CLARK: Building Conversational Intelligence for Knowledge Management in the Space Domain

Marcel Henriquez[1], Paulius Danėnas[2], Darius Dilijonas[2], Jose Ontiveros[3] and Andrew Herd[4]
[1]Red Data, Almere, The Netherlands
[2]Intellerts, Kaunas, Lithuania
[3]Immedia IT Software, Madrid, Spain
[4]ESA - European Space Agency, Noordwijk, The Netherlands

marcel.henriquez@reddata.nl
p.danenas@intellerts.com
d.dilijonas@intellerts.com
jmontiveros@immediait.com
andrew.herd@esa.int

Abstract: This paper presents the case study of the CLARK (*Captured Lessons applied for the Reuse of Knowledge*) project which is an evolution of the Lessons Learned (LL) portal at the European Space Agency (ESA). The SCARLET research activity has been a key for-runner to this project on an academic research level. It has provided us with important insights into the architecture of the system design. A knowledge graph (KG) has been developed to (re)search (for) knowledge from different angles through the established relationships by combining and structuring data from many sources. Additionally, a chatbot is trained to provide a conversational intelligence experience on top of the data provided by the different Knowledge Management (KM) activities at ESA. Hence, CLARK project enables to retrieve, search, or provide knowledge for the employees across ESA through different interfaces connected to the core systems. In this paper, we also report some findings from the initial testing which indicate very promising results in terms of user engagement and usability of CLARK system and the overall approach. The paper also discusses current and future challenges.

Keywords: Conversational intelligence, Chatbot, Knowledge graphs, Natural language querying, Space, Knowledge re-use

1. Introduction

This paper is a case study of the European Space Agency development project to create a major evolution of the existing Lessons Learned capabilities currently in use at the Agency through their existing SharePoint-based webtool. Through considering the application of knowledge graphs and chatbot technologies to achieve this major "next step" in better addressing user needs, it has been possible to realise a radical shift in capabilities offered to users. This adoption of multiple novel technologies has in turn provided a higher level of relevance to user needs and hence benefit regarding user themselves. As a result it is suggested that these capabilities can now be more broadly applied and be available as Knowledge Management activities start to yield real solutions to everyday problems in a way that users want these to be offered and presented. In placing the user front and centre to the proposed solution, CLARK - Captured Lessons Applied for the Reuse of Knowledge, a diverse range of user interface options are being developed to provide knowledge through image and text-based interfaces, including providing interactions between CLARK and the user in a human like, conversational way, through the conversational agent (commonly known as a Chat Bot).

This project is a follow up on the research project SCARLET - Space Conversational Agent for Retrieving

Lessons learned and Expert Training - performed by the team of Strathclyde University led by Massimiliano Vasile (Mihaylov et al. 2022). The research has been done as a joint undertaking with the European Space Agency by Strathclyde University from a perspective "which technologies can be used to create a conversational agent" and "how to prove this technology imparts a higher quality of lesson capture and dissemination". A minimum viable product has also been built by Strathclyde using static data and tested successfully. The CLARK project uses the experience from the Strathclyde University SCARLET research team to build a production system that collects and processes data continuously, robust enough for ESA engineers.

As to be expected from a large organization like the European Space Agency, particularly one that works on very complex problems and technological solutions, there is a lot of new techniques being generated and knowledge being acquired. Project teams working in the many different Directorates (of which ESA has eleven) focussing on different aspects of space missions are all working hard to get the (large) space programs aims realised (together with the space industry), operationalized and maintained.

For nearly two decades, ESA has had procedures in place to gather Lessons, captured and learned at the Space program level. For the last six years a centralized Lessons Learned (web-based) portal has been implemented that further helps ESA implement its process and to capture the important knowledge in a centralized system and in a common way. The lessons learned portal is the flagship of the knowledge management activities undertaken by ESA to safeguard its internal knowledge. The next step, and the main subject of this paper, is the implementation of CLARK on top of the centralized Lessons Learned repository. The main objective of CLARK is to directly address the identified use cases and user needs in order to allow the ESA workforce to access knowledge through different interfaces and through different approaches.

The technical setup of the CLARK system was configured in such a way to allow it to suit different types of situations users are looking for Knowledge. The simple search, for example, for a user who knows exactly what they are looking for and can find it with just a few keywords. The

Knowledge Graph is provided for the user who is visually inclined and knows how to navigate a number of data streams through their interrelationships (as presented in a knowledge graph). In turn, the chatbot being provided for the users that prefer to "interact" with the system, entering or extracting knowledge in an assisted manner, this particularly applies to the case where the user has an idea what he or she is looking for, but not how to look for it and wants a more "open" search environment.

For the end result of the CLARK project four components had to be developed.

- A data platform capability where data from different sources is held centrally and continuously updated from each source. The sources provide primarily structured data (content). At this point there are 9 sources that provide data to the CLARK data platform. The Lessons Learned portal being the primary source of configured information.
- A Knowledge Graph platform. In the Knowledge Graph platform different relations between the data (given in the graph as unique named nodes) are created and displayed to the user. Examples of a node and their relationships is a person and lessons, technology and a project and so on. This gives the user a possibility to browse through large, interconnected datasets through a completely visual interface where the node relations determine the outcome of the term being searched.
- A chatbot platform. The chatbot interface that is fed by the data platform and the Knowledge Graph platform, provides a conversational assistant to the user for both knowledge capture and retrieval. The chatbot also has a learning loop which allows it to learn from the interactions with ESA workforce. Additionally, the chatbot also has the capability to improve the quality of the capture a new lesson from a user, such as by identifying terms and acronyms not understood (i.e., not available) in the CLARK system.
- User Interfaces. For the user to use these new capabilities, two interfaces (offering the same capabilities) are built, a web interface (for desktop use) and a mobile app interface (for mobile device use). The key here is for the interfaces to be as simple and intuitive as possible, so the user does not have to focus on understanding what the interface does, but rather focussing on which capabilities she / he wants to start utilizing . In addition the user interfaces allow for a simple customisation by the user, to have the interface provide lessons matching the user interests, zas these lessons become available in the lessons learned portal.

2. Related Work

The fields of knowledge management and information processing historically have attracted a lot of attention from the scientific community; therefore, complete discussion of all the work which can be attributed to the success of this work is out of scope. However, there are several subfields which are considered more relevant in this context. Knowledge graph related research is one of these fields; in fact, knowledge graphs (also referred to as semantic networks) have been widely adopted in various domains. National Aeronautics and Space Administration (NASA) has also adopted KG technology for its Science Mission Directorate domains (Timmer et al. 2023); moreover, they have previously adopted it for scientific knowledge storage and organization (Zhang et al. 2020).

One of the most important recent developments in natural language processing (NLP) was the introduction of BERT transformer architecture (Devlin et al. 2018) and its successor RoBERTa (Liu et al. 2019). Transfer learning paradigm enables simplified reuse or adoption of previously pretrained models for different tasks or domains. Hence, these models are often applicable as parts of modern NLP systems, providing contextual embeddings with more relevance for particular domains. Several models exist for space related domains, namely SciBERT

(Beltagy et al. 2019), AstroBERT (Grèzes et al., 2021), SpaceBERT, SpaceRoBERTa and SpaceSciBERT (Berquand et al. 2021). In the context of CLARK project, they might be helpful to improve detection of named entities or concepts, as well as improve support for domain specific language or terminology.

Question answering (QA) is another influential topic; the most recent and well-known applications emerge from large generative language models (LLM), such as GPT-4 (OpenAI 2023), LaMDA (Thoppilan et al. 2022), LlaMA 2 (Touvron et al. 2023), combined with careful prompt engineering. This approach was first verified by creators of GPT-3 model showing that providing several examples can vastly improve generated output quality (Brown et al. 2020). However, although widely developed, such systems still struggle to provide precise answers. To overcome this issue, techniques like prompt chains or direct integration with external data sources or knowledge bases directly is generally considered. (Diefenbach et al. 2017) provide an extensive survey of traditional approaches for QA over knowledge bases; they show that a huge number of issues in question analysis, query construction, disambiguation and general natural language processing could be required to be solved to achieve high quality results.

Finally, natural language querying enables formulating queries in natural language which is more appropriate for human users. This is one of the core problems addressed by CLARK (and particularly chatbot) development. (Affolter et al. 2019) identified four types of NL database interfaces, particularly, keyword-based, pattern-based, parsing-based, grammar-based. Moreover, one can observe emerging trends of neural translation-based systems based on deep-learning approaches. Sequence-to-sequence (seq2seq) architecture, which was introduced for machine translation tasks, is also applicable for SQL-like query language generation (Zhong et al. 2017). (Bazaga et al. 2021) introduce Polyglotter translation approach based on Seq2Seq architecture which can be applied to generate queries for almost any database due to graph abstraction which is used to represent query candidates (in their work, MySQL and Neo4J are used as case studies).

3. The CLARK Approach

3.1 It all Starts with the Data

The primary target of the whole system is to let people learn from the experience(s) of their colleagues within ESA. For this target to be met, and the technology to be enabled, different types of data are required. Such data enables the user to perform contextual search for answers to his questions together with supporting objects or entities. In the context of CLARK, two categories of data are considered. The first category are primarily the lessons from the ESA Lessons Learned portal, associated with relevant documents. These lesson objects also contain information based on experience from all types of people (roles and functions) within the large projects. The second category of data is relevant to provide context in search, such as project list, technology tree and product tree. These all help the system to categorize the results and the user to put the input in perspective.

There are nine different sets of data sources in total which have been imported into the data platform. Some of them provide unstructured (e.g., text content) or partially structured data, hence they are processed using more advanced search technologies (such as Apache SOLR). Additionally, it is enriched with additional metadata to create cohesion between different data sources. As each source has its own schema for enrichment, this step can be performed automatically, resulting in structure within the unstructured data. Additionally, to enable change management (different metadata, change of sources, etc.) data is initially stored in data lake.

3.2 Knowledge Graphs

Knowledge graph (KG) technology enables organizing data as a set of entities and relations between them, thus forming a semantic network, which can be further traversed to mine more complex relationships (including complete hierarchies represented by *isPartOf* or *isParentOf* relations), filter them or visualize in a hierarchical way which is very convenient for such type of data. This inspired us to build a KG-based system based on two core components: a graph database (Neo4J) which stores all the relations and nodes, and a graph visualization component to render to representation of these nodes and their relationships. The core data comes from the data platform (described in Section 3.1), where it is further refined, preprocessed and a set of nodes are formed, together with their interrelations.

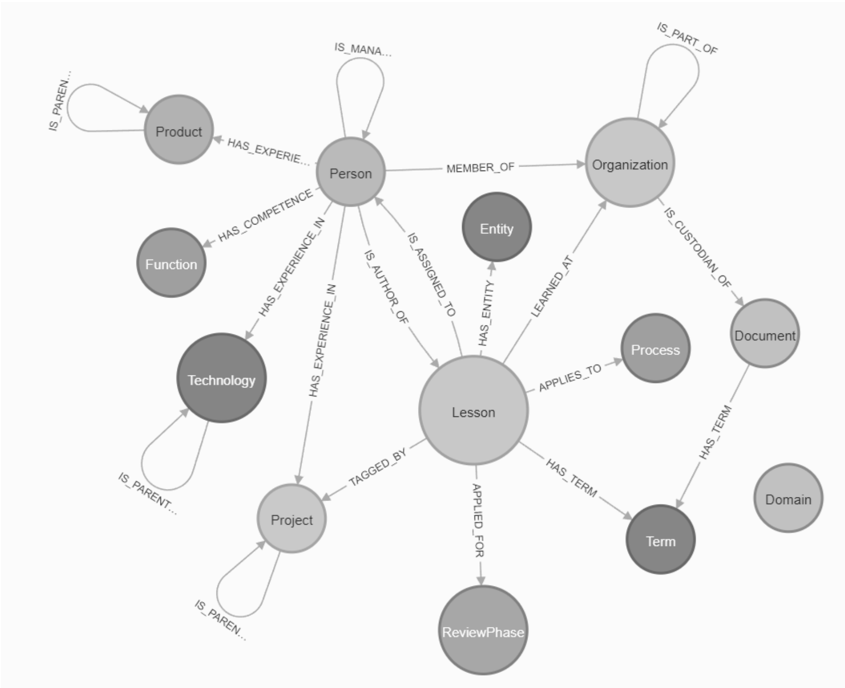

Figure 1: Knowledge Graph Structure

The KG system provides an interactive search tool for the users. This gives the user a possibility to browse through large datasets in a completely visual interface where the relations determine the outcome. Additionally, knowledge graph enables advanced search capabilities which can be employed by the chatbot to generate consistent and domain-focused answers to the system users, such as answering questions *Who is the author of this lesson?* or *List me lessons associated with project Sentinel*.

3.3 Conversational Intelligence with CLARK Chatbot

Arguably, the main and the most complex objective of CLARK project was to create a chatbot that can provide the user with the knowledge he or she is looking for through a conversation. Moreover, this is a domain-specific chatbot which should be capable to provide precise answers as available in the underlying KG database, which makes it more complex compared to recent trend of generative deep learning based chatbot technology (such as GPT) that tend to sometimes provide incorrect answers or even perform so-called “hallucinations”. Therefore, it also addresses several specific challenges as discussed further.

To communicate with the user and successfully assist him, CLARK chatbot implements a certain number of scenarios. They define how the user is expected to interact with the system. Currently, the following scenarios are implemented:

- List objects which can be searched: "Show me employees available in the database”, “Provide me the list of available technologies”
- Find number of objects in the database: "How many lessons learned are available in the database?”
- Find one or more objects by related term: "Show me information about Peter Peterson”, “Give me information related to Sentinel“
- Find objects that are related to other objects by some relation: “Show me people who worked in project Lunar Lander“, “Show me documents containing term flow control”
- Find property of some object: "Find me upper-level products for product Sensors", “Who is the supervisor of Peter Peterson?”
- Help on lessons learned similarity: “What is the similarity between lessons LL-0641 and LL-0650?”
- Submit new lessons learned: “I want to submit new lesson learned”
- Create specific recommendations to the user: “Could you give me some recommendations?”
- Add keywords and narrative to the lesson learned through the chatbot interface in the portal
- Provide help, sample queries and assistance to the user about the use of the chatbot: “How can I discover objects associated with a specific entity?”, “Show me how to search for properties of a specific object”

Additionally, it can provide recommendations for the user:

- Based on common terms from lessons in the previous user learning experience and lessons in the database
- Recommendations related to the working domain of the user
- Based by work of colleagues in the department of the user
- Based on similarity to a particular lesson

Recommendations are created using a specialized interactive menu-like scenario: the user is asked to select a recommendation type and, if required, provide additional input. The chatbot outputs several lessons learned according to the selected criteria. This scenario is illustrated in Figure 2.

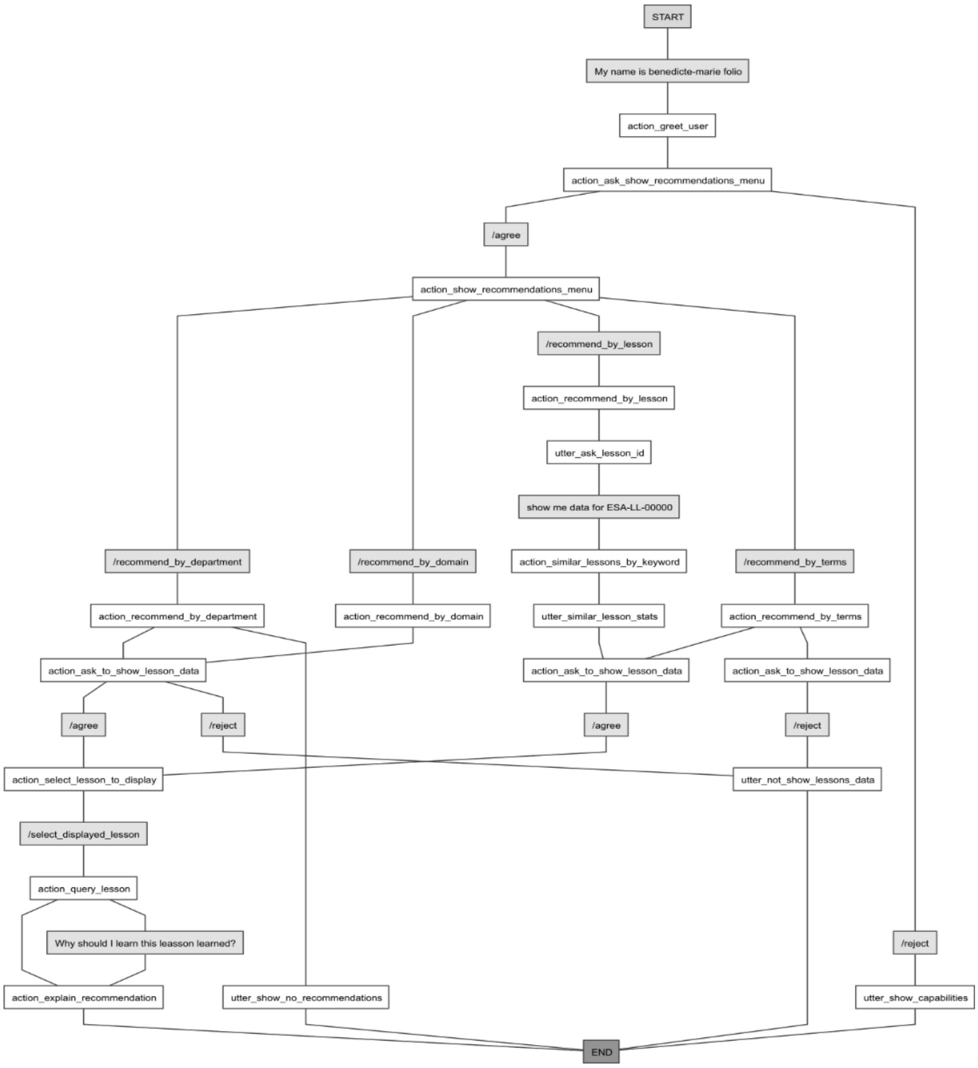

Figure 2: Recommendations Scenario

The final implementation is based on RASA framework which provides its own deep learning-based technology (DIET classifier) for training chatbot models and mapping them with specific knowledge-based actions. Moreover, it is also capable of integrating recent developments in language models, such as transformer models (BERT, RoBERTa) or generative GPT-like models, to provide their contextual embeddings as inputs. However, our experimental testing indicated that application of BERT or GPT-2 did not result in significant performance increase, hence, final model is based on RASA default DIET classifier with advanced processing pipeline setup.

3.4 The Complete View

For the user to use these new capabilities, two interfaces are built. From the usability perspective, interfaces should be as intuitive as possible, so the user does not have to focus on understanding what the interface does, but rather which capabilities are now available and can be used immediately.

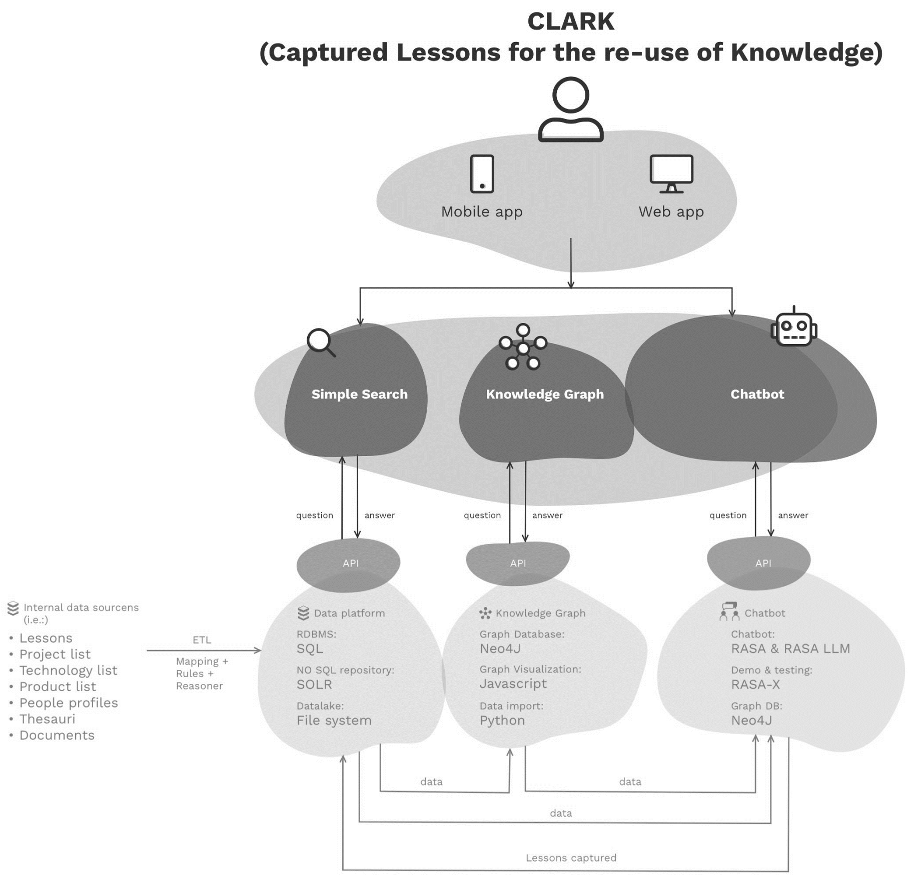

Figure 3: CLARK Architecture

Creating an elaborate system like CLARK does not mean the intended users will automatically start using the system. For the adoption of the CLARK system a webapp and mobile app (both compliant with ESA requirements) have been built. The webapp will by incorporated in the Lessons Learned portal and a mobile app will work only on ESA issued phones to comply to strict ESA security requirements. Both interfaces offer the user the chance to search for knowledge through the three possibilities.

4. Experimental Evaluation

To test the initial system, we ran an internal experimental evaluation at ESA. Eight users were selected to perform initial testing of the chatbot. They were provided with the documentation, including implemented scenarios, sample intents (questions for the chatbot). RASA-X, a visual frontend for conversation-driven development provided by RASA NLU framework, was used for interactive testing. We note that RASA-X does not represent the final user experience and should be viewed only as a temporary means for testing the overall approach; this is different actual usability testing performed for the CLARK webapp and mobile app.

The users were asked to provide their feedback in a free form. In this summary, we omitted feedback which was related to technical limitations which are to be solved after final integration of all components, minor technical inconsistencies (missing text, errors in text display, etc.) or out of scope functionality (mostly user interface related). While we consider them as less relevant at the initial stage, we admit that it could affect their first impression**.** Nevertheless, all these results proved to be beneficial for both the teams of chatbot and webapp/mobile development, as they clearly indicated the expectations from the end users and helped to solve some of them at initial stages of development.

Below we provide a summarized evaluation obtained from the feedback of the testers.

- *User feedback*. The main factors that users observed were related to the inconsistency of the output messages.

- *Conversation*. The overall chat experience seemed to be too formal or constrained for some of the users ("*the conversation had the feeling of being a set of commands input by the user..*"). This is not surprising since this is a chatbot focused on information retrieval, rather than communication; yet it provides hints for the future iterations. However, some of the users emphasized the intuitiveness but emphasized the limited feedback, especially when the chatbot failed to recognize intent and to continue scenario (*"unclear where you were at any on time and what process you were in..", "abrupt endings…"*). Clearly, this not only helped us to solve some unintended issues but provided clearer guidelines to handle such situations as fallback scenarios in the future.
- *Usability*. The feedback in this group is mostly related to the scenario. The key factors that users lacked were loss in conversation, absence of integrated chatbot help and user guidance, confirmation whether the chatbot understood their intents correctly, and inability to reset the chat.
- *Visualization*. This input is more relevant to the mobile app group, as it provided some expectancies for the user interface (bold field names, selection of black/white background, presence of reset section button), content layout or general output of the results.

Some important findings were observed after analyzing the obtained results:

- The feedback from the users who, according to their feedback, seemed to be familiar with the LL portal was more positive. This is not surprising as familiarity with the data present in the portal helped them to understand chatbot capabilities and formulate their requests to the chatbot more clearly.
- The actual expectations of the users were quite different – some of them were satisfied with the overall experience (although they had some comments) while the others seemed to have some skepticism about it. Such satisfiability distribution is not surprising at the initial stages, as it is expected to improve chatbot performance and usability in later iterations.
- From the system engineering viewpoint, the variability of user feedback clearly identified to need for multiple iterations of requirements re-elicitation and refinement, as well as application of agile development while developing conversational intelligence-based systems.

5. Current Challenges and Future Work

The initial testing has already provided us with valuable feedback from the users. However, there are multiple challenges or guidelines to be addressed in the future developments:

- *Integration of large language models (LLM).* LLMs, as the most recent developments in NLP, provide numerous application capabilities in solving various tasks. Despite their widely recognized issues, they can provide quite satisfactory solutions to various problems. In CLARK context, they could be applied as enhancements for multiple tasks, such as correcting mistypes or grammatical errors, improving general text quality, paraphrasing, summarizing text, extracting keywords or named entities, answering answers directly from LL text, and even generating structured representations such as visual diagrams or mind maps from text.
- *Longer chatbot memory context window*. This would enable more natural communication between user and chatbot. For instance, the chatbot could be able to extract various facts and store them in memory, while making use of them if asked. However, it might impose challenges related to advanced natural language processing and dialog management, such as co-reference resolution ("*it*" – which fact would it refer to?), word-sense disambiguation among others.
- *Better and more personalized recommendations for the user*. The user could get more personalized LL recommendations based on his preferences. This might be enabled by setting these preferences directly or by using integrated user feedback mechanisms (such as asking if the user liked this lesson, etc.). This requires the ability to set, store and manage these preferences. Moreover, lesson similarity and matching techniques could be improved as well, using novel semantic matching techniques.
- *(Semi) automated knowledge graph construction*. The knowledge graph could be automatically enriched with entities or relations extracted for LL text; this is also addressed in recent developments in the field of knowledge graph completion.
- *Dynamic alignment to changes in the underlying knowledge graph*. The chatbot should be capable of detecting new node types, entities or relationships. We note that the current implementation is already KG metadata-aware, yet, it still has some limitations, mostly due to restrictions or internal elements which must be excluded during search.

- *More advanced fallback implementation scenarios*. The proper implementation of fallback scenario, albeit seeming to be very simplistic, might not be trivial. Practical realizations may require handling different situations, such as: the text entered by the user, is garbage and should be skipped from further processing; the chatbot does not recognize any relevant entities or intents; handling and incomplete list various undefined exceptions.

6. Conclusion

The complete CLARK system is now in beta and will rapidly advance to production level. There are still a number of hurdles that need to be cleared as part of the planned integrated tests later in the year. An ESA team of testers with different backgrounds have used all three capabilities in their “stand-alone” form with positive and constructive feedback. The next step is to deliver the knowledge graph and Chatbot capabilities integrated in a single user interface and connect all platforms with the final “live” data (production).

The test results up until now are promising which are in line with the initial outcomes of the Strathclyde research that this technological solution (combining knowledge graph and chatbot in a single interface) is a unique and novel means to address user needs for quality knowledge capture and retrieval.. Some rapid progress has been made of the past months to realise this new capability and it should be acknowledged that none of the CLARK components existed 15 months ago when the project first started. Having said that, directly building upon the work of Strathclyde, and with the combination of the skills and capabilities of the knowledge graph and chatbots teams, the development project has been able to make great steps and a production version will be delivered within the next three months and deployed on the ESA corporate infrastructure as a standard capability.

The CLARK development is already creating a wave of interest within the Agency which is undergoing an internal transformation, and it is hoped that others (and other technology developments) can use this new exciting technology to achieve other significant advancements in lessons learned and any other application that might benefit from this capability..

Finally, it can be stated that CLARK basis is robust enough for further development, by extending it with additional entities (or nodes in KG), relations between them, concerning the chatbot – implementing new actions and scenarios. This can be in the form of additional sources and data, creating more context and scenarios. Moreover, as stated in the future work, there are numerous ways to improve it, making it even more accessible and personalized for the end user.

References

Affolter, Katrin, Stockinger, Kurt, and Bernstein, Abraham (2019), "A comparative survey of recent natural language interfaces for databases", The VLDB Journal, Vol. 28, No. 5, pp. 793-819.

Bazaga, Adrián, Gunwant, Nupur, and Gos Micklem (2021) "Translating synthetic natural language to database queries with a polyglot deep learning framework." Scientific Reports, Vol. 11, no. 1, pp. 1-11.

Beltagy, Iz., Lo, Kyle, and Cohan, Arman (2019) "SciBERT: A Pretrained Language Model for Scientific Text", Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP), pp. 3615-3620.

Berquand Audrey, Darm Paul, and Riccardi, Annalisa (2021). "SpaceTransformers: Language Modeling for Space Systems", IEEE Access, Vol. 9, pp.133111-133122.

Brown, Tom B., Mann, Benjamin, Ryder, Nick, Subbiah, Melanie, Kaplan, Jared, Dhariwal, Prafulla, Neelakantan, Arvind, Shyam, Pranav, Sastry, Girish, Askell, Amanda, Agarwal, Sandhini, Herbert-Voss, Ariel, Krueger, Gretchen, Henighan, Tom, Child, Rewon, Ramesh, Aditya, Ziegler, Daniel M., Wu, Jeffrey, Winter, Clemens, Hesse, Christopher, Chen, Mark, Sigler, Eric, Litwin, Mateusz, Gray, Scott, Chess, Benjamin, Clark, Jack, Berner, Christopher, McCandlish, Sam, Radford, Alec, Sutskever, Ilya, Amodei, Dario (2020) "Language models are few-shot learners." Advances in neural information processing systems Vol. 33, pp. 1877-1901.

Devlin, Jacob, Chang, Ming-Wei, Lee, Kenton, and Toutanova, Kristina (2019), "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding", In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers), pp. 4171-4186

Diefenbach, Dennis, Lopez, Vanessa, Singh, Kamal and Maret, Pierre (2017) "Core techniques of question answering systems over knowledge bases: a survey", Knowledge and Information Systems, Vol. 55, No. 3, pp. 529-569.

Grèzes, Félix, Blanco-Cuaresma, Sergi, Accomazzi, Alberto, Kurtz, Michael J., Shapurian, Golnaz, Henneken, Edwin A., Grant, Carolyn S., Thompson, Donna M., Chyla, Roman, McDonald, Stephen, Hostetler, Timothy W., Templeton, Matthew R., Lockhart, Kelly E., Martinovic, Nemanja, Chen, Shinyi, Tanner, Chris, and Protopapas, Pavlos (2021) "Building astroBERT, a language model for Astronomy & Astrophysics", CoRR, abs/2112.00590.

Liu, Yinhan, Ott, Myle, Goyal, Naman, Du, Jingfei, Joshi, Mandar, Chen, Danqi, Levy, Omer, Lewis, Mike, Zettlemoyer, Luke, and Stoyanov, Veselin (2019) "RoBERTa: A Robustly Optimized BERT Pretraining Approach", CoRR, abs/1907.11692.

Mihaylov, Dimitar, Broughton-Stuart, Graye, Vasile, Massimiliano, Moshfeghi, Yashar and Herd, Andrew (2022) "A Space Conversational Agent for Retrieving Lessons-learned and Expert Training". 73rd International Astronautical Congress, Paris, France, pp. 18-22.

OpenAI (2023). "GPT-4 Technical Report." ArXiv abs/2303.08774.

Thoppilan, Romal, Freitas, Daniel De, Hall, Jamie, Shazeer, Noam, Kulshreshtha, Apoorv, Cheng, Heng-Tze, Jin, Alicia, Bos, Taylor, Baker, Leslie, Du, Yu, Li, YaGuang, Lee, Hongrae, Zheng, Huaixiu Steven, Ghafouri, Amin, Menegali, Marcelo, Huang, Yanping, Krikun, Maxim, Lepikhin, Dmitry, Qin, James, Chen, Dehao, Xu, Yuanzhong, Chen, Zhifeng, Roberts, Adam, Bosma, Maarten, Zhou, Yanqi, Chang, Chung-Ching, Krivokon, Igor, Rusch, Will, Pickett, Marc, Meier-Hellstern, Kathleen S., Morris, Meredith Ringel, Doshi, Tulsee, Santos, Renelito Delos, Duke, Toju, Soraker, Johnny, Zevenbergen, Ben, Prabhakaran, Vinodkumar, Diaz, Mark, Hutchinson, Ben, Olson, Kristen, Molina, Alejandra, Hoffman-John, Erin, Lee, Josh, Aroyo, Lora, Rajakumar, Ravi, Butryna, Alena, Lamm, Matthew, Kuzmina, Viktoriya, Fenton, Joe, Cohen, Aaron, Bernstein, Rachel, Kurzweil, Ray, y Arcas, Blaise Agüera, Cui, Claire, Croak, Marian, Chi, Ed H., Le, Quoc (2022) "LaMDA: Language Models for Dialog Applications." ArXiv abs/2201.08239.

Timmer, Roelien C., Mark, Megan, Khoo, Fech Scen, Scoczynski Marcella, Martins, Ribeiro, Berea, Anamaria, Renard, Gregory, and Bugbee, Kaylin M. (2023) "NASA Science Mission Directorate Knowledge Graph Discovery", In Companion Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023 (pp. 795-799). ACM.

Touvron, Hugo, Lavril, Thibaut, Izacard, Gautier, Martinet, Xavier, Lachaux, Marie-Anne, Lacroix, Timothée, Rozière, Baptiste, Goyal, Naman, Hambro, Eric, Azhar, Faisal, Rodriguez, Aurélien, Joulin, Armand, Grave, Edouard and Lample, Guillaume (2023). "LLaMA: Open and Efficient Foundation Language Models." ArXiv abs/2302.13971.

Zhang, Jia, Pourreza, Maryam, Ramachandran, Rahul, Lee, Tsengdar J., Gatlin, Patrick, Maskey, Manil, and Weigel, Amanda Marie (2018), "Facilitating Data-Centric Recommendation in Knowledge Graph". In 4th IEEE International Conference on Collaboration and Internet Computing, CIC 2018, Philadelphia, PA, USA, October 18-20, pp. 207-216

Zhong, Victor; Xiong, Caiming and Socher, Richard (2017) "Seq2sql: Generating structured queries from natural language using reinforcement learning", arXiv preprint arXiv:1709.00103.