An Assistive System for Transferring Domain Knowledge to Novice Officers

Héctor López-Carral

Laboratory of Synthetic, Perceptive, Emotive and Cognitive Systems (SPECS), Institute for Bioengineering of Catalonia (IBEC), The Barcelona Institute of Science and Technology (BIST), Barcelona[1]

Paul FMJ Verschure

Donders Institute for Brain, Cognition and Behaviour – Donders Centre of Neuroscience, Radboud University, Nijmegen[2]

Abstract:

Instructional strategies in many operative fields, including law enforcement, have reached a high level of complexity due to dynamically changing task environments and the introduction of different technologies to help users in their operational work. In the last decades, a transition has been observed from dedicated trainers to the adoption of automated technologies to support the trainees. Based on a review of state-of-the-art literature and direct feedback from law enforcement agencies, we have developed an assistive system to aid in the knowledge transfer from expert to novice officers and, consequently, improve the time necessary to train new police practitioners. This system is grounded on the most relevant instructional principles derived from cognitive and learning theories. The result is a system that can dynamically deliver suggestions based on previous successful actions from other users and the current performance and state of the user. To validate our system, we implemented a knowledge graph exploration task. The novel knowledge transfer system is introduced here by presenting the results from our literature review, explaining the architecture of the assistive system, and discussing our observations from the validation task. With this work, we aim to facilitate the transfer of domain knowledge, which could have a significant impact on the training and education of law enforcement officials in and for the Digital Age.

Keywords: assistive system; knowledge transfer; training; recommender system; crime investigation; knowledge graph.

1 Author's email: hlopez@ibecbarcelona.eu

2 Author's email: paul.verschure@donders.ru.nl

Introduction

Instructional strategies in many operative fields have reached a high level of complexity due to dynamically changing task environments and the introduction of different technologies to help users in their operational work. In the last decades, a transition has been observed from dedicated trainers to the adoption of automated technologies to support the trainees. This paradigm shift makes transferring precise knowledge to novice users a challenging problem, which becomes especially relevant when the user is dealing with large and complex datasets from which to extract relevant information.

Supportive technologies, such as recommendation systems, have attracted a lot of interest in the last decades, both in the industry and the academia. The goal of such systems is to help the users to reduce the burden imposed by the high information load that is intrinsic to the exploration of large and complex datasets by providing valuable suggestions in the form of specific items or possible actions to choose from. Despite clear technical advances witnessed in the field in improving the accuracy of the recommendations, several challenges and open issues remain, especially regarding the specific role of various human factors.

Among the functionalities that were identified to provide Law Enforcement Agencies (LEAs) with a set of automated tools and systems to boost the investigative work in the fight against illicit trafficking activities, one was the capability to provide adequate solutions facilitating the transfer of the acquired expertise among experienced users and, consequently, boost the take-up time necessary to train new users. In order to accomplish this task, we decided to build a novel assistive system, which, combining practical knowledge from classical recommender systems with theoretical knowledge from cognitive systems, is able to aid in the transfer of domain knowledge to novice officers.

We will discuss, firstly, the recommender systems in general before outlining the recommender system for assisting knowledge transfer that reflects the best practices, approaches, and directions in the respective law enforcement domain. Our recommender system is conceptually grounded in a cognitive architecture, learning from interactions to later assist novice users by suggesting key pieces of information that other users have selected. Then, we describe the case used for validating this system in a knowledge graph exploration task based on a novel interface for LEAs to present the collected information in a criminal investigation. Finally, we will put forward our conclusions and outline possible next steps.

Introduction to recommender systems

Recommender systems have been used extensively in research and industry since the mid-1990s (Goldberg et al., 1992). The most common domain for their use is electronic commerce (e-commerce), the entertainment and media industry, and services. Many online businesses employ dedicated algorithms to provide recommendations to their customers based on inputs such as their history of items visualised and purchased or their demographic data. Another popular area in which recommender systems are used is multimedia applications (Ge & Persia, 2017). For example, many online music platforms use them to recommend songs or artists based on what each individual listens to (Song, Dixon & Pearce, 2012). Similarly, recommendation systems are common in online video platforms to provide personalised suggestions for TV shows, movies, and other videos (Asabere, 2012).

Several types of recommender systems have been proposed that, depending on the techniques employed, can be classified into different categories (Park et al., 2012; Villegas et al., 2018; Ricci et al., 2011; Adomavicius et al., 2011). In content-based recommendation schemes, the system learns to propose items similar to those that were preferred in the past by the same user. In contrast, collaborative filtering approaches recommend items that other users with similar profiles have preferred in the past. Knowledge-based systems recommend options based on specific domain knowledge about how certain features meet users' needs and preferences. Finally, hybrid systems are based on the combination of the techniques mentioned above to improve performance (Burke, 2002).

Despite providing varying degrees of support, overall, recommender systems are not always tailored to specific user needs and situations. It has been suggested that adaptive recommender systems should be modelled in terms of situations rather than knowledge structures (Adomavicius & Tuzhilin, 2005; Richthammer & Pernul, 2018; Adomavicius et al., 2011). Such a system would be capable of delivering better results to the user by taking into account contextual factors in the delivery of highly tailored information. Typically, these contextual factors include location, time, computing context, the activity of the user, or social relations (Verbert et al., 2012).

However, context can also refer to the motivational, cognitive, and emotional aspects that are inherent to the interaction between the user and the system. Most of the research on personalised recommender systems has been focused mainly on technical issues, neglecting the importance of the underlying psychological and implicit factors when exploring and analysing data (Buder & Schwind, 2012).

Thus, it is now considered relevant that for a recommender system to be effective, it should merge a variety of techniques and features in order to offer valuable support and reduce the demands imposed by information load. In this sense, systems have been developed that incorporate adaptive content presentation and adaptive navigation support (Brusilovsky, 2007). Content adaptation adjusts the presentation of the content to the user's goal, knowledge, and other information, which is stored in a model of the user to balance factors such as cognitive load, arousal, or learning style (Jin, Cardoso & Verbert, 2017).

Recommender system for domain knowledge transfer

This literature review on knowledge transfer systems reveals a multifaceted and active field where a plethora of technological approaches have been proposed and developed. It also becomes apparent that individual differences (such as motivational and emotional ones) have not received proper consideration when defining effective recommender technologies. This is mainly because of a lack of coherent principles derived from learning and cognitive sciences to guide the development of such systems.

Instead of working from a pure computer science perspective, the proposed recommender system will be grounded on cognitive theories, specifically, the Distribute Adaptive Control (DAC) theory of mind and brain (Verschure, 2012). This theory will serve a dual role in the theoretical framing and the implementation of the core functionalities of the system.

DAC considers humans themselves as adaptive systems that react and adapt to the changing demands of the environment by applying self-regulation strategies in response to intrinsic goals and motivations. The same principles play a foundational role in the implementation of more effective cognitive artificial systems.

Conceptually, this recommender system can be realised as an artificial agent whose reasoning and memory components need to extract relevant knowledge from sequences of interactions in a coordinated way. The proposed system thus emerges as the interplay of the Reactive, Adaptive and Contextual layers as defined in the DAC architecture (see Figure 1).

Figure 1: Abstract conceptualisation of the cognitive architecture of the knowledge transfer system based on the DAC framework.

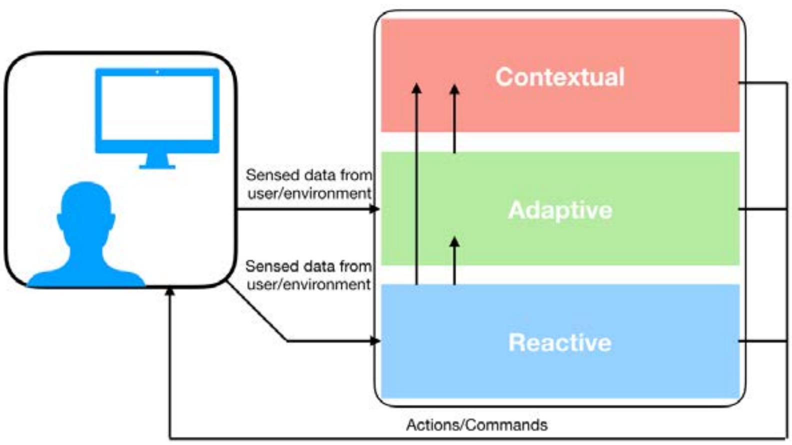

The recommender system emerges as the interplay between the three layers, which work at different timescales, with the fastest layer at the bottom and the slowest one at the top. In this architecture, the layer at the bottom (Reactive layer) provides the basic form of interaction, taking as input information from the environment and the user to facilitate the basic interaction.

Secondly, the Adaptive layer oversees adjusting the information given to the user, such as suggesting a specific piece of information or directing attention to a specific subset of information. Finally, the Contextual layer operates on longer timescales to learn from interactions from all the users, building profiles and detecting interaction strategies in order to create a knowledge base on which to optimise its behaviour to improve its capabilities in assisting the users. All in all, the system works hierarchically at different time scales, from the immediately reactive, to the medium to adapt to each user, to the long one across different interactions.

Next, one of the key aspects of a recommender system like this, which participates dynamically during an interactive task, is to decide when to provide a suggestion. There are many criteria that could be employed, depending on factors such as the specific task that the user is carrying out, how the interface has been implemented, or the number of user feedback sources available. Although we could include more complex features related to the user state (e.g., estimating stress, attention), here we present the interaction features that we have implemented in the current version of the system, to be used in an online task running on a web browser.

One of the interaction criteria is based on time. If the user has spent more than a specific amount of time without interacting with the system (by clicking somewhere), a suggestion is provided. This is done to stimulate interaction with the system, which is based on exploration to obtain information. This time threshold was fixed at 10 seconds.

Another criterion to provide a suggestion is based on the number of clicks that the user has performed without advancing in the given task. If the user has clicked a certain number of elements without getting closer to solving the task, a suggestion is provided with the goal of reorienting the user towards more relevant information.

If these criteria are not met, no suggestions are provided, as this would indicate that the user is carrying out the task successfully: with fluidity and accessing information that is relevant to solve the task at hand. This way, expert users, who already have successful strategies to accomplish the task, are not encumbered by unneeded recommendations, while novice users, who have not yet developed successful strategies, get the necessary guidance.

Another important aspect of the recommendation system is that not all the suggestions are equally revealing of the next action to take. Instead, there are different levels of recommendations, which are adapted dynamically based on the performance of the user. First, the system starts by providing general recommendations based on the content that just some users interacted with, but not most of them. As the users keep interacting with the system, if they have already received several suggestions at the current level, the recommendation level gets upgraded, and, consequently, the system recommends content of increasing popularity among the previous expert users who successfully solved the task.

To bootstrap the recommender system, some initial interaction data was needed. To achieve this, a custom synthetic data generator was developed. For a given task, the algorithm that was developed generates a random solution resembling one that an expert user would perform. This synthetic interaction data arrives at a solution by following a series of steps that are close to the optimal ones, by following some natural strategies that most users would develop after familiarising themselves enough with the system (i.e., becoming expert users).

The algorithm creates this synthetic data by working backwards from the solution of the task (i.e., starting at the end of the interaction). Then, it generates data corresponding to clicks of random pieces of information at different levels of separation from the solution. The result is a data file almost indistinguishable from the one obtained from actual interaction data.

Finally, the last step in the process is generating the recommendations from the interaction data collected or generated. To achieve this, a custom algorithm was implemented. It gathers all the existing interaction data for a given task and lists all the existing pieces of information. Then, it counts how many times each piece of information was selected by the users. The result is a data file that will later be processed by the main application to create a ranking of possible recommendations based on this information.

Use Case: Investigation knowledge graph exploration

As the initial use case of this recommender system, we chose the exploration of different knowledge graphs. These knowledge graphs represent, conceptually, one investigation. Each knowledge graph is composed of a number of interconnected nodes. The nodes represent a piece of evidence which is related to others. This is indicated by lines (edges) connecting the nodes bidirectional.

Thus, a knowledge graph here is an abstract graphical representation of all the information collected in an investigation. This modality of information presentation and exploration was designed in collaboration with Law Enforcement Agencies as part of a bigger system of state-of-the-art tools to assist officers in their investigative work by exploiting the latest digital technologies.

In this context, to validate the resulting recommender system that we implemented, we developed a simplified knowledge graph tool that does not use real investigation data, but a gamified and goal-oriented version of crime investigations. The users are asked to put themselves in the position of an investigator who must solve a series of investigations using a new visual interface. For this, they are invited to interact with the knowledge graphs, interacting with the nodes (again, each representing a piece of evidence) in search of a target node. This target node is the solution for each of the cases, representing the piece of evidence required to solve the investigation. Nodes around this target provide hints that allow participants to find out the solution.

Although this task uses the analogy of solving a case, it is important to emphasise that this is just the conceptual idea. As explained, the task is a simplified version, being closer to a game than an actual job of an officer investigating a real case. The way to solve each of the tasks is based on solving a series of logic puzzles, as explained below.

Figure 2 depicts the user interface that we implemented to present the task. The knowledge graph itself occupies the central part of the screen. Users can interact with the graph by clicking on the different nodes to obtain information about them (name and possible relationship to the target node). The name of the node also appears when hovering the mouse cursor over it. It is also possible to move the nodes by clicking and dragging, which might be helpful to get more clarity on the connection to other nodes, although this is never required. Users can also displace the graph by clicking and dragging on an empty space, as well as zooming in and out by using the controls provided in the top-right corner, although these actions are not required either. Finally, in the top centre, the category of the target node is displayed.

Figure 2: The user interface of the knowledge graph exploration task.

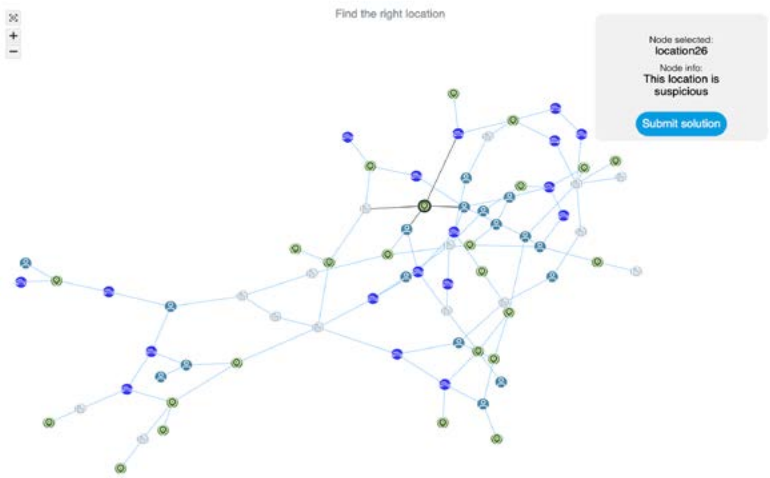

Interaction controls for the displacement and zoom of the graph are located in the top right corner (from the bottom up: zoom out, zoom in, restore view). On the top centre, the category of the current target is indicated. The panel on the top left provides information about the node that is currently selected, which appears with a black outline and black connections in the graph. This panel also has the button to submit the solution, corresponding to the node currently selected.

We decided to use four different categories of nodes to provide enough diversity without being too distracting or overloading. These four categories are: person, vehicle, text, and location. Each category is differentiated from the others by using a different iconic figure and colour (see Figure 2 and Figure 3).

As mentioned before, the nodes surrounding the target provide relevant information that is needed to the solving the case. Depending on their closeness and relevance, four levels are established and displayed in the node information:

- *"This [category name] is suspicious"*: This appears for the target node and for all nodes of the same category that are within three degrees of separation from it.
- *"This [category name] is directly related to the target"*: This appears for nodes that are directly related to the target (first-degree connection), of a different category from it.
- *"This [category name] is indirectly related to the target"*: This appears for nodes that are indirectly related to the target (second-degree connection, which is, connected through exactly one node in between), of a different category from it.
- *"Unclear"*: This appears for all other nodes not covered in the previous three categories, i.e., all nodes that are too distant and unrelated from the target.

Using this information provided by the different nodes close to the target, the solution is implied. In each knowledge graph, there is only one possible solution, and the information, when enough nodes are explored, points unambiguously to it. Users must integrate this information in a logical manner. It is a matter of logically inferring the solution by integrating some simple relationship data.

The complexity of the task is modulated by the size of the knowledge graph, determined by the number of nodes and connections. The higher the number of nodes and connections, the higher the difficulty, as the visual complexity increases and there are more nodes to explore. We created three difficulty levels according to this: 50, 100, and 200 nodes and connections. Two graphs of each difficulty level are presented, in increasing difficulty, for a total of six cases for each participant.

As indicated, one of the key aspects of this system is the presentation of suggestions to the users. These suggestions are provided in the form of recommended nodes based on the actions of other users. When a node is suggested, it gets selected with a thicker light-blue outline. Its connections to other nodes also appear in the same colour (see Figure 3). When a node is suggested, a panel appears on the bottom of the screen, alerting users of this fact and thus ensuring that they notice the suggestion. This message stays on the screen for 3 seconds.

Figure 3: Example of a suggested node.

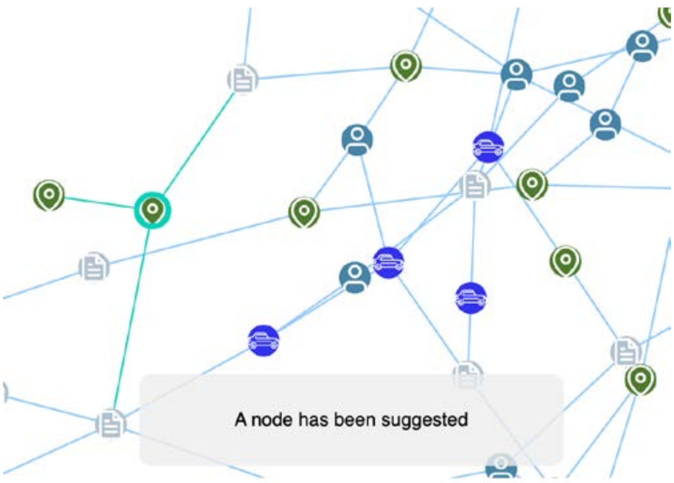

It appears with a thicker light blue outline, as well as its direct connections to other nodes. A temporary panel appears on the bottom, alerting the user that a node has been suggested.

Discussion and conclusion

Here, we have presented a novel assistive system capable of learning from interactions with users in order to provide relevant suggestions to other users, in the context of investigative work performed by law enforcement officials, with the aim of facilitating the learning of the use of a new system for the exploration of investigation information.

As explained, the assistive system developed, based on principles from recommender systems and cognitive science, is used in the exploration of a knowledge graph composed of different nodes and connections representing pieces of information collected during an investigation. This knowledge graph implemented here is analogous to one that could be used in a real investigation but customised to provide a goal-oriented task to users: exploring the graph to obtain information necessary to find a target node.

From a technical standpoint, in order to develop such a system, a multitude of components were implemented, as described, including a generator of knowledge graphs, the recommender system itself, a generator of synthetic interaction data used to bootstrap the recommender system, and a generator of recommendations for a given graph based on interaction data (either collected from actual users or generated synthetically).

For the experimental validation of this assistive system, two groups of participants are proposed: an experimental one, which receives suggestions as needed (based on different criteria), and a control one, which does not receive any assistance from the system. Thus, the expectation would be that participants who receive automated recommendations from the system perform better, both in objective metrics (such as performance), implicit features (such as mouse movements) and self-reports (in the questionnaire provided after the main tasks). In addition to an initial validation with civilian participants, a validation should be carried out in collaboration with LEAs and, especially, with the end users of the system.

As mentioned earlier, the system presented in this article works as a web application that runs on a web browser for it to be available online. Due to this, and as an initial implementation of the assistive system, the recommendations were triggered based on different interaction features like the time elapsed, the number of clicks, attempts to solve the task, etc., which were the most viable and appropriate sources, while still providing the necessary information for the knowledge transfer system that was used.

However, more sophisticated methods could be implemented to trigger these recommendations based on the internal states of the users, as inferred in real-time based on various signals. For example, suggestions could be provided based on the estimated cognitive load of the user using pupil dilation signals captured by an eye-tracker, or stress levels could be estimated from physiological signals such as the variation in heart rate using the appropriate sensor.

In conclusion, here, we have proposed a novel assistive system for transferring domain knowledge to novice officers, exploiting modern technologies to facilitate the training of officers in the use of new digital tools to be used in the field in the course of their work. With this work, the authors would like to highlight how state-of-the-art technologies can be applied by forward-thinking LEAs, with the aim of improving the training and education of current and future law enforcement officials in and for the Digital Age.

Acknowledgements

This research was conducted under the framework of the ANITA project, supported by the European Commission (H2020-787061). It was also supported by the European Commission project EIC 101058240 PHRASE. The authors would like to thank Riccardo Zucca for his contributions to this work, among other collaborators from the ANITA project.

References

- Adomavicius, G., Mobasher, B., Ricci, F. & Tuzhilin, A. (2011) What a Recommender System Knows About Contextual Factors. *AI Magazine* 32(3), 67–80.
- Adomavicius, G. & Tuzhilin, A. (2005) Towards the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge & Data Engineering*. 17(6), 734–749.
- Asabere, N.Y. (2012) A survey of personalized television and video recommender systems and techniques. *Information and Communication Technology Research*. 2 (7), 602–608.
- Brusilovsky, P. (2007) Adaptive navigation support. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. 4321 LNCS, 263–290. doi:10.1007/978-3-540-72079-9_8.
- Buder, J. & Schwind, C. (2012) Learning with personalized recommender systems: A psychological view. *Computers in Human Behavior*. 28 (1), 207–216. doi:10.1016/j.chb.2011.09.002.
- Burke, R. (2002) Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction*. 12 (4), 331–370.
- Ge, M. & Persia, F. (2017) A survey of multimedia recommender systems: Challenges and opportunities. *International Journal of Semantic Computing*. 11 (03), 411–428.
- Goldberg, D., Nichols, D., Oki, B.M. & Terry, D. (1992) Using collaborative filtering to weave an information tapestry. *Communications of the ACM*. 35 (12), 61–71.
- Jin, Y., Cardoso, B. & Verbert, K. (2017) How do different levels of user control affect cognitive load and acceptance of recommendations? *CEUR Workshop Proceedings*. 1884, 35–42.
- Park, D.H., Kim, H.K., Choi, I.Y. & Kim, J.K. (2012) A literature review and classification of recommender systems research. *Expert Systems with Applications*. 39 (11), 10059–10072. doi:10.1016/j.eswa.2012.02.038.
- Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. & Ricci, F. (2011) *Recommender Systems Handbook*. doi:10.1007/978-0-387-85820-3.
- Richthammer, C. & Pernul, G. (2018) Situation awareness for recommender systems. *Electronic Commerce Research*. (0123456789). doi:10.1007/s10660-018-9321-z.
- Song, Y., Dixon, S. & Pearce, M. (2012) A survey of music recommendation systems and future perspectives. In: *9th International Symposium on Computer Music Modeling and Retrieval*. 2012 p.
- Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachsler, H., Bosnic, I. & Duval, E. (2012) Context-aware recommender systems for learning: a survey and future challenges. *IEEE Transactions on Learning Technologies*. 5 (4), 318–335.
- Verschure, P.F.M.J. (2012) Distributed Adaptive Control: A theory of the Mind, Brain, Body Nexus. *Biologically Inspired Cognitive Architectures*. 1, 55–72. doi:10.1016/j.bica.2012.04.005.
- Villegas, N.M., Sánchez, C., Díaz-Cely, J. & Tamura, G. (2018) Characterizing context-aware recommender systems: A systematic literature review. *Knowledge-Based Systems*. 140, 173–200. doi:10.1016/j.knosys.2017.11.003.