and crowd mobilisation — as well as directing us to gather background knowledge on key likely actors (including politicians and journalists) and places of interest.

Pushing this contextual framing down into the foraging loop means that we can attach metadata to collected data that will be useful in the later stages of interpretation. In the case of social media data acquired from Twitter, as well as the textual content of the tweets we are also interested in the context: who is saying it, where and when. Whether a tweet issues from a politician's account, a journalist, an activist group, or a member of the public can be important in making sense of the signals available from social media.[17] A key feature of our approach is that all of this information — data and metadata — is represented in a uniform CE knowledge base and information architecture, even though processing data vs metadata may use different techniques.

Another example of the rapid loop-closing from hypothesis to data collection was a real-time realisation that we needed to capture additional contextual information regarding the stance of various key tweeters. When examining real-time reports from Twitter on events in relation to the NATO Summit in September 2014, a social scientist colleague offered a hypothesis that particular individuals were key influencers in shaping public reaction — positive and negative — to the event. The team had a brief discussion and settled on `stance` as an acceptable name for a CE concept to model this, with `pro-NATO` and `anti-NATO` as instances. Adding these elements to the model and knowledge base took a matter of minutes, including the time to discuss and agree the model. The non-technical social scientists were also deeply involved in the process, both in terms of deciding the names to be used and being showing the formal representation (in CE). This depth of engagement would not be expected with more traditional machine-friendly formats; we hypothesise that eventually these individuals would be able to directly make such extensions to the knowledge base without support from a technical team due to the human-friendly format. Telling the system that particular tweeters had one or other stance allowed us to enrich the knowledge graph and issue queries in terms of events, influencers, and stance. A simple example of the use of the `stance` concept via the Moira conversational interface is shown in Figure 5. (The example query is in relation to the official Twitter account of an individual who is a local political public figure.)

While proven useful in a number of pilot studies, our uses of the conversational approach in sensemaking have focussed on relatively simple cases so far. Having elements such as `stance` has allowed us to progress from simple who/what/when/where queries to begin to look at issues of why and how. We have ongoing work looking at situational understanding in the context of reactions to major crime events and are starting to develop richer models at the upper end of Pirolli and Card's process, but these are not yet formalised in CE.[20] Our present focus in this respect is to gain a better understanding of the analytic processes from a social science standpoint before encoding these into our knowledge representation architecture with a view to assisting the analysts with these sensemaking tasks. Nevertheless, concurrent research in applying CE-based representation and reasoning techniques to intelligence analysis tasks has shown promising areas where human-machine collaboration can assist human cognitive processes.[21]

The final step in Pirolli and Card's model is that of presenting a case, using the connected hypotheses, evidence, and data to tell a story to inform decision making. Our whole approach is founded on the human-consumability of the knowledge representation, and ability to ask *why* to uncover rationale for any inference, statement or connection in the knowledge base. These are useful features to support the story-telling step, but they are not enough: we also need to apply narrative framings to the assembled body of knowledge. In this context we have examined the use of techniques including comic strips and multiple-act story structures represented as CE metadata.[22] The concept of a `story` sits above the domain of interest (people, places, events, etc) and organises the domain information into an episodic sequence (`preface`, `act one`, `act two`, etc) with associated actors and key events mapped to the domain-level model and instances. As we have seen, the extensible and dynamic features of CE allow it to be used in any context and to enrich any domain-level information with a narrative. At the same time, the conversational ability to ask questions, including *why?*, permits the consumer of the narrative to reveal detail that the higher-level story deliberately omits. This may uncover unanswerable questions or provoke a user to tell the system something that has not yet been taken into account, causing feedback down Pirolli and Cards chain, and becomes new information in the hybrid human-machine environment; available to machine agents for inference and human readers for general awareness and potential insight or inspiration.

5. DISCUSSION AND CONCLUSION

In this paper we have endeavoured to draw parallels between the sensemaking process, conceptualised by Pirolli and Card as an interconnected set of feedback loops, and patterns of conversational interaction where machines are increasingly playing assistive roles. We have focussed mainly on interactions that attempt to connect what machines are good at — for example, large scale data collection and pattern identification — with what humans are best at: interpreting and hypothesising, while also promoting transparency and trust by allowing a user always to obtain the rationale for machine-produced information.

A major area where conversation has enormous value in human society is in enabling debate and argument, revealing areas of conflict and differing opinion, and helping to explore and (sometimes but not always) reconcile these differences. The importance of machine support for argumentation in sensemaking processes has been highlighted[23] and constitutes an important area for future work at the intersection of conversational and sensemaking systems.

The focus in this paper has been on textual communication through conversation to support human-computer collaboration but we acknowledge that users will often wish to consume images in addition to, or in place of, text.[24] In some of our previous work we experimented with the use of "visual gist" combining text and machine-generated images in both *confirm* and *gist/expand* interactions (Figure 2)[7] and we feel this remains a valuable area for further study.

Our current work focuses on evaluating the Moira agent in the context of experiments with human subjects performing collaborative crowdsourcing and sensemaking tasks in situ. Early results confirm that, as with our earlier experiments using still images of emergency response scenes referenced in Section 3, untrained participants are quickly able to deliver actionable information to the machine, and gain assistance from software in building situational awareness. Going forward, we intend to conduct multiple experiments with a mix of human and machine-based sensing, and to collaborate with researchers in the area of software-assisted argumentation to help in making sense of conflicting and uncertain information.

Acknowledgements

This research was sponsored by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. 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 US Army Research Laboratory, the US Government, the UK Ministry of Defence or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

Many of the examples in this paper were informed by collaborative work between the authors and members of Cardiff Universities Police Science Institute, http://www.upsi.org.uk. We especially thank Martin Innes, Colin Roberts, and Sarah Tucker for their valuable insights on policing and community reaction.

REFERENCES

1. P. Pirolli and S. Card, "The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis," in *Proceedings of International Conference on Intelligence Analysis*, 2005.
2. A. Preece, C. Gwilliams, C. Parizas, D. Pizzocaro, J. Z. Bakdash, and D. Braines, "Conversational sensing," in *Proc Next-Generation Analyst II (SPIE Vol 9122)*, SPIE, 2014.
3. D. S. Alberts and R. E. Hayes, *Power to the Edge: Command and Control in the Information Age*, CCRP, 2003.
4. A. Preece, T. Norman, G. de Mel, D. Pizzocaro, M. Sensoy, and T. Pham, "Agilely assigning sensing assets to mission tasks in a coalition context," *IEEE Intelligent Systems* **Jan/Feb**, pp. 57–63, 2013.
5. J. Z. Bakdash, D. Pizzocaro, and A. Preece, "Human factors in intelligence, surveillance, and reconnaissance: Gaps for soldiers and technology recommendations," in *Proc MILCOM*, 2013.
6. M. Srivastava, T. Abdelzaher, and B. Szymanski, "Human-centric sensing," *Phil. Trans. R. Soc. A* **370**(1958), pp. 176–197, 2012.
7. A. Preece, D. Braines, D. Pizzocaro, and C. Parizas, "Human-machine conversations to support multi-agency missions," *ACM SIGMOBILE Mobile Computing and Communications Review* **18(1)**, pp. 75–84, 2014.
8. T. Kuhn, "A survey and classification of controlled natural languages," *Computational Linguistics* **40**, pp. 121–170, 2014.
9. J. Austin and J. Urmson, *How to Do Things With Words*, Harvard University Press, 1975.
10. Y. Labrou and T. Finin, "Semantics and conversations for an agent communication language," in *Readings in agents*, M. N. Huhns and M. P. Singh, eds., pp. 235–242, Morgan Kaufman, 1998.
11. G. Pask, "A fresh look at cognition and the individual," *International Journal of Man-Machine Studies* **4**, pp. 211–216, 2014.
12. D. Mott, "Summary of ITA Controlled English," 2010.
13. "RDF 1.1 primer." World Wide Web Consortium, June 2014.
14. "OWL 2 Web Ontology Language document overview (second edition)." World Wide Web Consortium, Dec. 2012.
15. "SWRL: A semantic web rule language combining OWL and RuleML." World Wide Web Consortium, May 2004.
16. M. Osborne, S. Moran, R. McCreadie, A. Von Lunen, M. D. Sykora, E. Cano, N. Ireson, C. MacDonald, I. Ounis, Y. He, T. Jackson, F. Ciravegna, and A. O'Brien, "Facebook, Twitter and Google Plus for breaking news: Is there a winner?," in *Proceedings of the AAAI Conference on Web and Social Media*, 2014.
17. A. Preece, W. Webberley, and D. Braines, "Tasking the tweeters: Obtaining actionable information from human sensors," in *Proc Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR VI (SPIE Vol 9464)*, SPIE, 2015.
18. A. Preece, D. Pizzocaro, D. Braines, D. Mott, G. de Mel, and T. Pham, "Integrating hard and soft information sources for D2D using controlled natural language," in *Proc 15th International Conference on Information Fusion*, 2012.
19. I. Spasić, M. Greenwood, A. Preece, N. Francis, and G. Elwyn, "FlexiTerm: a flexible term recognition method," *Biomedical Semantics* **4(27)**, 2013.
20. C. Roberts, M. Innes, A. Preece, and I. Spasić, "Soft facts and spontaneous community mobilisation: the role of rumour after major crime events," in *Data for Good: How big and open data can be used for the common good*, P. Baeck, ed., pp. 37–43, Nesta, 2015.
21. D. Mott, D. R. Shemanski, C. Giammanco, and D. Braines, "Collaborative human-machine analysis using a controlled natural language," in *Proc Next-Generation Analyst III (SPIE Vol 9499)*, SPIE, 2015.
22. D. Braines, J. Ibbotson, D. Shaw, and A. Preece, "Building a "living database" for human-machine intelligence analysis," in *Proc 18th International Conference on Information Fusion*, 2015.
23. J. Llinas, "A survey of automated methods for sensemaking support," in *Proc Next-Generation Analyst II (SPIE Vol 9122)*, SPIE, 2014.
24. R. Crouser and R. Chang, "An affordance-based framework for human computation and human-computer collaboration," *IEEE Transactions on Visualization and Computer Graphics* **18(12)**, pp. 2859–2868, 2012.