Automated identification of insight seeking behaviours, strategies and rules: a preliminary study

Hepenstal, S., Zhang, L. and Wong, BL William

Accepted for publication.

Hepenstal, S., Zhang, L. and Wong, B. L. W. (2021) 'Automated identification of insight seeking behaviours, strategies and rules: a preliminary study', *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 65(1), pp. 1269–1273. doi: 10.1177/1071181321651348.

Automated identification of insight seeking behaviours, strategies and rules: a preliminary study

Sam Hepenstal[1], Leishi Zhang[2], and B.L. William Wong[2]
[1] Defence Science and Technology Laboratory, Porton Down, Salisbury, SP4 0JQ, UK
[2] Middlesex University London, The Burroughs, Hendon, London NW4 4BT, UK

In this paper, we demonstrate how insight seeking strategies and rules can be captured from analyst interactions with a question-answer system, as they perform an investigation. We present our analysis of an interactive investigation exercise undertaken by 14 experienced intelligence analysts. We propose that our approach to model the abstract higher order cognition involved in insight seeking provides a means to design intelligent systems that can reward and optimise potential lines of inquiry, ultimately creating the environment from which insights can be derived.

Key Words

Intelligent agents, Naturalistic decision making

INTRODUCTION

This paper provides an analysis of how insight seeking behaviours can be inferred from the questions asked by analysts, during investigations. We demonstrate how insight seeking strategies and the rules that lead to claims can be identified. We propose that these strategies and rules can be used by intelligent systems looking to explore lines of inquiry autonomously, to create the environment from which insights can be derived.

Criminal investigations involve a large amount of information and significant gaps where analysts have limited time to make critical decisions (Shaw, 2019). At the outset of an investigation an analyst may have very little understanding of the situation. They will need to ask many questions, the answers to each likely leading to further questions. Central to the investigation process is the aim of deriving insights that provide a deep understanding of a problem or situation, supporting analyst postulation about different ways to derive a solution. Insights enable analysts to make intelligent decisions and they often occur through exploration of the facts, by searching databases and identifying important information and information gaps, allowing the analysts to reach a 'claim'. Insights are difficult to observe and to document, as defined in the APA Dictionary of Psychology, insights are "the clear and often sudden discernment of a solution to a problem by means that are not obvious and may never become so, even after one has tried hard to work out how one has arrived at the solution." (Psychology, 2021). While it is not always possible to derive insights from the data available in a criminal investigation, it is reasonable to suggest that this is a desired outcome. Klein (Klein, A Naturalistic Study of Insight, 2011) notes the role of anchors and their manipulation to support insights. An anchor is an important data element or belief (Klein & Jarosz, A Naturalistic Study of Insight, 2011) and insights can emerge by spotting the implications of a new anchor, by finding contradictions within a set of anchors, or by discarding a weak anchor (Klein, A Naturalistic Study of Insight, 2011). While we cannot observe when an insight occurs, with the right system we can capture the ways that insight is sought.

In our system, we attempt to capture such insight seeking behaviours. When an analyst asks a question in an investigation to retrieve information, their questions contain anchors, and the manipulation of these anchors gives us an indication of their insight seeking behaviour. If we can capture these behaviours automatically while analysts interact with a system, we can learn about the insight seeking strategies they use in an investigation and the behaviours required to help them establish a plausible hypothesis or claim. Cognitive mimetics refers to the mimicry of higher order cognitive processes to help design intelligent technology (Kujala & Saariluoma, 2018). By understanding these higher order behaviours at an abstract level, independent to the underlying data and possibilities, we can inform an intelligent system to be able to reward or optimize lines of inquiry and to create the environment from which insight can be derived. The work described in this paper therefore provides the platform for future work in cognitive mimetics.

We have performed an experiment with 14 intelligence analysts where each analyst was asked to interact with a question-answer system to perform an investigation task. In this paper, we present the following contributions: (1) We describe qualitative analysis to interpret the insight seeking behaviours of analysts when they ask questions. (2) We describe a novel approach to automatically encode interactions with a question-answer system with the respective insight seeking behaviour. (3) We capture interesting and significant patterns in insight seeking behaviours, such as the rules that led to analyst claims, and investigation strategies used by analysts.

INTELLIGENCE ANALYSIS

There has been much attention given to the process of analysis and sensemaking. For this paper, we extend the research on decision making by criminal investigators presented by Gerber et al. (Gerber, Wong, & Kodagoda, 2016). Analysts and investigators are frequently confronted with fragmentary, out of sequence, missing, unknown, and ambiguous data. To make effective decisions amidst such uncertainty, the analyst looks to construct a narrative through abductive reasoning to describe the anchors that are perceived as being important, where anchors are supposed facts or entities. The analyst can then pose questions and gather additional data to help construct and challenge their developing narrative. In previous work, Hepenstal et al. demonstrated how possible lines of inquiry, supporting 'leap of faith', could be captured through interactions with a system (Hepenstal S. , Zhang, Kodagoda, & Wong, 2020). This work provided the states in an environment that could be explored incrementally by an intelligent agent, for example, possible follow-on questions, or levels of inquiry, based upon the information used in a question and the intent. However, the environment merely captured the possibilities and did not consider how lines of inquiry could be intelligently pursued, optimised or rewarded. A simple approach to select and explore possibilities in sequence, moving from one set of results to the next until no more information is found, would not provide an accurate reflection of information retrieval, and may not create the necessary conditions for insights. Analysts do not continuously seek new information by querying entities returned in results. Instead, they may look to cross-check new information with the important entities they are already aware of, to see if this fits within their narrative and to challenge their abductive reasoning. Analysts adjust their assumptions and utilise their expertise to actively seek insights, guided by cognitive strategies and their expectations, to provide the right conditions from which insights can be gained. A real investigation involves anchor manipulation to support abductive reasoning. The importance of the manipulation and management of anchors when triggering insight is described by Klein (Klein, A Naturalistic Study of Insight, 2011), where insight strategies are represented in a triple path 'Anchor model of insight'. The three triggers and associated activities are: (1) Contradiction: an inconsistency is found and a weak anchor is then used to rebuild a story. (2) Connection: an implication is spotted and a new anchor is added to the set. (3) Creative desperation: an anchor is discarded, to escape an impasse.

EXPERIMENTAL DESIGN

Purpose: The aim of this study was to understand how analysts queried data in search of insights, when performing an investigation. Specifically, we wanted to identify how insight seeking could be defined and what strategies and rules could be identified.

Participants: 14 operational intelligence analysts were recruited from a range of organisations, such as the police, National Crime Agency (NCA), military, and prison service. The only inclusion criteria was that the analysts had a minimum of 3 years full time experience in a role involving network analysis. Several of the analysts actually had more than 10 years' experience.

Equipment: In this study, analysts interacted with a prototype system called Pan. Pan is a question-answer system for retrieving information and exploring relationships between entities, described by Hepenstal et al. (Hepenstal, Zhang, & Wong, Pan: Conversational Agent for Criminal Investigations, 2020). A user could pose a natural language question to Pan, and the system attempted to provide an answer through classification of user intent. Analysts could observe and interact with the data retrieved by Pan through a network graph visualisation.

Task and Procedure: The task was for each analyst to individually perform an investigation to arrive at a hypothesis about a realistic scenario. The scenario data was synthetic, based upon an actual investigation to identify the owner of a mobile phone (IDMOB1) involved in an illegal firearm purchase. The analysts were first asked to review some briefing material, which described the scenario and the Pan application. They were then asked to complete an investigation, by posing questions to the system until they reached an initial hypothesis, or claim, that they were comfortable to describe. All were provided with the same introductory information to explain the scenario, including important entities (a known offender called 'Dan Govey', his phone 'IDMOB2', and an Organization called 'DGX Bodywork'). In an investigation it is usually the case that multiple plausible hypotheses can be reasoned from the available data. A participant was considered to have successfully completed the exercise when they were able to demonstrate an understanding of the fragmented and ambiguous pieces of information by being able to construct a plausible hypothesis and present their reasoning based upon information they had retrieved. As in the real-world, there were numerous plausible hypotheses. What was important for our study was how the investigator searched for the data that would help them reason about the situation and form their hypothesis. All the analysts were able to successfully complete the exercise. Face-to-face experiments were not possible due to the Covid-19 pandemic. Instead, a researcher shared their screen through virtual conferencing software and each analyst voiced their questions for the system. The researcher followed a checklist in each interview to ensure the same steps were covered for all analysts. The audio from the interviews was recorded, transcribed and analysed by a single researcher, who reviewed the data multiple times for accuracy.

Data Collection: Pan collected data as the analysts interacted with it, including the questions posed and the data retrieved. In this study, we draw from human factors research to understand and encode the phenomenon of insight seeking behaviours and apply analysis methods to enable the automatic discovery of insight seeking strategies and rules.

DATA ANALYSIS

Insight seeking behaviours: The analysts fluidly followed lines of inquiry and manipulated the anchors used in their questions. The results of one question sometimes led to further questions, if they found information that was interesting, and the results of those to even more. In this way, each investigation explored various levels of questioning. We have encoded the level of a question and its results automatically, by considering the anchors used in the question and the relative level at which they were first retrieved, as shown in Table 1. Anchors taken from the briefing material were assigned a level of 0. Whilst it was not possible to capture the insights gained by each analyst, by identifying the level of the anchors they were using, relative to one another, we could automatically encode the insight seeking behaviour of the analyst when they asked a question. If, for example, an analyst used a new anchor from the results of a question at the most current level of the investigation, this indicated that they were seeking implications for that new anchor. If in the next question, they reverted to ask about an anchor from an earlier level of their investigation, then they were escaping impasse and had removed the most recent anchors from their line of inquiry. In this way, we aligned analyst behaviours to the triple path model (Klein, A Naturalistic Study of Insight, 2011) through their manipulation of anchors in the questions they asked. Insight seeking behaviours were automatically encoded based upon the relative anchor levels used, as shown by the examples in Table 2. By automating the capture of insight seeking, it provided a platform to learn about analyst strategies and rules through their interactions with a system.

Table 1: Encoding anchor investigation levels.

Question	***Utterance***	*Anchor and level*
First question (Q1)	'*So, are there any individuals associated with DGX Bodywork?*' [Analyst 5; 18:35]	The analyst asked a question about DGX Bodywork, an organisation that was presented to the analyst as being potentially important in the briefing material. This was the starting point for the investigation, so all entities in the briefing material were encoded as level 0.
Second question (Q2)	'*What other information is there for IDDarrenSmith?*' [Analyst 5; 19:25]	Darren Smith was found to be employed by DGX Bodywork, in the results to Q1. In Q2, the analyst has used this entity in their question progressing the investigation to the next level. The entity for Darren Smith (IDDarrenSmith) is encoded as level 1.
Third Question (Q3)	'*What other information is there for IDShaunSpence?*" [Analyst 5; 19:35]	Shaun Spence was also found to be employed by DGX Bodywork, in the results to Q1. The entity IDShaunSpence is therefore also encoded as level 1.

Table 2: Insight seeking utterances.

Insight seeking behaviour	***Utterance***	***Use of anchors***	***Insight encoding***
Seeking contradictions	'*OK, so the next obvious question is, is Paul Richards connected to Dan (Govey)?*' [Analyst 1; 12:34]	Check if new anchor (Paul Richards) is consistent with developing narrative.	Question involves anchors from highest recent level and lower level(s).
Seeking implications	'*What mobile phones are associated to Shaun Spence?*' [Analyst 2; 10:00]	Check implications / connections for new anchor (Shaun Spence).	Question involves anchors from highest recent level.
Creative desperation	'*If we go back to Dan Govey. Who has IDMOB2 been in contact with?*' [Analyst 6; 10:25]	Revisit earlier stage of investigation (remove / ignore most recent anchors).	Question involves anchors from lower level than highest recent level. Highest recent level resets.

Insight seeking rules: In this paper, we have looked to identify rules, in terms of the necessary sequences of insight seeking behaviours that led to analyst claims. These rules are important to help us to understand what sequences of behaviours may help an analyst to reach a conclusion when performing an investigation. There are numerous approaches for automatic sequential rule mining, as described by (Fournier-Viger, et al., 2016). For example, the RuleGen algorithm (Zaki, 2000) can take a sequence, which is a list of itemsets, and output rules (sequential relationships between two patterns X and Y) based upon two user-specified thresholds, support and confidence. The support of a rule X==> is the number of sequences containing Y, and the confidence is the number of sequences containing X divided by the number of sequences containing Y. The results of the RuleGen algorithm outputs all sequential rules that have a support and confidence greater or equal to the respective thresholds. We wanted to identify all sequence rules that led to a claim, and have used the RuleGen algorithm (calculated with SPMF (Fournier-Viger, et al., 2016)) to find rules in the investigation sequence data. Insight seeking behaviours were recorded in the same itemset if they involved anchors from the same level of the investigation, in sequence. If anchors from a different level were considered, then we created a new itemset. The rules did not consider whether data was retrieved. They described the necessary insight seeking steps, and anchor manipulations, that were common across investigations prior to a claim. Rules were generated with a minimum support of over half the analysts (8) and 80% confidence.

Table 3: Examples of analyst claims, demonstrating completion of the investigation exercise.

Analyst	***Utterance demonstrating claim***
A2	'*Well I'm starting to lean towards Paul Richards. But that is based purely on the domestic assault situation, rather than necessarily going through all of this in finer detail. From the information at hand, who has motivation for violence. The other way would be to look at Susan Leech, but I find it unlikely she would be communicating with herself, having two phones to that extent.*' [27:50]
A11	'*You would assume therefore that because there was a lot of contact between IDMOB1 and IDMOB3 that they were in a relationship and therefore potentially IDMOB1 is owned by Paul Richards.*' [19:00]

Insight seeking strategies: An analyst's knowledge builds gradually throughout an investigation to culminate in a claim. The claim can be described as a narrative that underpins the relationships between important anchors, as identified by the analyst. Analysts completed the exercise when they could make a claim about the potential owner of 'IDMOB1'. Some examples of claims are shown in Table 3. The questions that guide an analyst to build their narrative and reach a claim are not a hurdle to remove to speed insight; they are the mechanism by which the environment for insight is constructed. While RuleGen helps identify insight seeking behaviours that exist prior to a claim, this does not fully capture the nature of relationships between insight seeking behaviours. In particular, the different strategies that analysts may apply to gather and construct a narrative around anchors, which are also influenced by whether or not data is retrieved. We looked to identify what strategies were applied by analysts. In this paper, we analysed common sequences of three linked behaviours. We chose three behaviours, because this created a large enough sequence to be interesting while also allowing for significant patterns to be found in our small dataset.

The network of behaviours was formed by linking questions together as they occurred in sequence. We created separate subgraphs for every sequence of three behaviours and the relationships between them. Table 4 shows how analyst questions were combined to form sub-graph strategies, capturing the relationships between anchors used, insight seeking behaviours, and whether data was returned. We then constructed a randomised network of insight seeking behaviours, with the possible relationships between nodes derived from the real data. We created random strategies until at least 80% of strategies found in the real data were represented in the randomised data. Finally, we identified subgraphs that were significantly overrepresented in the actual network compared to the randomised network, otherwise known as network motifs. To do this, for each motif, we tested the difference between the proportions reflected within the randomised network and the actual network. We were only interested in finding those motif strategies that were used by analysts for a higher proportion of the network than random, so we calculated a one-sided p-value and selected significant motifs where $p < 0.01$.

Table 4: Defining three-part insight seeking strategies.

Analyst	*Utterance*	*Part*	*Strategy level*	*Insight seeking*
A8	‘*Ok. So, I guess further questions are, do we know who owns MOB3?*’ [12:00]	1	The start of the strategy is at level 0. Data was returned.	Seek implications from level 0
A8	‘*and do we know who owns MOB4?*’ [12:10]	2	The anchor at this level was found at the same level of the investigation as part 1, therefore is still strategy level 0. No data was returned.	Seek implications from level 0
A8	‘*does she have any link to Dan Govey?*’ [12:35]	3	The anchor in this question is at a higher level that part 2, so is strategy level +1.	Seek contradictions from level 1

Table 5: Comparison of the 7 analysts who met the most rules and the 7 who met the fewest. Where cells are green, there is a significant difference between means ($p<0.05$). Where they are red, the difference is not significant ($p>0.05$).

Analyst group	*No. rules met (avg [st. dev])*	*No. questions (avg [st. dev])*	*Data retrieved (avg instances [st. dev)*	*Levels explored (avg [st. dev])*
Top 7	299.43 [7.72]	16.43 [3.42]	12.57 [1.99]	6.71 [1.48]
Bottom 7	114.29 [73.37]	13 [2.33]	9.86 [0.99]	4.86 [1.25]

RESULTS AND DISCUSSION

Our rule analysis identified 1034 rules in total and 313 claim rules (that included a claim within the sequential rule). These claim rules were our focus, given that we wanted to understand the behaviours that appeared in a sequence when a claim also appeared. The rules identified interesting sequences of insight seeking behaviours that were common across analysts before they reached a claim. An example of a rule, with utterances and insight seeking behaviours that meet it, is shown in Table 7. Some rules were simple, for example, that all the analysts sought implications at some point before reaching a claim. Other rules were more complex, for example, that when 12 analysts escaped impasse before reaching a claim, they had previously sought implications three times, moving to a different level in the investigation for each (confidence 0.92). We compared the 7 analysts who met the most rules with the 7 who met fewest (Table 5). We found that those analysts who met the most rules had a greater depth and breadth of information from which to draw a claim. We have also compared the number of questions asked by the two groups of analysts, to check whether more rules were met simply by asking more questions. In Table 5 we can see that the number of questions asked by the two groups was not significantly different. However, those who met more rules did cover significantly more breadth and depth in their investigations. Therefore, we propose that the number of rules provides an abstract metric to assess whether the information retrieved by following lines of inquiry will create the necessary environment for insights. We have captured significant insight seeking patterns in the analyst investigations. These represent deliberate insight seeking strategies, rather than random behaviours. There were 29 significant strategies gathered from our investigation data. On inspection, these make logical sense. For example, the significant subgraph that analysts: (1) Seek implications from the current level of the investigation (level 0), returning data. (2) Seek implications with the results (level +1), returning data. (3) Seek inconsistencies with the results (level +2). This insight seeking strategy is reflected in the qualitative data in Table 6.

Table 6: Example of a significant motif strategy used by 4 analysts.

Analyst	*Utterance*	*Part*	*Response*	*Insight seeking*
A4	‘*What is Susan Leech linked to?*’ [10:35]	1	Data is found, results include a domestic assault activity.	Seek implications from level 0
A4	‘*So, I'd want to know who the offender was (in the domestic assault activity).*’ [11:25]	2	Data is found, results include a person called Paul Richards.	Seek implications from level +1
A4	‘*How is Paul Richards linked to Dan Govey (a known anchor)?*’ [12:10]	3		Seek contradictions from level +2.

Analysts face a wide variety of situations, and we propose the insight seeking strategies we have captured are applicable across investigation domains. We can see that the strategy described above is not specific to our experiment scenario, by comparing with a real investigation provided by Hepenstal et al. (Hepenstal S. , Zhang, Kodagoda, & Wong, 2020). In the real scenario an analyst explained that, once they receive a suspect phone number, they “go and find other phone calls” involving that number. They seek further information from which they can draw connections or implications for new anchors. When they find phone calls, and new phone numbers, they then “get call data for others in the network”. They use the results to add further anchors. Finally, for any results they “also check all the numbers additional people have phoned against all other numbers” in the investigation databases. They seek inconsistencies or contradictions for the new anchors they have gained, in relation to their current understanding of the investigation.

CONCLUSION

In this paper, by bringing together qualitative and quantitative approaches, we have developed a method to capture insight seeking behaviours, strategies and rules in an investigation scenario. We propose that our findings provide the foundations to develop an intelligent system that can create the environment for insights and supports abductive reasoning, by understanding the context of how insight is sought, and claims are made. For example, the strategies and rules could provide an abstract means to optimise information search and retrieval processes. By understanding the way that analysts seek insights, through manipulating anchors, we can build upon previous work (Hepenstal S. , Zhang, Kodagoda, & Wong, 2020) to explore investigation paths autonomously and create the conditions required for an analyst to derive insights. Our initial findings are limited, having been derived from a single investigation scenario exercise. We have considered one fixed structure for insight seeking strategies and we have not explored the significance of other lengths of subgraph pattern. We will look to explore this further in future work.

REFERENCES

Fournier-Viger, P., Gomariz, A., Gueniche, T., Soltani, A., Wu, C.-W., & Tseng, V. S. (2016). The SPMF Open-Source Data Mining Library Version 2. *In Proceedings of the 19th European Conference on Principles of Data Mining and Knowledge Discovery (PKDD 2016)* (pp. 36-40). Springer.

Gerber, M., Wong, B. W., & Kodagoda, N. (2016). How Analysts Think: Intuition, Leap of Faith and Insight. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (pp. 173-177). Sage.

Hepenstal, S., Zhang, L. K., & Wong, B. W. (2020). Pan: Conversational Agent for Criminal Investigations. *Proceedings of the 25th International Conference on Intelligent User Interfaces Companion (Cagliari, Italy) (IUI '20)* (pp. 134-135). Cagliari: ACM.

Hepenstal, S., Zhang, L., Kodagoda, N., & Wong, B. W. (2020). Providing a foundation for interpretable autonomous agents through elicitation and modelling of criminal investigation pathways. . *In Proceedings of the 64th Human Factors and Ergonomics Society Annual Meeting.* Sage.

Klein, G. (2011). A Naturalistic Study of Insight. *Journal of Cognitive Engineering and Decision Making*, 335-351.

Klein, G., & Jarosz, A. (2011). A Naturalistic Study of Insight. *Journal of Cognitive Engineering and Decision Making*.

Kujala, T., & Saariluoma, P. (2018). Cognitive Mimetics for Designing Intelligent Technologies. *Advances in Human-Computer Interaction*.

Psychology, A. D. (2021). *https://dictionary.apa.org/*. Retrieved from APA Dictionary of Psychology: https://dictionary.apa.org/insight

Shaw, D. (2019). *Crime solving rates 'woefully low', Met Police Commissioner says*. Retrieved from www.bbc.co.uk: https://www.bbc.co.uk/news/uk-48780585

Zaki, M. J. (2000). Scalable algorithms for association mining. *IEEE Transactions on Knowledge and Data Engineering*, 372–390.

Table 7: Example sequential rule and utterances for insight seeking behaviours that occurred prior to a claim.

	Example rule, supported by 13 analysts with 100% confidence: When analysts sought implications twice at the same level of the investigation (same itemset), they later escaped impasse (from a different itemset) prior to reaching a claim. Examples of utterances are shown below, for 4 analysts.		
Insight seeking description.	Sought implications twice at same level of investigation (the entities used in questions were encoded with the same investigation level).	Escaped impasse at a later stage. After pursuing further levels of the investigation, the analyst came back to an anchor they had identified at an earlier stage of the investigation:	Reached claim. After further questions, the analyst said:
Analyst 6	“*Right so who is the owner of IDMOB3 then?*” [8 mins 10 seconds]… “*and who is the owner of IDMOB4?*” [8:45]	“*If we go back to Dan Govey. Who has IDMOB2 been in contact with?*” [10.25]	“*So probably my hypothesis would be Mark Watson would be said person. Probably who I'm leaning towards.* [22.10]
Analyst 14	“*Who lives at ID3MitchinDrive?*” [10:30]… “*What do you know about IDPaulRichards?*” [10:45]	“*Are there any direct links between IDDanGovey and IDPaulRichards?*” [12:10]	“*My assumption would be that IDMOB1 is owned by Paul Richards*.” [17:10]
Analyst 2	“*So, can I ask what mobile phones are associated to IDShaunSpence.*” [10.00]… “*Ok and same question for Darren Smith*” [10.30]	“*What people are associated to IDMOB3?*” [21:25]	“*I'm starting to lean towards Paul Richards.*” [33:00]
Analyst 10	“*Can we see who lives at 3 Mitchin Drive, if there is anyone else*?” [12:20] … “*OK. What information is available on Richards*?” [14:55]	“*What information is available on Shaun Spence*?” [18:30]	“*My top suspect would be Shaun Spence, because of his connection to the OCG as well as Richards.”* [23:40]