How analysts think: A preliminary study of human needs and demands for AI-based conversational agents

Sam Hepenstal[1], B.L. William Wong[2], Leishi Zhang[2] and Neesha Kodogoda[2]

[1] Defence Science Technology Laboratory, Porton Down, Salisbury SP4 0JQ, England
[2] Middlesex University London, The Burroughs, Hendon, London NW4 4BT, England

For conversational agents to provide benefit to intelligence analysis they need to be able to recognise and respond to the analysts intentions. Furthermore, they must provide transparency to their algorithms and be able to adapt to new situations and lines of inquiry. We present a preliminary analysis as a first step towards developing conversational agents for intelligence analysis: that of understanding and modeling analyst intentions so they can be recognised by conversational agents. We describe in-depth interviews conducted with experienced intelligence analysts and implications for designing conversational agent intentions.

Key Words

Algorithmic transparency; conversational agents; cognitive task analysis; intention capture; intelligence analysis

INTRODUCTION

In this paper we present findings from a preliminary Cognitive Task Analysis (CTA) investigation into identifying the 'intentions' of intelligence analysts when retrieving information during live investigations. These findings will be used to model human intentions that can be developed into concepts for an AI-based technology called a conversational agent (CA).

Artificial intelligence (AI) technologies which fulfil the role of personal assistants are becoming a norm in people's everyday lives. The popularity of AI CAs is increasing (Kinsella 2018, Kinsella 2019). This type of CA is defined as a 'spoken dialogue system' (McTear 2002), which uses spoken language to interact with users to accomplish a task. Such systems are particularly attractive due to the ease with which mundane yet otherwise time-consuming tasks can be performed. By applying rules to define the task required by a user the CA identifies requirements and responds accordingly.

In general, the tasks performed by commercial personal assistants are low risk and there are limited consequences if the CA gets it wrong. For example, Google Assistant advertises that it can be used to manage user defined tasks, plan a user's day (finding directions, booking meetings, and other similar tasks), query for media or knowledge, or manage connected devices in the home. These tasks can be easily validated against user expectations. There is no need to understand how information has been found or the method.

We believe CAs can also benefit areas which require high-risk and high-consequence decision making, for example, by speeding up complex queries to provide information retrieval and analysis in a police investigation or through shared human-machine reasoning across large datasets. If, for example, an analyst could ask a CA a question like "what vehicles are owned by associates of [known offender]?" and a CA could interpret this, perform the necessary searches, and present results to the analyst. It would deliver results significantly faster than an analyst manually building, performing and collating queries. We have identified a number of problems, however, which are critical to address if AI CA's are to be used for intelligence analysis. Firstly, analysts must trust the system. To foster trust there needs to be a mechanism to achieve common understanding between human and machine of the goals, strengths and constraints of each party. We believe this can be provided by designing for algorithmic transparency, beyond current approaches to 'explainability'. Explanations which focus upon the mathematical models used in machine learning algorithms are only one part of the requirement. Algorithmic transparency must also provide visibility of the goals and constraints of the AI system (Hepenstal et. al. 2019). In this way we can develop applications which allow for shared learning and harness the capabilities of both human and machine. Secondly, there is a problem with the brittleness of a CA due to the need to design actions in advance. In terms of the Law of Requisite Variety brittleness occurs when the technology fails to cope with the variety of demands that it has to cope with when in use. For example, if we develop a CA with Artificial Intelligence Markup Language (AIML), a commonly used approach to develop chat interactions which supports many chatbot platforms and services (Radziwill and Benton 2017). A pattern is described for a task category (intention) together with a template response if the users' text matches with the pattern. There is a need to define or learn the intentions (and subsequent actions) which the CA can fulfil, and to ensure that they are consistent and distinct so as not to confuse either the pattern matching algorithm or the user. In intelligence analysis investigations there is a need for flexibility, where the direction of investigation and the information required may not be known beforehand. Analysts require a CA which can learn from them and evolve to identify new intentions and new tasks, whilst still providing algorithmic transparency.

Our aim is to report on a study to elicit the 'intentions' of intelligence analysts when retrieving information during investigations (i.e. the questions asked and attributes required to answer), to assist the design of CA capabilities. We use CTA interviews and capture insight (i.e. clear and deep understanding) of the questioning, elaborating, reframing, and connecting (i.e. sense making) strategies that occur during early to late stages of investigation. We believe this understanding is necessary for developing AI CA's that can recognise and interpret analyst questions correctly, and also for meeting transparency needs. A preliminary analysis of interview data is discussed with some key findings, specifically, on the importance of hypothesis scope for enabling recognition and the consolidated information requirements for CAs to answer analyst questions. We identify that while CAs must be cognisant of the scope of an investigation, they present a significant opportunity to help mitigate key problems within current approaches to intelligence analysis such as cognitive bias and availability bias.

We propose that CA intentions should be underpinned by CTA data and that this approach can enable us to design CAs which overcome problems encountered by earlier AI systems, such as brittleness and transparency.

RELATED WORK

An initial step to understand what a CA should be able to do is to capture what analyst intentions look like and ways to structure them so they can be recognised by a CA. Previous research has been conducted to investigate inference making within intelligence analysis (Wong and Kodagoda 2016), however we are specifically interested in finding the investigation requirements to retrieve information. There are a variety of cognitive models which could help model analyst intentions. In this paper we have considered Toulmin's model for argumentation (Toulmin 1958), Klein's data-frame model for sensemaking (Klein et. al 2006), with a focus on the Recognition-Primed Decision (RPD) model (Klein 1989) as an analytic framework.

The RPD can be used to explain how experienced people make rapid decisions, including how they recognise a situation. Klein describes "four important aspects of situation assessment (a) understanding the types of goals that can be reasonably accomplished in the situation, (b) increasing the salience of cues that are important within the context of the situation, (c) forming expectations which can serve as a check on the accuracy of the situation assessment (i.e., if the expectancies are violated, it suggests that the situation has been misunderstood), and (d) identifying the typical actions to take." (Klein 1993) These allow people to compare patterns from their experience of past situations with emerging situations, thus enabling quick understanding, predictions and decisions. We propose that these four aspects are also crucial to the recognition of a given situation by a CA. The experience of a CA is captured by the training data used to train the intention classification model. This is how a CA can recognise user queries. When a user asks a question the CA first needs to find the desired intention, this uses cues in the users language and extracted entities to find a match. An intention will have associated goals (it can only recognise and respond if it can fulfil the predicted task) and corresponding actions, for example to apply a function to explore the data and find connections between two entities of interest. There may also be expectations for lines of inquiry based upon past similar questions and data patterns which influence the CA's choice of action. Recognition is not only important when building CAs, but also crucial to provide a user with an understanding of CA cognition. If we are able to make the situation assessment performed by the CA visible to a user for any intention which matches the user's input, then they will be able to see what the CA is reasoning and doing.

METHOD

We conducted four in-depth interviews with experienced intelligence analysts. Three analysts have worked in policing, across various police forces, and the other analyst has a background in defence intelligence. Each analyst has over 3 years of experience working on a range of operational investigations, with a focus upon major crime. We have chosen CTA because our aim is to understand intentions which underpin how an analyst thinks and reasons. Each interview lasted an hour and applied the Critical Decision Method (CDM) (Klein 1989; Wong 2003) to elicit analyst expertise, cues, goals and decision making on a memorable investigation they were involved with from start to end. The CDM interview technique was used to ensure important information was captured. Of particular interest were the nature and requirements of analyst questions at critical stages in investigations, specifically, their cues, goals, expectations and actions. These stages are typically time-pressured and are therefore prime situations in which CAs could assist analysts. Interviews addressed the analyst's experience, such as the conditions which allowed them to use their prior knowledge and the recognition of situations which a novice analyst may have missed. A timeline of key events was sketched out by the analyst and explored in detail.

Preliminary data analysis has looked to identify themes across critical decision points. Figure 1 shows a timeline of critical stages of the kidnapping scenario. By delving into these stages during interviews we drew attention to the analyst's argumentation, sensemaking and recognition. We use a similar approach to Wong (2003) to move from an incident summary, to a decision chart, and finally to develop a decision analysis table. The evidence trail has been preserved by linking each passage in the table to detailed transcript statements, with reference to the time and the interview audio file in which statements were said.

EARLY ANALYSIS AND RESULTS

In the CTA we studied four cases: a kidnapping case, an attempted murder case, target analysis, and firearm dealing. Three key themes have emerged from the interviews with analysts. Firstly, in a live investigation analysts often lack time to conduct in depth analysis to refute (verify and validate) their hypotheses. As one analyst put it, "*you don't always get comfort to do that, you respond and validation comes afterwards.*" [A1, 40:30] Incident timelines show how each piece of insight leads to intense periods of manual information gathering. For example, in Figure 1 we see that a single insight (that the victim has been assaulted in the past) leads to a variety of time consuming information retrieval tasks.

Figure 1: Kidnapping Scenario Timeline

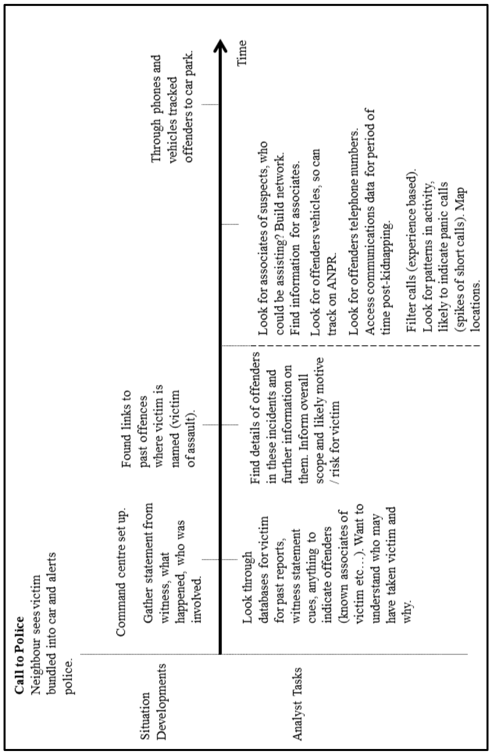

Secondly, information is not always available, either because it simply does not exist, cannot be accessed, or there is so much of it an analyst cannot filter and explore it in time. To correct for these limitations experience is needed, for example to crudely filter large volumes of data based upon expected patterns, or to apply abductive reasoning to predict new lines of inquiry from small amounts of information. The final theme is, therefore, that commonly held assumptions are significant in guiding the direction and boundaries of investigation paths by influencing the hypothesis which explains the overall investigation scenario. We have termed this hypothesis the investigation 'scope'. Investigation scope is crucial to direct intelligence analysis and enable recognition, by creating a basis from which expectancies can be drawn. In all interviews it has been a key priority for analysts to identify a scope so that investigations can advance.

In the kidnapping incident the analyst was able to identify that the victim "*had been assaulted over a period of time by a fluid group of local boys.*" [A1, 16:00] This led them to narrow the scope of the incident to a local gang and nothing larger. Hence, when later the analyst was tasked to investigate communications data they discounted calls outside the local region, "*if they are phoning someone in another part of the country they are not likely to be part of the kidnapping. We were looking for local calls on a frequent basis and a cluster around the time he was kidnapped.*" [A1, 32:00] In the case where the analyst was looking for a firearms dealer, scope was also narrowed based upon experience for the type of call which the analyst expected from a firearm dealer i.e. that the duration would not be really short (less than a few seconds), it would not be a system call, and "*we can rule out text messages, based upon experience that criminals (when purchasing firearms) normally call about this kind of thing.*" [A4, 16:00] The analyst also "*knew this was a local criminal group, through their lifestyle and surveillance on people involved. It is a picture you build over time.*" [A4, 15:30] They were therefore able to narrow the scope for the rest of the investigation and in subsequent questions looked to build information against expected patterns within the hypothesis scope. Without a reasonable definition for scope, analysts can struggle to recognise patterns and identify lines of inquiry in time pressured situations. The possibilities are too broad and analysts need a basis on which to make sense of information and identify expectancies. The past experiences which have informed the analysts mental patterns and enable abductive reasoning are valuable and can deliver results quickly by stripping out extraneous data. The scenario described by an analyst of an attempted murder is a good example where scope could not be defined and subsequently the investigation could not progress. The husband of the victim had been a suspect but was cleared through verification of his statement, via call data and CCTV. "*No other evidence was available, so no lines of inquiry. Expertise, such as the burglary expert, felt it looked like staged burglary. Had expected pattern to the way draws had been pulled out (if burglar), that it was not a burglary pattern. I felt this was too tenuous. It could be a novice burglar.*" [A2, 36:00] The investigation continued, but could not progress any further and the case remains unsolved.

There is a danger that cognitive bias will be introduced to investigations when abductive reasoning is used to commit to a hypothesis scope which subsequently directs investigation and intelligence development. If lines of inquiry seek expected patterns only within the hypothesis scope then analysts are effectively looking to confirm what they already believe, rather than considering alternatives. Decisions which are informed by experience can be misled given the influence of narrative (Pennington and Hastie 1992), particularly in intelligence analysis where there are typically strong narratives behind unusual scenarios. If the true pattern of activity is outside the scope for an investigation it may be missed. Scope is therefore crucial to understand and refine throughout an investigation, but this must be done with care. The creation of scope and expected patterns for emerging information is perhaps a perfect example where human and machine should collaborate. We propose that the design of CA intentions and interaction must be cognisant of the hypothesis scope of the investigation and, if possible, able to review and reveal alternative paths and patterns. Where there is a lack of information the CA must be able to utilise analyst expertise to direct investigations, whilst providing validation and verification. We propose that information retrieval and analysis should take a shared approach where the analyst can learn how to interact with the CA and vice versa.

STRUCTURING INTENTION ATTRIBUTES

We wish to develop CAs which can aid intelligence analysts by understanding and responding to their intentions. To achieve this we look to provide contextual structure to the points of inquiry when an analyst interacts with a CA. Figure 2 presents three structures which could capture the pattern of analyst intentions. We describe a specific stage in the firearms scenario as an example, where the analyst wants to find contact between a person of interest (POI) and a firearm dealer.

Figure 2: Models for structuring extracted intention attributes

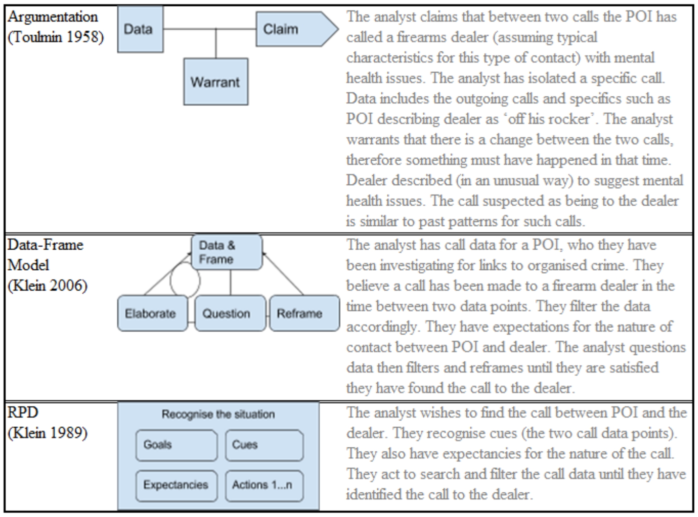

Figure 3: Snippet of Decision Analysis Table (Kidnapping)

HYPOTHESIS SCOPE	**WHO (VICTIM), WHEN, HOW**	**WHO (VICTIM, OFFENDER), WHEN, HOW, WHERE, WHY**
CUES	Man gone missing. Thought he had been kidnapped due to witness report. Known to be vulnerable.	Identified in police records man had been victim of assaults by fluid group of youths. Not linked to NCA or serious organised crime.
GOALS	Understand what could have happened and more about the victim	Understand how dangerous youths are, find out what vehicles they use and telephone numbers and where they live/operate.
EXPECTANCIES	Unknown - scope too broad	That those involved are local known bullies and not OCG. Expectation that this was an incident which had gone too far and offenders had panicked.
ACTIONS	Searched known associates of victim, looked for previous convictions, spoke to neighbours and witnesses, looked at telephone information. looked for victims name.	Search databases for offenders looking for vehicles, telephone numbers and associates
WHY?	To reduce scope of investigation and assess level of risk	To assess risk to victim (danger posed by offenders) and to trace possible locations through vehicles, telephones, addresses
WHAT FOR?	To direct next steps of investigation and better use experience to recognise patterns	To locate the victim and provide support

The attributes defined by any of these models could be used to derive intentions for CAs. Recognition is integral to identifying the appropriate query to retrieve information and we have therefore focused upon the RPD. Figure 3 provides some example snippets from our decision analysis table which draws upon aspects of recognition from the RPD model. We have included the addition of scope to the table, defined as the overarching elements of a situation for which the analyst has an accepted hypothesis about what is occurring. For example, in the kidnapping incident, the analyst described some local offenders and their relation to the victim, that "*they would break in and attack him. He was continuously assaulted until he gave them money. I think everything had gone too far on this occasion so they bundled him into the car.*" [A1, 11:30] This captures the overall hypothesis scope of the investigation. Expectancies can then be described for the specific patterns which occur within the scope, for example that the victim will have been taken somewhere within the local area. Cues are the snippets of specific pieces of information which have directed critical stages of the investigation, such as the identification of a suspect's vehicle, or the report of the firearm sale where "*we had heard (POI) was looking to get a firearm from someone 'off his rocker'*" [A4, 9:00]. Actions are the information retrieval tasks, including the methods, which the analyst carried out to achieve their intended goals. For example, one analyst described an action to find "*key individuals to the network. We were looking for pinch points, for key facilitators.*" [A3, 9:00] Using the Decision Analysis Table (Figure 3) we have consolidated themes (Figure 4). These themes provide the pattern of information which a CA needs to extract from a question. This pattern can be recognised by a CA and trigger information retrieval.

Figure 4: Consolidated Decision Analysis Table

HYPOTHESIS SCOPE	**ASSESS 5WH**
CUES	Inputs for 5WH (persons name, vehicle reg, time span etc...) and relationships where necessary.
GOALS	To retrieve summary information, or specific details
EXPECTANCIES	Expected event pattern for scope informed by past events with similar scope (experience).
ACTIONS	For information retrieval these include: adjacent information (i.e. who is registered to phone number), connected information (i.e. what associates linked to a telephone number called by an offender live in a particular location), common connections (i.e. in what locations have both phone numbers been together) amongst others.
WHY?	To build on, refute, or confirm scope and associated pattern.
WHAT FOR?	To advance the investigation

Scope is the Who, What, Where, When, Why and How (5WH) elements which are present within an analyst's hypothesis to describe the overarching situation. The more elements for which there is a hypothesis, the narrower the scope of the investigation. Expectancies are the patterns which can occur within the hypothesis scope. Cues are specific inputs which will be used to build queries to retrieve information, such as the name of a victim, or their semantics. Goals are the type of information which is required to provide an answer to the analyst, such as specific details of a victim and their history, or a summarised picture of call patterns. Actions are the types of search required, for example to look at adjacent links to a specific piece of information (i.e. the vehicles owned by a person) or broader connections (the addresses associated with telephone numbers contacted by a phone number of interest). In the context of graph data we can think of possible actions as graph tasks (Lee et. al. 2006). The RPD Consolidated Decision Analysis Table presents a basis for query attributes which could answer the information retrieval questions described in interviews. For example, to ask the CA to find vehicles which are registered to a POI the CA would only require the POI identifier (name), the item of interest (vehicles) and the relationship (registered to) as cues, to understand that adjacent searches is the action, and that specific information is the goal, to return the information desired. More complicated questions could utilise other models of cognition, such as argumentation or sensemaking models. It is important that AI-enabled technologies can apply an appropriate model and extract necessary attributes. One analyst described a scenario where a software tool could clean phone data and provide the top ten call results. This was not particularly useful, as the "*top ten probably isn't very interesting if it is their mum and sister*". [A4, 27:30] The expectation is that, given the scope, the POI is considered unlikely to have their firearm dealer as a top contact.

We propose that, by capturing analyst behavior using CTA methods and modelling intentions within the context of an appropriate model for analyst cognition, we can model a 'brain' for a CA which can address issues of transparency and brittleness.

FUTURE WORK

Our initial interviews with analysts have helped to confirm our suspicions that CAs can provide a useful aid to intelligence analysts and help mitigate the three problems of lack of time, access to information, and common assumptions. For example, in situations which require support to live operations and frequent tasks to retrieve information, such as during a kidnapping, saving even just a small amount of time by reducing the need to form and write query syntax could have a huge impact. One analyst described that "*if in a threat to life situation that process (to find and map phone numbers linked to a location) might take me twenty minutes. If a computer can answer me that in one minute, literally that nineteen minutes could save someone's life.*" [A4, 47:00] We feel, therefore, that there is value in pursuing research to address the problems identified for CA interfaces for use in intelligence analysis. CAs must, however, overcome the problems of transparency and brittleness. Analysts articulated the need for transparency, when described succinctly that "*an analyst always has to justify (in court) what they have done, and so should the system.*" [A4, 35:00]

Through our preliminary analysis we have developed a contextual structure of query attributes for analyst intentions when retrieving information (Figure 4). A recommendation for further work is to conduct more interviews and refine attributes within a specific model or set of models. We also propose to identify distinct intentions from our CTA interview data and consider how this model can learn and evolve. We should also consider how a CA provides responses to users, particularly given the effect that narrative descriptions can have on decision making (Pennington and Hastie 1992).

ACKNOWLEDGEMENTS

This research was assisted by experienced intelligence analysts from the Defence Science Technology Laboratory (Dstl).

BIBLIOGRAPHY

"Google Assistant." <https://assistant.google.com/>.

Hepenstal, Sam, Kodagoda, Neesha, Zhang, Leishi, Paudyal, Pragya and Wong, B. L. William. "Algorithmic Transparency of Conversational Agents." In Joint Proceedings of the ACM IUI 2019 Workshops, Los Angeles, USA, March 20, 2019.

Kinsella, Bret. "Amazon Echo Device Sales Break New Records, Alexa Tops Free App Downloads for iOS and Android, and Alexa Down in Europe on Christmas Morning." 26 Dec 2018.

---. "NPR Study Says 118 Million Smart Speakers Owned by U.S. Adults." 7 Jan 2019.

Klein, G. A., R. Calderwood, and D. MacGregor. "Critical Decision Method for Eliciting Knowledge." *IEEE Transactions on Systems, Man, and Cybernetics* 19.3 (1989): 462-72.

Klein, G., B. Moon, and R. R. Hoffman. "Making Sense of Sensemaking 2: A Macrocognitive Model." *IEEE Intelligent Systems*21.5 (2006): 88-92.

Klein, G. "A Recognition Primed Decision (RPD) Model of Rapid Decision Making." (1993).

"Language Understanding (LUIS)." <https://www.luis.ai/home>.

Lee, Bongshin, et al. *Task Taxonomy for Graph Visualization*. ACM, 2006.

Mctear, Michael. "Spoken Dialogue Technology: Enabling the Conversational User Interface." *ACM Computing Surveys (CSUR)* 34.1 (2002): 90-169.

Pennington, Nancy, Reid Hastie, and Abraham (editor) Tesser. "Explaining the Evidence: Tests of the Story Model for Juror Decision Making." *Journal of personality and social psychology* 62.2 (1992): 189-206.

Radziwill, Nicole M., and Morgan C. Benton. "Evaluating Quality of Chatbots and Intelligent Conversational Agents." *arXiv preprint arXiv:1704.04579* (2017).

Toulmin, Stephen E. *The Uses of Argument*. Cambridge University Press, 1958.

Wong, B. L. W. *Critical Decision Method Data Analysis*. Lawrence Erlbaum Associates, 2003.

Wong, B. L. W., and Ann Blandford. *Field Research in HCI: A Case Study*. New Zealand Chapter of ACM SIGCH, 2003.

Wong, B. L. W., and Neesha Kodagoda. *How Analysts Think: Anchoring, Laddering and Associations*. 60 Vol. Los Angeles, CA: SAGE Publications, 2016.