Middlesex University Research Repository

An open access repository of
Middlesex University research

<http://eprints.mdx.ac.uk>

Hepenstal, Sam, Zhang, Leishi ORCID logoORCID: <https://orcid.org/0000-0002-3158-2328>,
Kodagoda, Neesha and Wong, B. L. William ORCID logoORCID:
<https://orcid.org/0000-0002-3363-0741> (2020) Pan: conversational agent for criminal
investigations. IUI '20: Proceedings of the 25th International Conference on Intelligent User
Interfaces Companion. In: IUI '20: 25th International Conference on Intelligent User Interfaces,
17-20 Mar 2020, Cagliari, Italy. ISBN 9781450375139. [Conference or Workshop Item]
(doi:10.1145/3379336.3381463)

Final accepted version (with author's formatting)

This version is available at: <https://eprints.mdx.ac.uk/29593/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this work are retained by the author and/or other copyright owners unless otherwise stated. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge.

Works, including theses and research projects, may not be reproduced in any format or medium, or extensive quotations taken from them, or their content changed in any way, without first obtaining permission in writing from the copyright holder(s). They may not be sold or exploited commercially in any format or medium without the prior written permission of the copyright holder(s).

Full bibliographic details must be given when referring to, or quoting from full items including the author's name, the title of the work, publication details where relevant (place, publisher, date), pagination, and for theses or dissertations the awarding institution, the degree type awarded, and the date of the award.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

See also repository copyright: re-use policy: <http://eprints.mdx.ac.uk/policies.html#copy>

Pan: Conversational Agent for Criminal Investigations

Sam Hepenstal

Defence Science Technology Laboratory, UK

Neesha Kodagoda

Middlesex University London, UK

Leishi Zhang

Middlesex University London, UK

B. L. William Wong

Middlesex University London, UK

ABSTRACT

We present an early prototype conversational agent (CA), called Pan, for retrieving information to support criminal investigations. Our approach tackles the issue of algorithmic transparency, which is critical in unpredictable, high risk, and high consequence domains. We present a novel method to flexibly model CA intentions and provide transparency of attributes that is underpinned with human recognition. We propose that Pan can be used for experimentation to probe analyst requirements and to evaluate the effectiveness of our explanation structure.

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. Pan: Conversational Agent for Criminal Investigations. In *25th International Conference on Intelligent User Interfaces Companion (IUI '20 Companion)*, March 17–20, 2020, Cagliari, Italy. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3379336.3381463>

1 INTRODUCTION

Criminal intelligence analysis requires repeated questioning of available information as analysts look for patterns, strengthen their interpretation of modus operandi, question specific details and predict paths in the data [3]. This can be time consuming, particularly when manual formulation of query syntax is restrictive, cumbersome and complex. 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” [10]. A more natural interaction with data through a conversational agent (CA) could speed up this process and provide significant benefits. One analyst interviewed in our previous work commented that “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

This material is licensed under the terms of the Open Government Licence except where otherwise stated. To view this licence, visit <http://www.nationalarchives.gov.uk/doc/open-government-licence/version/3> or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk
IUI '20 Companion, March 17–20, 2020, Cagliari, Italy
© 2020 Crown copyright (2019), Dstl.
ACM ISBN 978-1-4503-7513-9/20/03
<https://doi.org/10.1145/3379336.3381463>

Figure 1: System Components

can answer me that in one minute, literally those nineteen minutes could save someone’s life.” [5]

Criminal intelligence analysis is a high risk and high consequence domain where decisions can have serious consequences, thus ethical considerations are important [1]. We focus on addressing the issue of algorithmic transparency, defining transparency as a combination of (i) the explainability of results and (ii) the ability of a user to inspect and verify system goals and constraints [4]. Our prototype can perform specific tasks related to information retrieval and we can therefore use it in experimentation to further identify requirements and evaluate the effectiveness of our explanation structure, which is aligned to human recognition.

2 SYSTEM OVERVIEW

Intelligent artificial agents come in a variety of forms and levels of complexity [11]. To support information retrieval in a criminal investigation our focus is on developing a reflex CA that can respond based upon some rules, for example by applying an appropriate search method. Typical CAs understand users by matching their input pattern to a particular task category (intention), for example using ‘Artificial Intelligence Markup Language’ [9]. A commercial example is Language Understanding (LUI) [8] which matches ‘utterances’, which are statements made by a user, to ‘intentions’, where a matched intention triggers a specific set of rules.

We use a similar intention matching approach, however, rather than trying to manually identify all possible analyst intentions and appropriate rules, they are derived dynamically. We use Formal Concept Analysis (FCA), a mathematical method to transform tabular data into lattice structures [2], to construct intentions from the attributes (rules) required to answer questions. This allows us to pick apart discrete modules of code, which can then be explained to the user. We have encoded modules with the recognition aspects of

Figure 2: Concept Lattice for RPD Model Intentions (computed and drawn with Concept Explorer [12])

Table 1: 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.

the Recognition-Primed Decision (RPD) model [6] in an attempt to mirror analyst thought when recognising situations related to information retrieval tasks. We propose that the use of the RPD gives an explanation structure that will aid recognition and understanding of system processes within a CA intention. The system overview is displayed in Figure 1, where our database is a knowledge graph and attribute models build and process query syntax and results.

3 TRANSPARENCY DESIGN

We conducted Cognitive Task Analysis (CTA) interviews, applying the Critical Decision Method [7], with four intelligence analysts to delve into a particularly memorable incident for each [5]. Throughout the interviews analysts provide many questions which they ask in an investigation to retrieve information. In this study we identified over 500 differently phrased questions with a variety of information retrieval steps required.

3.1 Intention Modelling

From each question we picked out the functional attributes, or rules, required to provide an answer. These functional processes are structured against the RPD model to aid intention recognition and understanding (Example in Table 1). We then performed FCA across all questions to identify intention concepts, where each question is the object and each functional module required is an attribute. The lattice is shown in Figure 2. The final layer of concept circles

represents complete concept intentions, where all parts of the RPD are considered. The circles are sized based upon the number of associated questions.

3.2 Intention Transparency

Pan provides both explainability and system visibility. Explanations are given in responses to user questions and convey information such as the underpinning data. Users can explore this data interactively. The FCA approach to concept definition provides a clear structure to give visibility of a complex system by underpinning CA intention concepts and processes with the context of human recognition i.e. the RPD model. 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 textual descriptions 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 inform how the analyst interprets any information returned or how to rephrase their question.

4 CONCLUSION

We have developed an early stage prototype CA with a flexible intention architecture which is designed for human recognition. This approach places transparency at the core of the system, rather than as an add on. We intend to use our prototype as a probe to explore transparency issues for complex decision aids and potential solutions, with users working in the domain of criminal intelligence analysis. We will seek to evaluate the effectiveness of our RPD explanation structure.

REFERENCES

- [1] Penny Duquenoy, Donald Gotterbarn, Kai Kimppa, Norberto Patrignani, and B.L. Wong. 2018. *Addressing Ethical Challenges of Creating New Technology for Criminal Investigation: The VALCRI Project*. 31–38.
- [2] Bernhard Ganter, Gerd Stumme, and Rudolf Wille. 2005. *Formal concept analysis foundations and applications*. Springer, Berlin.
- [3] Matylda Gerber, B.L. William Wong, and Neesha Kodagoda. 2016. How Analysts Think: Intuition, Leap of Faith and Insight. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 60, 1 (2016), 173–177.
- [4] Sam Hepenstal, Neesha Kodagoda, Leishi Zhang, Pragya Paudyal, and B.L. Wong. 2019. Algorithmic Transparency of Conversational Agents. *Joint Proceedings of the ACM IUI 2019 Workshops (ACMUI-WS 2019)*. In: *IUI 2019 Workshop on Intelligent User Interfaces for Algorithmic Transparency in Emerging Technologies*.
- [5] Sam Hepenstal, B.L. Wong, Leishi Zhang, and Neesha Kodagoda. 2019. How analysts think: A preliminary study of human needs and demands for AI-based conversational agents. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 63 (11 2019), 178–182.
- [6] Gary Klein. 1993. *A Recognition Primed Decision (RPD) Model of Rapid Decision Making*.
- [7] 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.
- [8] Microsoft 2020. *Language Understanding (LUIS)*. Microsoft. <https://www.luis.ai/home>
- [9] Fernando A. Mikic, Juan C. Burguillo, Martin Llamas, Daniel A. Rodriguez, and Eduardo Rodriguez. 2009. CHARLIE: An AIML-based chatterbot which works as an interface among INES and humans. *EAAEE Annual Conference, 2009*, 1–6.
- [10] Danny Shaw. 2019. Crime solving rates ‘woefully low’, Met Police Commissioner says. <https://www.bbc.co.uk/news/uk-48780585>
- [11] Michael Winikoff, Lin Padgham, and James Harland. 2001. Simplifying the Development of Intelligent Agents.
- [12] Serhiy A. Yevtushenko. 2000. System of data analysis “Concept Explorer”. (In Russian). Russia, 127–134.