What are you thinking? Explaining conversational agent responses for criminal investigations

Sam Hepenstal
Defence Science Technology Laboratory, UK

Leishi Zhang
Middlesex University London, UK

Neesha Kodagoda
Middlesex University London, UK

B. L. William Wong
Middlesex University London, UK

ABSTRACT

The adoption of complex artificial intelligence (AI) systems in environments that involve high risk and high consequence decision making is severely hampered by critical design issues. These issues include system *transparency* which covers (i) the explainability of results and (ii) the ability of a user to inspect and verify system goals and constraints. We present a novel approach to designing a transparent conversational agent (CA) AI system for information retrieval to support criminal investigations. Our method draws from Cognitive Task Analysis (CTA) interviews to inform the system architecture, and Emergent Themes Analysis (ETA) of questions about an investigation scenario, to understand the explanation needs of different system components. Furthermore, we implement our design approach to develop a preliminary prototype CA, named Pan, which demonstrates *transparency* provision. We propose to use Pan for exploring system requirements further in the future. Our approach enables complex AI systems, such as Pan, to be used in sensitive environments introducing capabilities which otherwise would not be available.

CCS CONCEPTS

• **Computer systems organization** → **Human-centered computing**.

KEYWORDS

explainability, criminal intelligence analysis, conversational agents, transparency

ACM Reference Format:
Sam Hepenstal, Leishi Zhang, Neesha Kodagoda, and B. L. William Wong. 2020. What are you thinking? Explaining conversational agent responses for criminal investigations. In Proceedings of the IUI workshop on Explainable Smart Systems and Algorithmic Transparency in Emerging Technologies (ExSS-ATEC'20). Cagliari, Italy, 7 pages.

1 INTRODUCTION

Artificial intelligence (AI) based conversational agent (CA) technologies are complex systems which are increasing in popularity [5, 6], because they provide more intuitive, natural, and faster access to information. They could benefit criminal investigations, where repeated information retrieval tasks are performed by analysts and the volume of data that requires filtering and processing is significant. In June 2019 Cressida Dick, the Commissioner of the Metropolitan Police, explained that "sifting through vast amounts of phone and computer data is partly to blame (for low solved crime rates) as it slows down investigations"[16]. A more natural interaction, which removes the requirement for analysts to translate their questions into restrictive syntax or structures, could speed up this process significantly. If an analyst were able to communicate with their data in the same way as they do with their colleagues, through natural language, then they could achieve significant time savings and speed up investigations.

However, for complex applications to be used in high risk and high consequence domains, transparency is crucial. If an analyst misinterprets system processes and information caveats when retrieving information in a live investigation then the impacts can be serious, for example leading to errors such as directing resources to the wrong location, or failing to find a vulnerable victim. Misinterpretation is a particular risk where there are subjectivities, such as when using a CA to interpret human intentions. We define transparency as the ease with which a user can (i) explain results provided by a system, in addition to (ii) being able to inspect and verify the goals and constraints of the system within context [3]. Without transparency, including appropriate levels of audit, complex systems cannot be used by intelligence analysts to support their investigations.

The domain of intelligence analysis is broad and diverse, therefore we have focused upon a narrow spectrum of criminal intelligence analysis and information retrieval tasks. To develop the prototype we first gathered and analysed data from CTA interviews with operational police analysts to identify the way they recognise, construct, and develop their questioning strategies in an investigation. We captured important attributes within the interview data linked to the Recognition-Primed Decision (RPD) model [9] and applied Formal Concept Analysis (FCA), a mathematical method to transform question objects and associated functional attributes into lattice structures, to identify intention concepts (*contribution 1*). We can therefore provide an explanation structure for each intention, and the underlying system processes, which mirrors the way in which humans recognise situations. We propose that this approach enhances the ability to inspect system behaviour and deliver transparency.

We also present findings from scenario based interviews with a different set of operational analysts. In these we looked to identify what information is required in explanations of the various components that form a CA. The interview data is distilled to distinct statements made by the analysts and further refined using Emergent Themes Analysis (ETA), to form an explanation framework

ExSS-ATEC'20, March 2020, Cagliari, Italy

covering CA system components (*contribution 2*). We describe a novel CA prototype, named Pan, (*contribution 3*) designed to address transparency issues, using our findings from the two sets of interviews.

The work discussed in this paper provides a preliminary investigation of transparency issues for information retrieval with complex systems, in the specific domain of criminal investigations. In future work we plan to probe further and to evaluate the prototype through experimentation with intelligence analysts.

2 RELATED WORK

Analysts play an important role in criminal investigations, as the results of their analysis underpins decision making by police commanders. For example, intelligence analysis directs the prioritisation of lines of inquiry in an investigation and assessments of key suspects. The process of intelligence analysis involves repetitive and intellectually non-trivial information retrieval tasks where "each piece of insight leads to intense periods of manual information gathering"[4]. For example, if a new lead is provided about a suspicious vehicle, analysts would ask questions such as 'who owns the vehicle?' and 'is the vehicle linked to any previous incidents?' If an intelligent system can improve this process the impact could be significant.

Manual formulation of query syntax or interactions with traditional analysis tools can be cumbersome and time consuming. A more natural interaction, which removes the requirement for analysts to translate their questions into restrictive syntax or structures, could speed up this process significantly. If an analyst were able to communicate with their data in the same way as they do with their colleagues, through natural language, then they could achieve significant time savings and speed up investigations.

We define typical CAs as being able to understand users by matching their input pattern to a particular task category (intention), for example through 'Artificial Intelligence Markup Language' (AIML) [15], where the intention triggers a set of functional processes. For banal tasks, such as playing a music playlist, the risks of an incorrect or misleading response are low and the resulting consequences limited. As a result, traditional CAs have not been built with algorithmic transparency in mind. If you ask Google Assistant, for example, why it has provided a particular response it will not be able to tell you and instead responds with humour, such as 'Let's let mysteries remain mysteries.' This is not appropriate for use in criminal investigations where decisions can have serious impact, for example to direct resources towards the wrong suspect.

2.1 Criminal investigations need explaining

Some research to date has touched on the need for a CA to be able to explain its responses. Preece et al. describe the ability to ask a CA 'why' they have provided a particular response, so an analyst can obtain the agent's rationale. An explanation could be "a summary of some reasoning or provenance for facts"[13]. This understanding of explanation is consistent with research into explainable machine-learning, where the focus is placed upon the specifics of the data retrieved, or the internals of a model. Gilpin et al. [2], defines eXplainable AI (XAI) as a combination of interpretability and completeness, where interpretability is linked to explaining the internals of a system and completeness is to describe it as accurately as possible.

Intelligence analysis is a field where analysts operate in complex, subjective, uncertain and ambiguous environments, and a simple explanation of the data or a model which defines a response is not enough to satisfy their needs for understanding. For example, if the method applied by the system presents significant constraints of which the analyst is not aware. Previous research has looked at this issue and developed a design framework for algorithmic transparency [3]. This describes the necessity to go beyond XAI when designing intelligent systems, to include visibility of the system goals and constraints within context of the situation. Context relates to the usage and user, including a user's mental model for the ways in which the CA system works. Users who have a different mental model to the realities of the system can encounter difficulties and are prone to error [12].

2.2 Structuring Human-Machine Recognition

In a policing scenario, when an analyst is presented with a situation they immediately look to make sense of it. They apply experience to recognise aspects of the situation and construct a plausible narrative explanation with supporting evidence. Klein [9] presents the Recognition-Primed Decision (RPD) model to characterise how humans recognise and respond to situations, including their cues, expectancies, actions and goals. The RPD model was first developed to understand how experienced people can make rapid decisions using a case study on fire ground commanders, another high risk and high consequence domain.

We desire a CA that can recognise situations in a similar fashion and respond to analyst questions appropriately. We also need analysts to recognise the behaviour of a CA in each situation when it attempts to understand and respond to the analyst. We propose that the RPD model provides a useful foundation to designing CA intentions so a CA can recognise analyst inputs, in addition to an explanation structure so that its behaviour can also be recognised and understood by the analyst.

3 MODELLING CA INTENTIONS

3.1 Participants and Method

We conducted Cognitive Task Analysis (CTA) interviews, applying the Critical Decision Method [9], with four intelligence analysts to delve into a particularly memorable incident for each. The analysts have a minimum of 5 years operational experience. In this study, we analyse interview data to identify the thought processes of analysts, including the questions they asked during their investigations and requirements for responses.

3.2 Analysis and Results

For each interview we attempted to understand how analysts identified what was happening and the information they needed to advance their investigations. Critical to this process is how analysts recognise and respond to situations. We analysed analyst interview statements, structuring them against the Recognition-Primed Decision (RPD) model [7], and found that the model is appropriate to capture and explain their processing of information in an investigation (Table 1). We propose that the RPD model, therefore, provides

Table 1: RPD Mapping from Interview Statements (Example from Interview 1)

Transcript Statement [CTA: A1, 11:30]	Goals	Cues	Expectancies	Actions	Why?	What for?
"*We had no idea initially what the kidnap was for. We were searching associates, we looked for any previous criminal convictions, we spoke to neighbours, and telephone information for his phone. One of the neighbours had suspected he had been kidnapped, and a witness had seen him being bundled into a car and alerted the police because they knew he was vulnerable.*"	Understand the motive, the risk to the victim, and possible suspects	Man gone missing. Thought he had been kidnapped due to witness report. known to be vulnerable	There is information for victim within existing databases	Searched known associates, looked for previous convictions, spoke to neighbours and witnesses, looked at telephone information.	To reduce scope of investigation and assess level of risk	To direct next steps of investigation and better use experience to recognise patterns
Extracted Questions:	Goals	Cues	Expectancies	Actions	Why?	What for?
What people are associates of victim?	Find associates	Victim name	The victim knows the offenders	Search for people connected to victim name	To find potential suspects	So that inquiries can be made into suspects
Does the victim have any previous convictions?	Find convictions	Victim name	The victim has been targetted before	Search for convictions directly linked to victim name	To understand past victimisation or offending	To assess risk and inform prioritisation
What calls have involved the victims phone?	Find calls	Victim phone number	The victim has been involved in recent calls	Search for calls involving phone number	To find recent communications	To identify possible leads or location

a concise and clear representation of an analyst's behaviour when retrieving information, and thus can be used to give an explanation structure for their intentions. We can design system processes that mirror this representation.

In Table 1, we also show how we extract individual questions asked by analysts from interview statements and can structure them against the RPD model. Furthermore, we can interpret the RPD attributes more generically to suit multiple questions of the same type. During the interviews each analyst provided many examples of their information needs and the questions that they asked when performing an investigation. For example, one analyst stated that "I looked through every database for the victim's name, custody records, PNC (Police National Computer), stop and search, vehicles he drove, to see if he had been stopped and searched with other people in the vehicle and if they had been named." [CTA: A1, 15:00]. From this statement, we can extract a number of questions posed by the analyst that could be directed towards a CA, including "how many vehicles have travelled to the victims address?" To answer this question the analyst provides cues for 'vehicles', 'travelled' and 'victims address'. Their goal is to retrieve summary information i.e. 'how many', and they are interested in finding a specific pattern of data in the database, which connect the cues. Table 2 provides a different example, with generic RPD attributes.

In this paper, we present how RPD attributes can be used to dynamically model analyst intentions for searching and retrieving information, through Formal Concept Analysis (FCA). FCA is an analysis approach which is effective at knowledge discovery and provides intuitive visualisations of hidden meaning in data [1]. FCA represents the subject domain through a formal context made of objects and attributes of the subject domain [14]. By breaking down analyst questions and structuring their components against the RPD model we extract attributes which can be used by a CA to process a response. In this study we identified specific RPD attributes which address over 500 analyst questions, akin to those described

Table 2: Example FCA-RPD Objects and Attributes

Recognition-Primed Decision Aspect	FCA Object: "Has [victim name] been reported in any activity?"
Cues	Pass specific input details (Victim Name, Activity)
Goals	Present confirmation
Expectancies	Expected that input details and pattern exist
Actions	Perform adjacent information search for entities extracted
Why?	Retrieve list for further exploration.
What for?	To find new lines of inquiry.

by analysts in interviews. We then performed FCA to identify intention concepts. In our case, the subject domain comprises the intentions of an analyst when they ask questions in an investigation. Therefore, FCA objects are questions including “Has [victim name] been reported in any activity?”. FCA attributes are the RPD model specifics which the CA must recognise and act upon in order to answer each question, such as the action ‘Perform adjacent information search’. Attributes can be simple methods, for example looking for single shortest paths or a pre-defined SPARQL pattern, or they can be more advanced capabilities, such as clustering similar instances. Importantly, each generic RPD attribute corresponds to a functional process and therefore can be developed as a module. FCA allows us to group modules together to form intentions, with question objects that can be used to train text classification for the user input to the CA.

The lattice, as shown in Figure 1, presents distinct object groupings. The final layer of concept circles are complete concept intentions, where all parts of the RPD are considered. The circles are sized based upon the number of associated questions. We can see that three questions in our set can be answered by combining the highlighted attributes. These attributes can answer the question, ‘how many vehicles are in our database?’, with ‘vehicle’ as a cue. The CA looks for adjacent information i.e. where there are instances of the class ‘vehicles’, presents a summary count, and retrieves a list. To provide transparency we propose we can simply present what attributes, and therefore functional processes, underpin a concept with their descriptions. Our model-agnostic and modular approach is akin to what Molnar [11] describes as the future of machine learning interpretability. We have used the concept lattice to define the intentions that an analyst can trigger through a CA interface, where each intention reflects our explanation structure; the RPD model.

4 UNDERSTANDING CA RESPONSES

4.1 Participants and Method

We interviewed four intelligence analysts with more than 10 years operational experience, from a different organisation to those interviewed previously. We aimed to explore their requirements for understanding the responses and processes of a CA in the context of a criminal investigation. Each interview lasted an hour and we presented interviewees with a series of questions and corresponding CA responses with two explanation conditions, switching the order of presentation. For one condition, responses described the data alone (1) and, in the other condition, the data and system processes (2). We were not attempting to test the differences between conditions, rather we used them as a starting point from which we could explore additional needs. Throughout interviews a single researcher took extensive notes from which individual statements were extracted. In total there were 114 distinct statements extracted, with counts for each analyst ranging from 24 to 34.

4.2 Analysis and Results

To analyse the statements we used an approach called Emergent Themes Analysis (ETA), as described by Wong and Blandford [19, 20], where broad themes, which are similar ideas and concepts, are identified, indexed and collated. ETA is useful for giving a feeling of what the data is about, with structure, and is fast and practical [10]. A single researcher analysed the statements and identified that they could be coded against the core functional components of a CA, for example ‘System Processes’ as shown in Table 3. From these components, we have drawn out the specific understanding needed for CA responses as sub-themes. The sub-themes are further categorised to form a general framework (Table 4) for explanation needs from an intelligent CA system.

Table 3: ETA Snapshot for Clarification of System Processes

Broad Theme	Sub-Theme	Framework Area	Statement
System Processes	Clarification of system inputs. Clarification of system processes	Clarification	I am concerned that info is missing because of search criteria. Understanding as a tool is also important for the whole system, such as when and where to use it. How have the results been worked out and what methods have been applied?

Table 4: CA Explanation Area Framework and Sub-Themes

Framework Area	ETA Sub-Themes
Clarification	Clarification of data attributes and structure, entity details, system input variables, metrics, question language, system processes, response methods, response language.
Continuation	Provide information to support continuation of investigation, including use of past interactions to move to next.
Exploration	Associated/additional data in responses or on periphery, intention match, system processes, source documents.
Justification	Provide information to justify selected system processes and the data defining the response.
Verification	Additional details for entities, correct intention match and impact/constraints of system processes. Check data reliability.

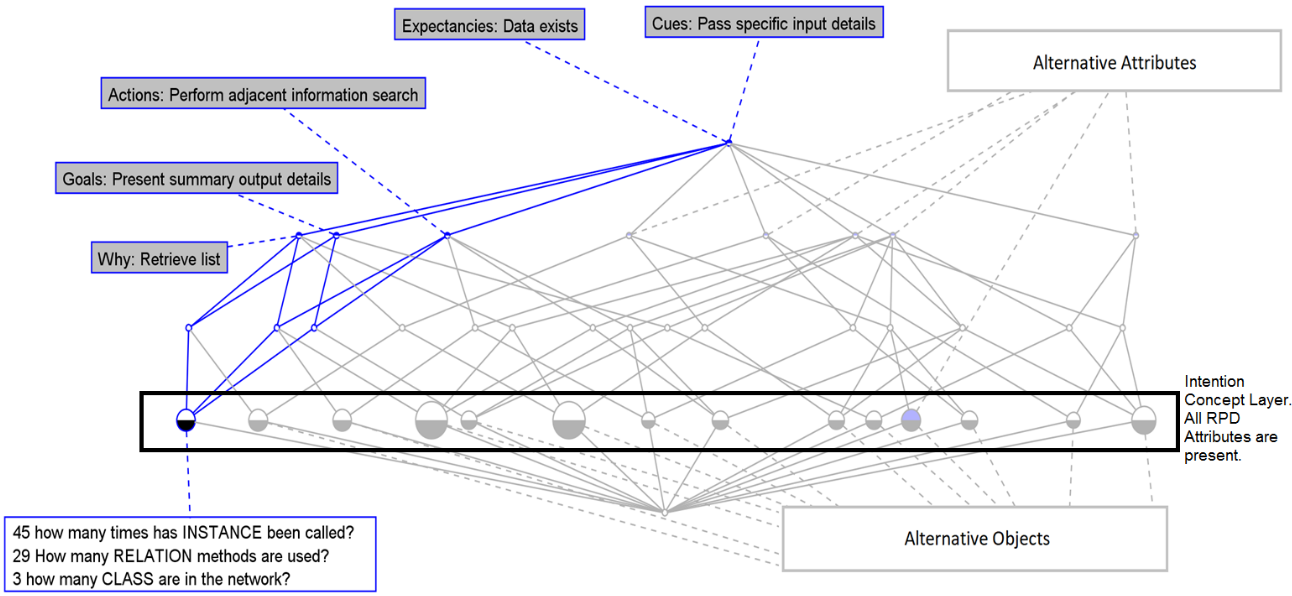

Figure 1: Concept Lattice for RPD Model Intentions (computed and drawn with Concept Explorer [21])

Exploring the interview data through the ETA method and structure is helpful when we come to design CA components. For example, examining Table 3 again, we can see that to provide understanding of system processes to an analyst we need to allow for clarification of both input variables and processes. Drawing upon details in the statements we can see that it is important to clarify any constraints related to the search inputs, the general capabilities of the system as a whole, and specific processes applied in any instance. We can incorporate explanations that provide clarification of these aspects, in addition to solutions for other themes extracted through ETA, into the design of our prototype application.

An analyst's ability to have clarification, verification and justification of system processes is crucially important, as identified by all analysts interviewed. This finding supports the framework for providing algorithmic transparency presented by Hepenstal et. al [3] and reiterates the need to go beyond traditional approaches to explainable AI (XAI) which focus upon explanations of the important features for a model and accuracy measures. Specific concerns included a need to justify follow up questions and the underlying rationale of the system for use in court (ETA: A1; Q2; C1). Additionally, an understanding of the system processes selected by the CA, including descriptions of the methods applied (all analysts, multiple statements), and inherent constraints, such as the questions which cannot be answered by the CA and information which has been omitted by the process (ETA: A2; Q1; C2 | A3; Q2; C2; | A4; Q4; C2). Essentially, analysts need to be able to justify, clarify and verify the CA intention triggered by their query and the related functional attributes. We believe our RPD explanation structure provides a neat mechanism to pick apart the system processes and provide, for each, the understanding required.

In Table 4, we display the framework areas and related sub-themes that emerged from ETA. Specific areas in the explanation framework can be linked to existing models for sensemaking, such as the Data Frame Model [8] for elaborating and reframing questions, or Toulmin's model for argumentation [17] to provide justification. Table 5 presents the key framework areas for each component theme, where at least two analysts made associated statements, together with a summary of sub-themes specific to both CA component and framework area. Different CA components draw more heavily on particular aspects of the framework and therefore our ETA analysis helps us to design and tailor explanations for each component.

5 CA PROTOTYPE

We have developed an initial prototype CA application called Pan, which uses FCA to define the different intention concepts to which it can respond. The objects (questions) which are attached to a concept are used as training data for machine learning text classification, so that a user's question can be matched to an appropriate intention. Each intention concept has associated attributes and we have developed methods to handle these as individual models, which create query syntax and interact with the database. In this way, FCA can combine multiple distinct combinations of attribute models flexibly to meet different analyst intentions. We propose that by combining model-based attributes with FCA to define intention concepts we provide a highly flexible approach to developing CA intentions. The objects and corresponding RPD attributes are critical for providing visibility to an analyst for the responses given by a CA and are akin to explainability scenarios i.e. "narratives of possible use that seek to answer the questions: who will use this

Table 5: CA Component Core Understanding Needs

CA Component Theme	Framework Area (common for multiple analysts)	Summary of Sub-theme(s)
Extracted Entities	Clarification + Verification (3)	More information of entities extracted for clarification and verification.
CA Intention Interaction	Clarification (3), Continuation (2)	Clear language to understand classification (i.e. no confusing response metric) and information to support continuation of investigation.
System Processes	Continuation (4), Verification (4), Clarification (3), Exploration (2), Justification (2)	User wants system understanding to support continuation of investigation, to allow them to verify processes are correct and explore them in more or less detail and justify their use/approach and constraints.
Data	Clarification (3)	Clarification of data updates and source, and data structure to aid forming questions.
Response	Clarification (4), Justification (4), Exploration (2)	Justification of response with underlying data, clarification of language (not trying to be human) and terminology, ability to explore results in more detail.

system and what might they need to know to be able to make sense of its outputs?" [18]

Our work to identify the core understanding needs for CA components has helped to inform the design of explanations for different parts of the system, for example, when the CA matches user input to an intention concept, triggers associated attribute models, and responds. The explanation provides information required for an analyst to understand the CA component themes of 'Data', 'Extracted Entities', and 'Response'. As an analyst types their query and entities are extracted, they are provided with identifier information where possible. We have also designed for the ability for an analyst to inspect and verify system goals and constraints. In our prototype we allow the user to step into the intention concept which has been triggered through a dialog window, so they can inspect and verify clear textual descriptions with our explanation structure, of the cues, goals, actions, expectancies and purpose of the intention. For example, when a concept triggers the action for finding single shortest path connections between instances, the analyst is presented with a description that includes any constraints to be wary of. Specifically, that it will not find longer paths or consider multiple routes. These caveats will impact how the analyst considers any information returned or how to rephrase their question. The attribute descriptions for each RPD module hang together as a narrative, akin to explainability scenarios. We intend to run experiments with Pan and operational intelligence analysts to validate our understanding of explanation needs and our RPD explanation structure for CA intentions.

6 USE CASES AND INITIAL FEEDBACK

In order for AI systems to be used for high risk and high consequence decision making they must provide transparency of their reasoning. As put by one analyst, "[the principal analyst] said none of my analysts would stand up in court where the beginning point of their evidence is an algorithm." [CTA: A4, 32:30] and that "You have to be able to trace it (your reasoning) all the way back to evidentially explain why you did each part... an analyst always has to justify what they have done, so does a system." [CTA: A4, 35:00] We believe that Pan addresses these issues by providing algorithmic transparency of its reasoning, using an architecture that aids recognition and explanations that meet our explanation framework. Early feedback from analysts on our approach is positive, opening routes for Pan to be tested in high risk and high consequence application domains where traditional CAs would not be deployed.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we describe our approach to capture and model analyst thought when retrieving information in a criminal investigation. We also present analysis to understand their needs for explanations from a complex CA system. Finally, we describe a prototype CA which incorporates FCA and RPD to build intention concepts and is therefore, we believe, transparent by design. We plan to evaluate the transparency impacts of our approach to intention concept design, gather additional requirements, and to validate our explanation framework through experimentation with operational analysts. To date we have not explored how a CA should present its responses to an analyst. Thus, we will look to explore how explanations are communicated, such as the specific textual or visual method.

The role of investigation scope was prominent in CTA interviews with analysts, where their questions were framed by the initial scope, thus introducing the risk that important information beyond the scope is missed. We will consider how CAs can help mitigate the constraints of investigation scope, through machine reasoning for example. Analysts expressed the desire to avoid obvious follow up questions, so it would be helpful for a CA to predict and explore additional questions autonomously. One approach for this is to model investigation paths as a Bayesian network. Transparency is a critical issue in autonomous systems and our explanation structure could help understanding by aiding the explanations of system behaviours across model states.

ACKNOWLEDGMENTS

This research was assisted by experienced military and police analysts who work for the Defence Science Technology Laboratory (Dstl) and the National Crime Agency (NCA) in the United Kingdom.

REFERENCES

[1] Simon Andrews, Babak Akhgar, Simeon Yates, Alex Stedmon, and Laurence Hirsch. 2013. Using formal concept analysis to detect and monitor organised crime.

[2] Leilani H. Gilpin, David Bau, Ben Z. Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining Explanations: An Approach to Evaluating Interpretability of Machine Learning.

[3] Sam Hepenstal, Neesha Kodagoda, Leishi Zhang, Pragya Paudyal, and B. L. William Wong. 2019. Algorithmic Transparency of Conversational Agents.. In *IUI ATEC*. Los Angeles.

[4] Sam Hepenstal, B. L. William Wong, Leishi Zhang, and Neesha Kodagoda. 2019. How analysts think: A preliminary study of human needs and demands for AI-based conversational agents.

[5] Bret Kinsella. 2018. Amazon Echo Device Sales Break New Records, Alexa Tops Free App Downloads for iOS and Android, and Alexa Down in Europe on Christmas Morning. https://voicebot.ai/2018/12/26/amazon-echo-device-sales-break-new-records-alexa-tops-free-app-downloads-for-ios-and-android-and-alexa-down-in-europe-on-christmas-morning/

[6] Bret Kinsella. 2019. NPR Study Says 118 Million Smart Speakers Owned by U.S. Adults. https://voicebot.ai/2019/01/07/npr-study-says-118-million-smart-speakers-owned-by-u-s-adults/

[7] Gary Klein. 1993. A Recognition Primed Decision (RPD) Model of Rapid Decision Making. (1993).

[8] Gary Klein, B. Moon, and R. Hoffman. 2006. Making Sense of Sensemaking 2: A Macroeconomic Model. *IEEE Intelligent Systems* 21.5 (2006), 88–92.

[9] G. A. Klein, R. Calderwood, and D. MacGregor. 1989. Critical decision method for eliciting knowledge. *IEEE Transactions on Systems, Man, and Cybernetics* 19, 3 (1989), 462–472.

[10] Neesha Kodagoda and William Wong. 2009. Cognitive Task Analysis of Low and High Literacy Users: Experiences in Using Grounded Theory and Emergent Themes Analysis. *Human Factors and Ergonomics Society Annual Meeting Proceedings* (2009).

[11] Christoph Molnar. 2019. *Interpretable Machine Learning. A guide for making black box models explainable.*

[12] Donald Norman. 1983. Design rules based on analyses of human error. *Commun. ACM* 26, 4 (1983), 254–258.

[13] Alun Preece, William Webberley, Dave Braines, Erin G. Zaroukian, and Jonathan Z. Bakdash. 2017. Sherlock: Experimental Evaluation of a Conversational Agent for Mobile Information Tasks. *IEEE Transactions on Human-Machine Systems* 47, 6 (2017), 1017–1028.

[14] Nadeem Qazi, B. L. W. Wong, Neesha Kodagoda, and Adderley Rick. 2016. *Associative search through formal concept analysis in criminal intelligence analysis.* Institute of Electrical and Electronics Engineers (IEEE).

[15] Nicole M. Radziwill and Morgan C. Benton. 2017. Evaluating quality of chatbots and intelligent conversational agents. *arXiv preprint arXiv:1704.04579* (2017).

[16] Danny Shaw. 2019. Crime solving rates 'woefully low', Met Police Commissioner says. www.bbc.co.uk

[17] Stephen E. Toulmin. 1958. *The Uses of Argument.* Cambridge University Press.

[18] Christine Wolf. Mar 17, 2019. Explainability scenarios: towards scenario-based XAI design *(IUI '19)*. ACM, 252–257.

[19] BL William Wong and Ann Blandford. 2004. Describing Situation Awareness at an Emergency Medical Dispatch Centre. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 48. SAGE Publications Sage CA: Los Angeles, CA, 285–289.

[20] William Wong and Ann Blandford. 2002. Analysing ambulance dispatcher decision making: trialing emergent themes analysis. *Vetere, F and Johnson, L and Kushinsky, R, (eds.) Ergonomics Society of Australia: Canberra, Australia.* (2002).

[21] Serhiy A. Yevtushenko. 2000. System of data analysis "Concept Explorer". (In Russian). Russia, 127–134.