A Conceptual Model for Mixed-Initiative Sensemaking

CHI 2018 Sensemaking Workshop, April 21, 2018, Montreal, Canada
ACM xxx.
http://xxx.

Chris Baber
University of Birmingham
Birmingham, UK.
c.baber@bham.ac.uk

Kristin Cook
Pacific Northwest National Laboratory
Richland, WA
Kris.Cook@pnnl.gov

Simon Attfield
Middlesex University
London, UK
s.attfield@mdx.ac.uk

Leslie Blaha
Pacific Northwest National Laboratory
Richland, WA
leslie.blaha@pnnl.gov

Alex Endert
Georgia Institute of Technology
Atlanta, GA
endert@gatech.edu

Lyndsey Franklin
Pacific Northwest National Laboratory
Richland, WA
lyndsey.franklin@pnnl.gov

Abstract
This paper is concerned with situations in which humans and machines work together in data-intensive sensemaking. Viewing humans and machines as intelligent and collaborating agents, we suggest that such collaboration can be most effective when either agent is able to take the initiative and also to interpret the activities of the other. To this end, we propose that it is necessary for one party to be able to observe the behavior of the other in sufficient detail to be able to infer the goals, purpose, or values that are guide this behavior. We present a conceptual model which describes how such collaboration might occur. The model emphasizes the need for awareness of the others' behavior to support common ground in relation to intentions represented at different levels of abstraction. Ultimately, we are asking what does an agent need to know about another agent's behavior in order to make such inferences?

Author Keywords
Sensemaking; mixed-initiative; human-machine teaming; common ground; cognitive work analysis.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; See http://acm.org/about/class/1998 for the full list of ACM classifiers. This section is required.

Table 1: Relating Abstraction Hierarchy terms to our conceptual model

CWA terms	How we interpret this term…
Functional purpose	The overall purpose of the system; what is the primary mission that is being pursued.
Abstract function	The competing objectives that the system values, which acts as constraints on path to follow
Purpose-related function	The main reasons why object-related functions are performed, i.e., the goals, targets or outcomes that agents are seeking to achieve
Object-related function	The actions that agents can perform on the physical objects
Physical objects	The information, data, artefacts upon which agents in the system can act

Introduction

Sensemaking is a process of reasoning about information in order to construct a coherent view, belief, or understanding about a situation of interest. It typically involves structuring information into some kind of representation and in doing so giving meaning or interpretation to the information. It can also involve challenging an existing interpretation or viewpoint and often leads to the evolution of new questions, in order to elaborate or test a point of view, resulting in cycles of information seeking and interpretation.

The interaction between people and the tools they use to augment cognition during sensemaking is an important factor in how sensemaking unfolds. The kinds of tools that might be relevant range from simple pieces of paper to complex visual analytics and artificial intelligence (AI) systems. Regardless of the form that this cognitive augmentation takes however, we believe that sensemaking can usefully be regarded as an instance of distributed cognition ([1], [2], [16]).

As the tools available for data-intensive sensemaking become increasingly sophisticated and intelligent, so too they become able to take on increasingly complex aspects of the cognitive work. In this respect, they may increasingly be regarded as cognitive agents, and hence as ‘collaborators’ rather than simply as tools. Accordingly, we might think of interaction with them as a dialogue between agents. Thus, we are interested in ways of describing this collaboration and dialogue as it relates to sensemaking.

We present a conceptual model to formalize the roles and responsibilities of human and machine agents in the mixed-initiative sensemaking process. Our model (Figure 1) highlights different levels of abstraction, ranging from low-level tasks using information and data, to higher-level values and interests.

What is the Value of this Conceptual Model?

Conceptual models are foundational to establishing theory against which we can develop metrics and methodologies for evaluating the performance of systems. Importantly, a conceptual model serves to unify observations and empirical findings into a representation of a process or system [4]. Our proposed conceptual model for mixed-initiative sensemaking brings together the following key observations and ideas, to begin establishing missing theoretical frameworks supporting joint human-machine intelligent systems collaborating on sensemaking tasks:

- Articulates connection between human and machine agents in sensemaking;
- Identifies touchpoints and communications needs between human and machine agents;
- Suggests common ground as a foundation for building human-machine interactions to support sensemaking;
- Provides levels of abstraction over which sensemaking activities occur and can be articulated appropriately;
- Emphasizes that goals and purposes must be articulated between human and machine entities for mixed-initiative sensemaking to be successful.

Motivations for Mixed Initiative Sensemaking

Current theories and models of sensemaking emphasize human-centric processes of finding meaning, or distribution of processes between a team of people. Little attention has been paid to describing coordination with intelligent systems as a form of collaborative sensemaking through mixed-initiative approaches, where ‘mixed-initiative’ enables *“users and intelligent agents to collaborate efficiently”* ([13, p.165]). As an example, MAPGEN (Mixed-initiative Activity Plan GENerator) [3] was developed to support the exploration of Mars. It involved a set of desired activities (specified by scientists and engineers) that had resource requirements (in terms of power, communications etc.) and a set of 'flight rules' (in terms of an agreed set of constraints that were used to prioritise activities and manage dependencies). What made the problem challenging was the need to balance multiple constraints, goals and resource-demands. Consequently, one can see that 'sense' is not a matter of getting the 'best' answer to the scheduling problem, but rather than it involves appreciating the competing factors that can have an impact on the plan. Thus, to achieve sensemaking through mixed-initiative collaboration:

1. human and machine agents must operate with shared purpose (i.e., make the outcomes that they are pursuing clear to each party);
2. human and machine agents must have a common frame on which to operate (i.e., they must agree how constraints operate in their decision space); and
3. human and machine agents must have the ability to share both their progress and their insights (i.e., a means of sharing and negotiating purpose, constraints, and dependencies).

If 'initiative' means taking the first step, then one needs to split tasks between humans and machines. If initiative means taking the lead, then this raises bigger questions concerning authority and responsibility. We see the former as relating to the tasks performed by the human-machine system, and the latter as reflecting the values that the system embodies. The idea that initiative is mixed implies an ebb and flow of initiative, depending on the activities to perform, the constraints being applied, and the data available. One way of conceptualising this is through the system-level descriptions offered by Cognitive Work Analysis (CWA).

Levels of Abstraction

We borrow the concept of 'Abstraction' from CWA ([14], [15]). A 'system' can be described using a variety of perspectives, e.g., focusing on the manner in which the system achieves its mission, or on the strategies that individual agents working in that system might employ for their own goals. Typically in CWA, one begins with a Work Domain Analysis (WDA). WDA takes the form of an Abstraction Hierarchy which provides a high-level overview of the system. This begins with a Functional Purpose (i.e., the state of the world that the system is intended to achieve). The Functional Purpose, in this analysis, highlights the 'mission' of the system, i.e., not necessarily a defined end-state but a direction of travel. Beneath the Functional Purpose, Abstract Functions define the objectives that the system values. These could be thought as stopping-functions, in that they would define measures of performance for that system. These represent constraints (or bounds) on system activity, partly because they are measurable and partly because they could conflict with or otherwise interact with each other. Combined with these high-level descriptors, the Abstraction Hierarchy includes the physical objects that the system employs and the actions that can be performed on these objects. In this way, the Abstraction Hierarchy shows *how* the system seeks to achieve its Functional Purpose (reading the diagram top-down) and also *why* the system is organised as it is (reading the diagram bottom-up). Table 1 shows how the Abstraction Hierarchy can describe how a human-machine system can pursue a Functional Purpose. This points to three issues:

1. in what ways does this human-machine system have to have a shared (functional) purpose?
2. what values does the system exhibit? and
3. how does this system establish and maintain 'common ground' to enable these values and purposes to be negotiated?

Common Ground

We think of the problem of managing collaborative sensemaking as partly a question of establishing and maintaining common ground in relation to (minimally at least) sense and intent. By this we mean aspects of the sense that is being made of a situation and the intent that the agents have towards it. Here we draw on Clark's theory of common ground. For Clark, common ground is *"the mutual knowledge, beliefs, and assumptions shared by the speaker and addressees."* ([8, p. 247]). Common ground comprises knowledge, beliefs and assumptions that agents share and know that they share. The primary reason for establishing and maintaining common ground is to help joint action flow as seamlessly as possible, where this can include communication itself [8].

Common ground is never complete, but only needs to be sufficient to address for some prevailing endeavor or interests. Imagine a conversation with someone (say at a bus-stop or at a party) who is an enthusiast for a sport that you know little or nothing about. You could seek a detailed explanation of the sport before starting the conversation, to get a 'shared' knowledge-base from which to build understanding. But this is likely to only give some rudimentary knowledge and will take time and effort on both parts. Far more likely, you will fall back on your general knowledge of what constitutes 'sport' and use this to create assumptions to inform the ongoing conversation. It is when you demonstrate misunderstanding through what you say or do or ask for explanation that the conversation might shift to these details. Such shifts are likely to be ad hoc and only occur when necessary. Of course, much depends on the 'goal' of the conversation, e.g., to be entertained by the other person's enthusiasm, to enjoy their knowledge, to be friendly, to pass the time, or to learn a little. But each 'goal' carries with it a different way of defining and managing 'common ground'.

As our sports example suggests, the word 'mutual' and 'shared' in Clark's definition can cover many different forms. To establish and maintain common ground, people draw on three sources of information:

- perceptual evidence (the experience to which people have access);
- linguistic evidence (the words that people are hearing);
- community evidence (knowledge which they might believe is shared within a given community, perhaps as the result of training or enculturation).

From this list, one can appreciate the challenges that establishing common ground between human and machine will face.

The 'perceptual evidence' could relate to the information available to a person or the data available to a machine. However this is interpreted (by either party) will lead to expectations as to what might be the most appropriate form of analysis. One could imagine that, in this instance, common ground could be supported through allowing the human and machine to agree on what constitutes the data and the form of analysis to apply. This is fundamental to developments in visual analytics [6].

Visual analytics is also concerned with the manner in which people interact with computers. This picks up on the idea that the 'linguistic evidence' is a core aspect of developing common ground. However, such evidence is not just constructed through explicit communication but can also be inferred by observation. For example, Heath and Luff [9] observed how staff controlling a London underground line used subtle and complex co-monitoring practices to achieve close coordination. In fact, it was relatively unusual for these staff to communicate explicitly about what they are doing.

Rather, tacit cues of behavior available in virtue of their being collocated in the control room are sufficient. This process is enabled by their awareness and maintenance of a body of practice (procedures and conventions) relating to coordinated action, which *"...informs the production, recognition, and coordination of routine conduct within the line control room"* ([9], p. 102). Although Heath and Luff do not discuss this in terms of common ground, it clearly must depend upon a shared background understanding about controlling an underground line (and perhaps even that line in particular) as well as physical proximity supported by the layout of control room.

This latter point illustrates 'community evidence', and it is a moot point as to how humans and machines can develop and share such evidence. However, consider collaboration between agents investigating a possible cyber-attack. Both human and machine will develop beliefs and working hypotheses that inform, and are informed by, their analysis of available information and data. In turn, these beliefs and working hypotheses reflect expectations about how cyber-attacks operate (based on prior experience of such events) and the motives of the people behind such attacks. Understanding each agents' values and interests, we propose, is a good basis on which to predict what agents may do next and how to coordinate actions. This means that there are role-based, as well as domain-specific, modes of conducting sensemaking. Not only in terms of the information and data that would be available to the agents (human or machine) in different situations, but also the beliefs (in terms of values and interests) that are brought to bear different agents assuming different roles in the system. The point is that the same data can make 'sense' in different ways to different agents in the system.

Mixed-Initiative Conceptual Model

Figure 1 shows the Mixed-Initiative Sensemaking conceptual model that we have been developing. It consists of swimlanes for the agents that here represent the *human* and the *machine* parties in the system. On the left of the diagram, each entity has a set of *Values and Interests*. This concept, inspired by the Functional Purpose of Work Domain Analysis, is used to indicate those overarching beliefs about the mission of the system. These could be articulated in the culture a of the humans (e.g., in terms of Policing, there might be a recognised duty to reduce crime, there might a target to reduce specific types of crime, there might be a desire to maintain good relations with the local community, there might be a need to gain increased funding or recruit more officers, etc.).

For the machine swimlane, the issue of what values and priorities it has been designed to pursue could be embodied in the design of the machine or the procedures and work practices of the system. In other words, it is not something that it is always articulated or shared. This means that the *Common Ground* (in the space between these two swimlanes) will involve either entity using the values and interests that it is favouring as a means of ordering and prioritising possible *Purposes* (which can be thought of in terms of the Abstract Functions in the WDA). In this case, the Purpose has associated measures of performance, which implies a stopping function, but also implies a trade-off space in which different purposes could be emphasised at the expense of others. This implies a process by which Purposes can be prioritised. We suggest that this could easily create confusion and ambiguity, unless it is possible for each agent to appreciate the other's position. This is analogous to the 'flight rules' in the MAPGEN example earlier. However, it might also reflect the shifting 'goals' hinted at in the 'sports' conversation -- as the activity develops, so agents might diverge in what they are seeking to achieve and which purpose they are prioritizing. In this case, common ground would relate to acknowledgement of the similarities and differences in

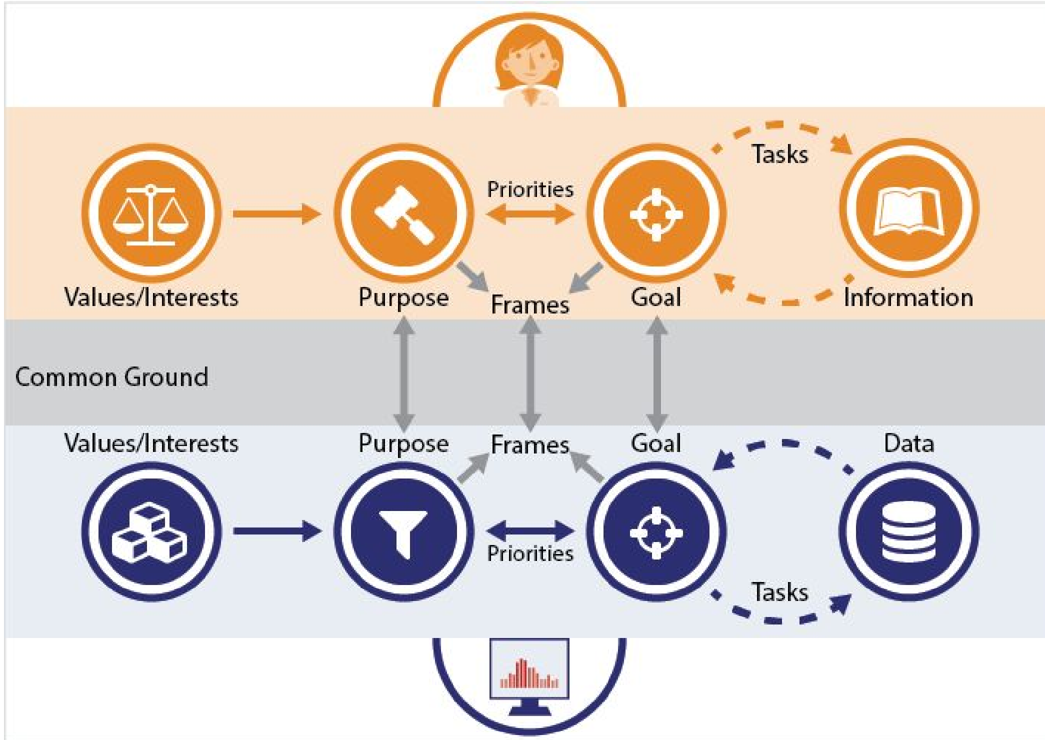

Figure 1: Conceptual Model of mixed-initiative sensemaking depicting both human processes (on the top) and machine processes (on the bottom).

these purposes. On some occasions, one might want both parties to have the same purpose but on other occasions, it might be more useful to operate with different purposes (either to make best use of the capabilities of the different entities, or to create opportunity for debate and analysis).

We see the challenge of prioritizing purposes as involving the definition and elaboration of frames (as in the Data-Frame model [11]). The suggestion is that a 'top-down' definition of a frame could be derived from a Purpose. For instance, one Purpose might be to have 'minimal intruder activity'. This sets up a frame in which parameters related to intruder detection become important, which, in turn, sets goals (analogous to Purpose-related functions in te WDA) for the agent in terms of the tasks (analogous to Object-related functions in WDA) that need to be performed on or with the information / data. From reading Figure 1 left-to-right, one can see how the values and interests of the agent relate to the purposes that they seek to achieve, and how these are prioritized and lead to goals and tasks.

Figure 1 can also be read from right-to-left. In this case, the 'information / data' that an agent encounters can suggest, or afford, a task. This task then leads to a reframing of the situation (in terms of the Data-Frame model [11]), and this leads to the definition of a purpose, which can then be assessed in terms of values and interests.

Changes in each stage along the horizontal axis for one or both agents, imply a possible need to resolve common ground across the vertical dimension. If common ground is not achieved, the agents may need to revert to higher levels of abstraction until common ground is found, before delving more deeply into the later stages of sensemaking. This raises the question of how to manage common ground across these stages.

Common Ground between Humans And Machines

As noted previously, we recast Clark's [5] 'evidence' to focus on human-machine interactions. 'Perceptual' evidence would encompass the information and data that either agents can access. In this case, common ground is a matter of interactive visualization, so that the human is able to perceive, and act upon, the data that the machine is analysing [8]. In this way, the 'frame' that is applied to the data can be agreed between human and machine. However, even this objective becomes problematic when the data become 'big' (or varied or ambiguous). Under these circumstances, common ground will need to be established in terms of the Purpose of the analysis. There are interactive techniques that allow people to provide feedback to analytic models and AI agents. For example, semantic interaction is an approach that couples model steering with user interactions native to the spatial organization of text documents [7]. This allows common ground to operate on 'linguistic evidence' or the 'purpose stage'. What is less apparent is how one might extends techniques to operate at the values and interests stage. We believe that this represents a new challenge for interaction design and can have implications for human-machine teaming and for explainability of AI output.

Conclusions

In this paper, we have proposed a conceptual model of mixed-initiative sensemaking which emphasizes the importance of common ground across different stages of sensemaking. The broad argument follows the suggestion that it is not sufficient to follow a 'human in the loop' approach to designing future sensemaking support, but to recognise that the 'human is the loop' [8]; or rather, to propose that the 'loop' involves the tightly coupled collaboration between human and machine in the establishment and management of common ground.

Not only does our conceptual framework hint at ways in which a research agenda could be fashioned to support this search for common ground, but we also propose that there are practical challenges that can lead to new approaches to design and development. For example, understanding the relationship between tasks and purpose could allow inference of the higher level values and interests that the agent is pursuing -- and these could then be negotiated with other agents in the system to ensure appropriate prioritization of purpose for a given mission. Ultimately, we are asking what does an agent need to know about another agent's behavior in order to make such inferences?

Acknowledgements

This effort was sponsored by the Analysis in Motion Initiative at the Pacific Northwest National Laboratory. 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. Government.

References

1. Simon Attfield, Bob Fields, B., Ashley Wheat, Rob Hutton, Jim Nixon, Andrew Leggatt & Hannah Blackford. Distributed sensemaking: A case study of military analysis. 12th International Naturalistic Decision Making Conference, McLean, Virginia, USA. 201
2. Chris Baber & Richard McMaster. Grasping the moment: sensemaking in response to routine incident and major emergencies. Avebury: Ashgate. 2017.
3. John L. Bresina, Ari K. Jonsson, Paul H. Morris & Kanna Rajan. Mixed-Initiative Planning in MAPGEN: Capabilities and Shortcomings. Workshop on Mixed-Initiative Planning and Scheduling. ICAPS05. 2005.
4. Min Chen, Georges Grinstein, Chris R. Johnson, Jessie Kennedy, & Melanie Tory. 2017. Pathways for theoretical advances in visualization. IEEE Computer Graphics and Applications 37, 4: 103-112.
5. Herbert H. Clarke. Using Language, Cambridge: Cambridge University Press. 1996.
6. James J. Thomas and Kristin A. Cook, eds. Illuminating the path: the research and development agenda for visual analytics. Pacific Northwest National Laboratory, Richland, WA. 2005.
7. Alex Endert, Patrick Fiaux, and Chris North. "Semantic interaction for visual text analytics." *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2012.
8. Alex Endert, M. Shahriar Hossain, Naren, Ramakrishnan, Chris North, Patrick Fiuaux & Christopher Andrews. The human is the loop: new directions for visual analytics. Journal of Intelligent Information Systems. 43. 411-435. 2014.
9. Christian Heath & Paul Luff. Team work: collaboration and control in London Underground line control rooms, In C. Heath and P. Luff (Eds.) Technology in Action, Cambridge: Cambridge University Press. 88-124. 2000.
10. Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM.
11. Gary Klein, Brian Moon and Robert R. Hoffman. Making Sense of Sensemaking 2: A Macrocognitive Model. IEEE Intelligent Systems, 21, 88-92. 2006.
12. Andrew Monk. Common ground in electronically mediated communication: Clark’s theory of language use. *HCI models, theories, and frameworks: Toward a multidisciplinary science*, pp.265-289. 2003.
13. Peter Pirolli and Stuart Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. Proceedings of International Conference on Intelligence Analysis. 2005.
14. Jens Rasmussen, Annelise M. Petersen and L.P. Goodstein. 1994. Cognitive Systems Engineering. New York, NY: Wiley.
15. Kim J. Vicente. 1999. Cognitive Work Analysis. Mahwah, NJ: Lawrence Erlbaum Associates.
16. Ashley Wheat, Simon Attfield & Bob Fields. Developing a model of distributed sensemaking: a case study of military analysis. *Informatics*, 3, 1. 2016.