

OCAI: the Operations CompAnIon to support decision making of flight control teams

Gabriele De Canio*, James Eggleston, Artur M. Palowski, Micha Schmidt, Peter Collins, Bruno Sousa, Daniel Mesples, Rui Santos, Evridiki Ntagiou

European Space Operations Centre, European Space Agency, Robert-Bosch-Str. 5, 64293 Darmstadt, Germany

* Corresponding author, gabriele.decanio@esa.int

Abstract

With the vision to provide a multi-mission infrastructure allowing Flight Control Teams (FCTs) to leverage standard systems, the European Space Operations Centre of the European Space Agency is committed to system harmonization. The existing systems landscape used by FCTs for anomaly investigation is somewhat fragmented. As a result, FCTs must operate multiple systems to retrieve, combine and analyse data in their daily decision-making processes. This creates inefficiencies and challenges for newcomers during onboarding and initial training. In the frame of the Artificial Intelligence for Automation (A²I) Roadmap, recently developed by ESOC in close collaboration with industry, an AI-based solution to target this very pain point was conceived, developed, and operationalised. This is OCAI, the Operations CompAnIon that supports decision making of FCTs by providing consolidated results and intelligent insights between related occurrences through a user-friendly interface. In this paper, we discuss the rationale and pain points, the application, and the benefits produced to those missions that supported its development and operationalisation. OCAI creates a data centric approach to codify knowledge from FCTs through a knowledge graph that engineers can navigate through an interface to explore and learn relations, also hidden and non-trivial, between data from different sources. This helps with training new engineers while also increasing the efficiency of FCTs. This work demonstrates how AI-based recommendation engines can prove useful for mission operations.

Keywords: Artificial intelligence, intelligent assistant, knowledge graph, recommendation engine

1. Introduction

With the ever-increasing number of spacecraft, new technologies should be spun in to the conservative realm of mission operations to support the work of the engineers and operators. In 2021, ESOC, together with industry, developed the Artificial Intelligence for Automation (A²I) Roadmap to define a strategic and coherent direction to those future research and development activities that aim to increase automation of mission operations by leveraging AI-based solutions [1]. Since the roadmap was developed to address the pain points of current processes, the benefits of implementing it are plentiful, as are the associated challenges [1].

One of the use cases identified in the A²I Roadmap relates to developing an intelligent AI assistant that can support decision making of Flight Control Teams (FCTs). This emerged as the most promising use case in terms of expected impact. Currently, FCTs have to operate multiple systems to retrieve, combine and analyse data in their day-to-day decisions. For example, important information is recorded in an electronic logbook (Überlog) but there is no existing integration with the Mission Control System (MCS) which contains information on out-of-limit parameters, or the application that provides offline access to historical telemetry (WebMUST), telecommanding and events with rapid plotting and data display capabilities. Ground stations activities are stored in other logs and tools while anomalies are contained in the Anomaly Report Tracking System (ARTS), another independent system. Flight operations procedures are accessible as PDF files in the Flight Operations Plan (FOP). Clearly, navigating through the different tools and systems is cumbersome, time-consuming and requires deep knowledge of the systems. The first step towards achieving a trusted assistant that FCTs can query using natural language, and that tackles this very pain point, is a software application that provides consolidated results and intelligent insights between related occurrences through a user-friendly interface. This is OCAI, the Operations CompAnIon. This paper describes how OCAI was developed in Sec. 2 and its architecture, features, and benefits in Sec. 3. Conclusions and next steps are distilled in Sec. 4.

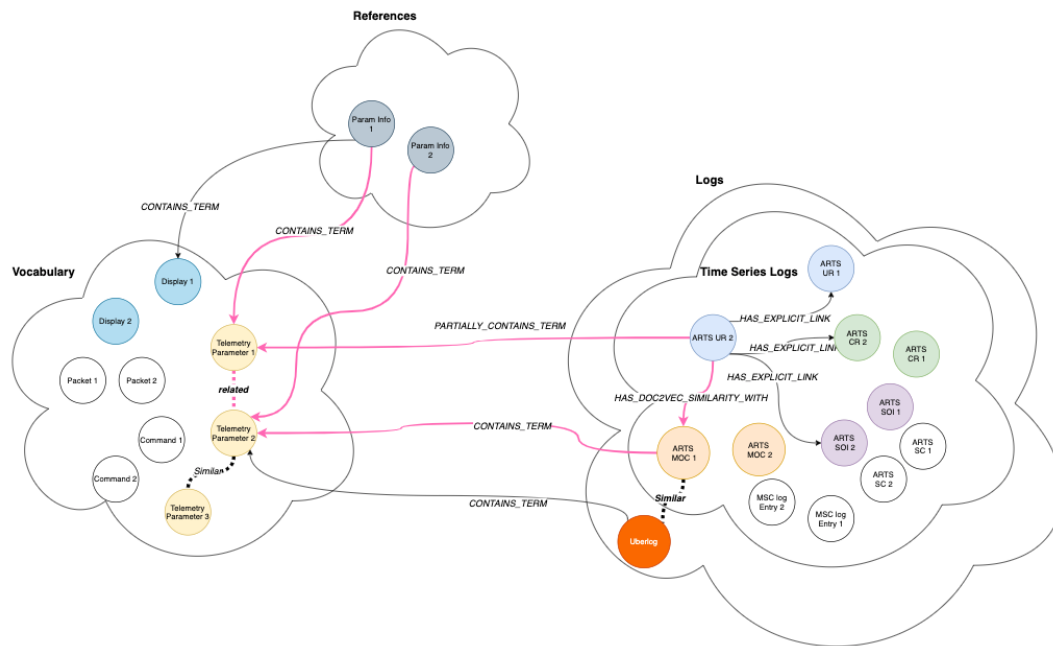


Figure 1. Schematic representation of the OCAI Knowledge Graph (KG)

2. Methodology

As mentioned in [1], two prototypes, were used to further validate the A²I Roadmap. The first, OCAI is discussed here, and the second, ESTIM, an investigation tool for ground stations passes is discussed separately [2]. The development of these two prototypes using an agile approach served to demonstrate the technical and cultural change benefits that implementing the roadmap would produce.

2.1 Technical solution

One of the main challenges of the current legacy systems are the underlying formats in which information is stored. AI can be used to model the knowledge contained in the underlying data sources by building taxonomies or ontologies of concepts along with the types of discovered semantic relationships. Contemporary industrial knowledge modelling approaches are mainly based on building a so-called knowledge base that is implemented either by following the database theoretical foundations (i.e. graph databases) or the foundations established by the Semantic web initiative (i.e. ontologies). Following best practices and considering the maturity of technological solutions, we represented the knowledge base as a graph, knowledge graph, that consolidates various heterogeneous data sources along with discovered concepts and relationships in a consistent format.

The OCAI knowledge graph contains information from different heterogeneous sources – operator and system logs, dictionaries, anomalies, failures, telemetry data, out-of-limit parameters, telecommands, to list a few – for different missions and models non-trivial semantic relationships between them. It makes query answering possible by simply querying and navigating the graph to discover new information relevant to the query. Our knowledge graph is a property graph that consists of a set of nodes (or vertices) and a set of edges (also called links or arcs). A node represents a data object along with its meta information. It can be, for instance, an entry in the anomaly reports logs, a flight operations procedure, or a telemetry parameter. An edge connects nodes and reflects certain relationships between them. For example, an edge can link semantically similar log entries and telemetry parameters together or connect more specific concept to a more generic one (e.g. “Procedure X” subsumes “General Flight Procedure”). Fig. 1 illustrates a schematic example of the OCAI Knowledge Graph (KG) with different types of concepts (e.g., “References”, “ParamInfo”, “Logs”, “ARTS UR”) and relations between them (e.g. “contains term”, “has doc2vec similarity with”). New relations can be discovered by navigating the graph (e.g., similarity between telemetry parameters 1 and 2 via ARTS entries “UR 2” and “MOC 1”). As such, part of the ‘intelligence’ is

contained in the back-end of OCAI, while the front-end provides a query answering capability by querying and navigating the generated knowledge model.

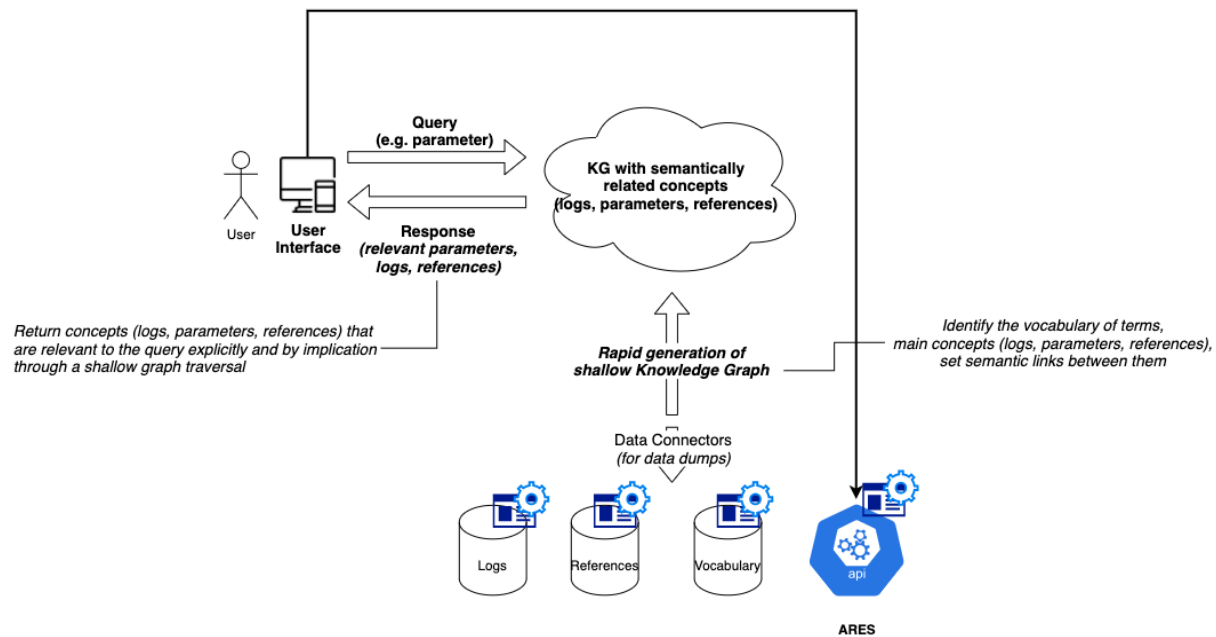


Figure 2. The architecture of OCAI

A query answering capability is implemented in a form of a user-friendly interface that translates users input to a query over the KG. These queries leverage discovered concepts and semantic relations between them and return the most relevant information to the user. The current implementation of OCAI supports written queries.

In summary, as depicted in Fig. 2, OCAI is based on three key building blocks:

- the centralization of data in one location to allow end users to view the information via one unique interface;
- the connection of information across different sources using AI;
- the interaction between end-users and the interface.

Such building blocks define the architecture of OCAI, which is organised in four main modules orchestrated by kedro [3] workflow management framework:

- data wrangling: it processes the raw data (tokenization, stop word removal, lemmatization, stemming) and generates a base representation of the knowledge graph;
- data intelligence: it enriches the knowledge graph with additional semantics through statistical and AI methods, such as TF-IDF, Word2Vec [4], Doc2Vec [5] for similarity relations between documents;
- user interface: it provides user-friendly access to the knowledge graph via query answering capabilities;
- data export: it exports data, reports, and the knowledge graph into the user interface.

Three mission operations teams supported the development of OCAI: GAIA [6], Sentinel-5p [7], and Cluster [8] by providing data and expertise, to develop and validate the tool.

2.2 Implementation

By using an agile approach, an integrated ESA-industry, cross-functional team, combining AI, software engineering, and mission operations domain experts, composed of the authors of the present paper and professionals from McKinsey & Co. and QuantumBlack, developed OCAI in only 10 weeks. This setup ensured close collaboration with the end-users who provided continuous feedback throughout. Since their comments were taken into account and helped shape the tool, their involvement and commitment remained very high during the whole project.

Main milestones of the prototype development included:

- deep analysis of underlying data, its structure and consistency;
- early identification of key datasets;
- data pipeline development;
- early convergence on the requested top-priority functionalities;
- early release of the user interface;
- development of a sound, scalable back-end.

During the second half of the project, the team worked on building the knowledge graph and the query-answering interface for end-users. The prototype functionalities were then tested by end-users with feedback being provided on a regular basis, which was essential to prioritise the backlog and adjust the direction of the overall development process.

2.3 Impact metrics

The impact of OCAI was evaluated both quantitatively and qualitatively. Metrics were defined together with Spacecraft Operations Managers of the three missions that supported the development of OCAI. These were used to assess the benefits of OCAI through interviews and anonymous surveys involving the FCTs. Results are presented in the following section.

3. Results and Discussion

OCAI was developed to support the FCTs with navigating and finding critical information needed to make decisions. It supports query answering for three different missions: GAIA, Sentinel-5p, and Cluster. The query answering is realized through querying and navigating the underlying knowledge graph that consolidates multiple heterogeneous data sources.

3.1 Knowledge graph

The OCAI knowledge graph contains five main groups of concepts and relations:

- Vocabulary to model all discovered terminology along with their mnemonics, textual descriptions, and units of measurements, e.g. TM packets, types of TM parameters, commanding sequences, procedures, and anomaly reports IDs.
- References containing so-called reference tables that make a correspondence between terms (e.g. Out-Of-Limit parameters) and recommended actions.
- Procedures with all their textual content extracted.
- Logs accumulating data from different sources, such as ARTS, MCS, and Uberlog.
- Relations to link relevant concepts representing information objects (i.e., terms, references, procedures, and logs) with each other. Five main types of relationships were used: CONTAINS_TERM and PARTIALLY_CONTAINS_TERM to relate information objects sharing same or similar terminology, HAS_EXPLICIT_LINK to connect object explicitly linked in the original data source, HAS_TFIDF_SIMILARITY_WITH, and HAS_DOC2VEC_SIMILARITY_WITH to set semantic relationships.

3.2 User interface

The user interface was implemented as an analytical web application and consists of 3 main elements:

1. Query input bar (Fig. 3). The header of the webpage allows users to configure their query by selecting a mission, entering their query and specifying the type of the semantic relationship that the user wants to be considered when the knowledge graph is navigated under the hood.
2. Query Answering tab (Fig. 3-6). The main tab providing the information retrieved from the knowledge graph for the specified query. The layout contains several blocks of information corresponding to different data sources:
 - Basic information about the queried term
 - Actions from the cross-reference table
 - Procedures
 - Logs

“Copyright 2023 by European Space Agency. Published by the Mohammed Bin Rashid Space Centre (MBRSC) on behalf of SpaceOps, with permission and released to the MBRSC to publish in all forms.”

- On-board events (including telecommand history, spacecraft events)
 - Related telemetry parameters and the time-series plots for the selected parameters underneath
3. Knowledge Graph tab (Fig. 7). Knowledge Graph tab allows users to get the visual representation of the part of the knowledge graph relevant to their query. Knowledge graph visualisation consists of nodes representing known terms from the vocabulary (TM parameters, packets, procedures, logs, entries from cross-reference table) and links between them representing relationships between them. Different colours of the nodes correspond to different types of the terms, and the labels on the links correspond to the type of the relationship (“CONTAINS_TERM”, “HAS_TFIDF_SIMILARITY”, etc.).

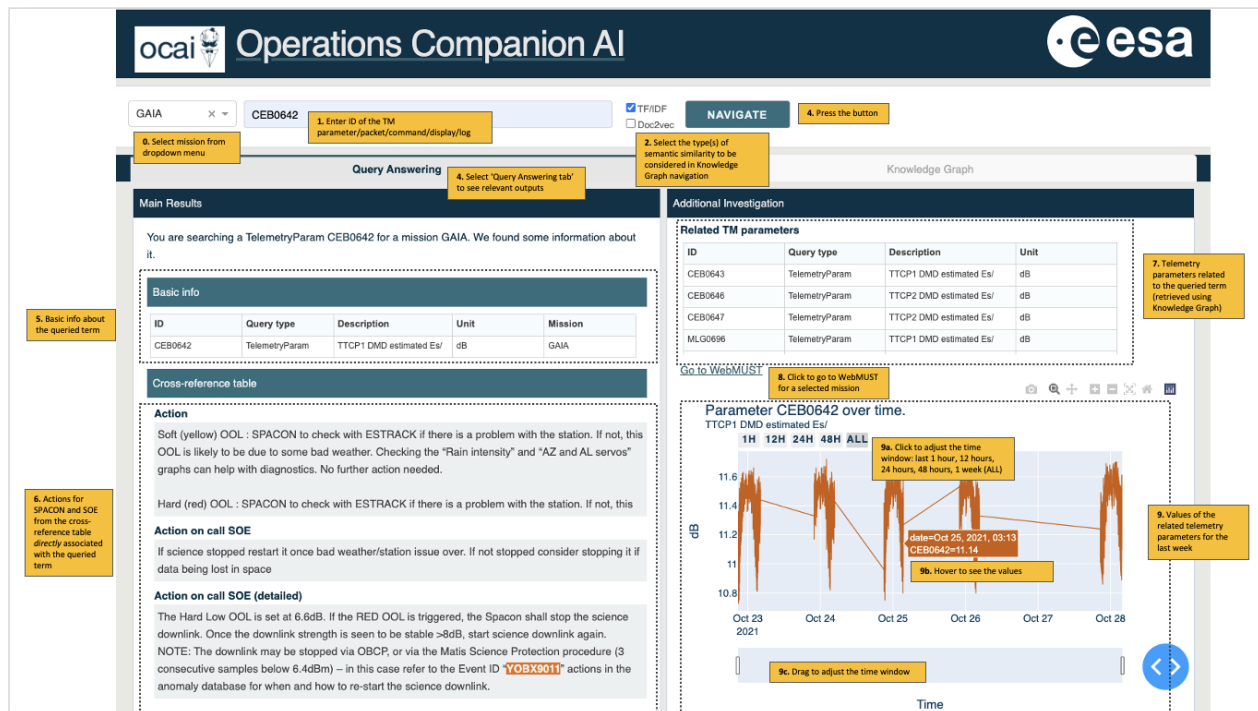


Figure 3: A screenshot of the user interface demonstrating the input bar with formed query and the retrieved information in the Main Results tab: Basic Info, Cross-reference table entries and related telemetry (TM) parameters with time-series plots

“Copyright 2023 by European Space Agency. Published by the Mohammed Bin Rashid Space Centre (MBRSC) on behalf of SpaceOps, with permission and released to the MBRSC to publish in all forms.”

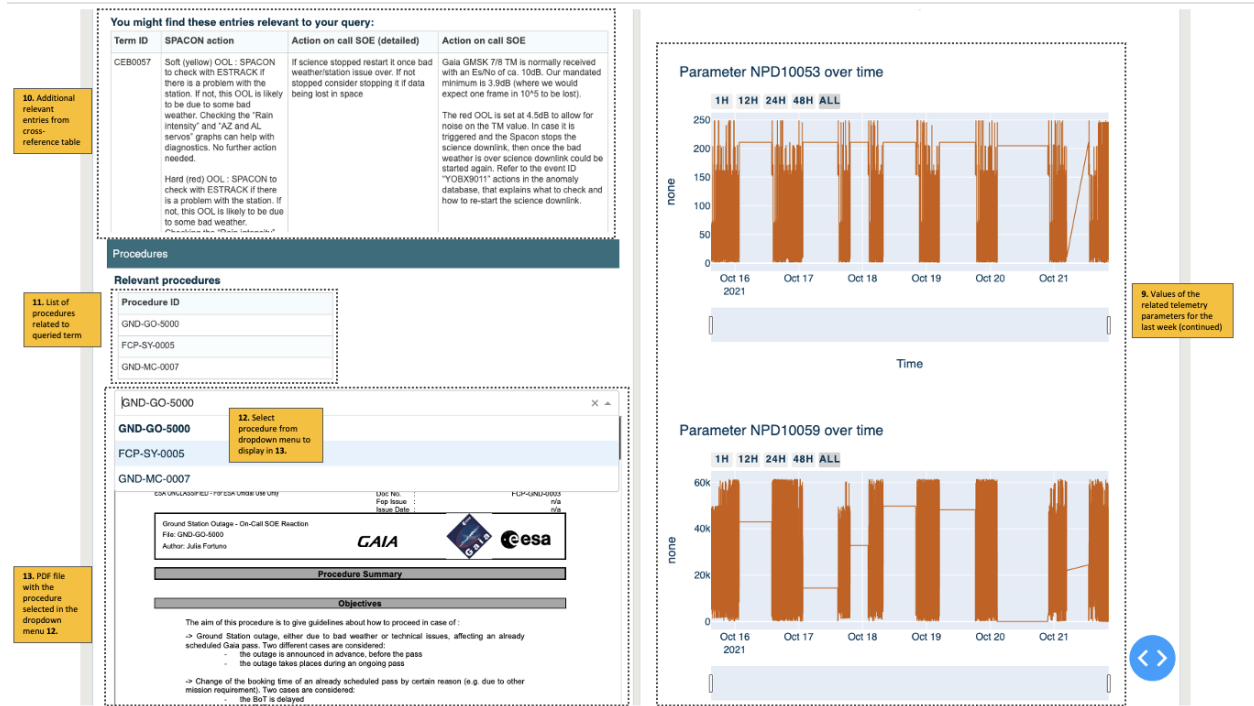


Figure 4: A screenshot of the user interface demonstrating the retrieved information from cross-reference table and procedures as well as time-series plots of related TM parameters

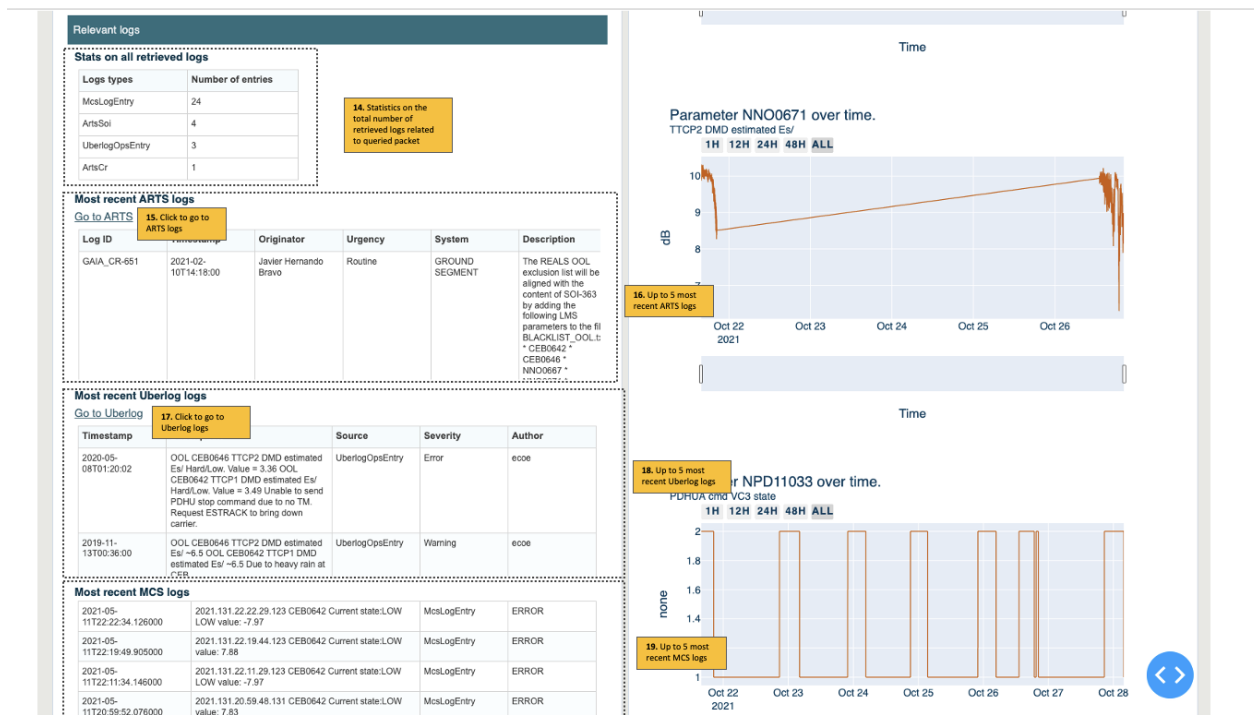


Figure 5: A screenshot of the user interface demonstrating the retrieved information from ARTS, Uberlog and MCS logs as well as time-series plots of related TM parameters

On-board events

Spacecraft events in the last 24 hours

Timestamp	SC event	Description
2021-10-21 17:36:23.640	YDXX2734	5_1 INFO YDXX2734 TC Sequence starts
2021-10-21 17:36:28.264	YDXX2735	5_1 INFO YDXX2735 TC Sequence ends
2021-10-21 20:25:15.567	YPLX2F07	5_1 INFO YPLX2F07 PDHU event DOWNLINK_MODE_TRANSITION
2021-10-22 14:59:36.428	YACX2E7C	5_4 HIGH YACX2E7C ADCS Loss of NM Convergence

26. Spacecraft events in the last 24 hours on the selected mission

TC history in the last 24 hours

Execution timestamp	TC ID	Description	APID
2021-10-21 16:56:24.500	ZPL12801	PL Change value HK param	204
2021-10-21 16:56:25.500	ZV725710	VPu7 SIF sampling PDHU	204
2021-10-21 16:56:26.500	ZPL12801	PL Change value HK param	204
2021-10-21 16:56:35.250	ZPL12801	PL Change value HK param	204
2021-10-21 16:56:36.250	ZV725710	VPu7 SIF sampling PDHU	204
2021-10-21 16:56:37.250	ZPL12801	PL Change value HK param	204
2021-10-21 17:26:04.800	ZSY60011	FILE Mgr request mapping	172
2021-10-21 17:27:38.895	ZSY60009	FILE Mgr request report	172

21. Telecommand history from the last 24 hours on the selected mission

Figure 6: A screenshot of the user interface demonstrating the retrieved information on the on-board events including Spacecraft events and Telecommand history

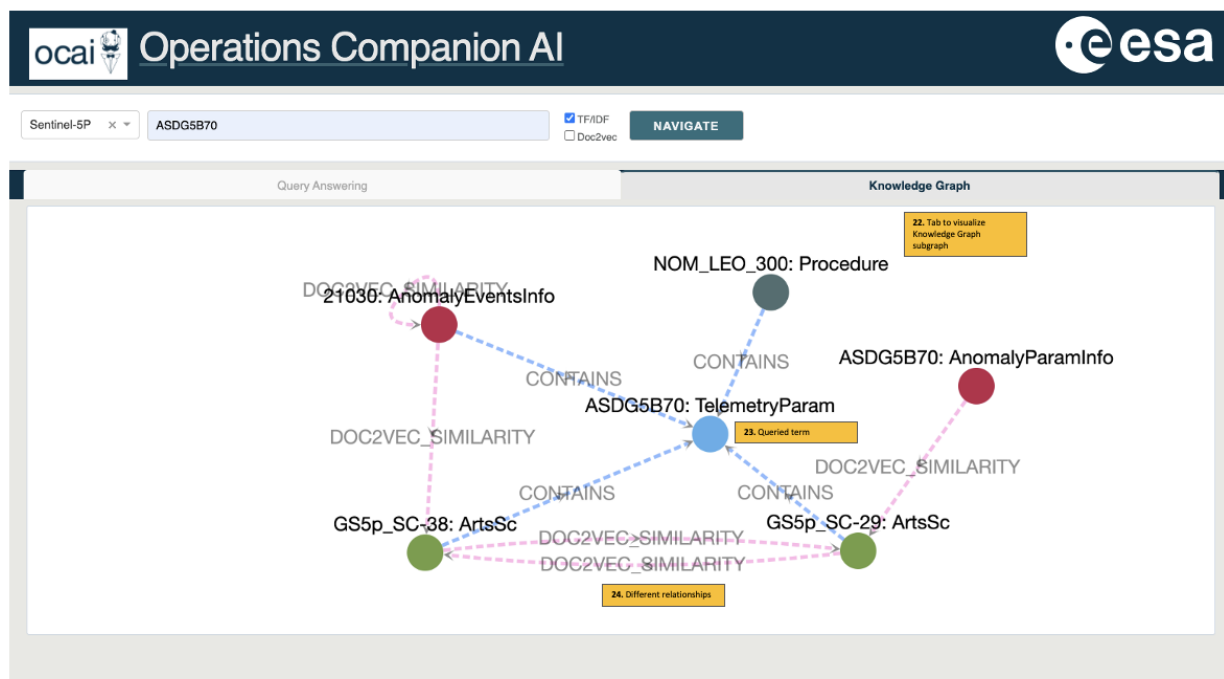


Figure 7: A screenshot of the part of the knowledge graph visualisation for the formed query

3.3 Impact and benefits of OCAI

The results of the interviews and surveys showed that end-users identified clear benefits in having all the resources integrated in a single tool as the time needed to perform the same tasks reduced significantly. End-users were also positively impressed by the quality and relevance of the answers produced by OCAI. The tool even discovered relations that expert engineers were not aware of, demonstrating the potential to enhance their capabilities. The areas with highest potential were related to satellite health monitoring – anomaly handling and follow on, on-call monitoring duties, operations supervision – and more general operations management, including reporting, meetings, maintenance of tools and of flight procedures. Moreover, OCAI was seen as an excellent tool to train new engineers.

4. Conclusions

The scattered landscape of data sources and tools that FCTs need to consult every day leads to inefficiencies, particularly for new members of the team. By consolidating the different data sets and modelling the underpinning relations between them through a knowledge graph, OCAI, an AI assistant with a friendly user interface, was

“Copyright 2023 by European Space Agency. Published by the Mohammed Bin Rashid Space Centre (MBRSC) on behalf of SpaceOps, with permission and released to the MBRSC to publish in all forms.”

developed by a cross-functional team using an agile way of working. The tool demonstrated the impact that AI can have in mission operations, tangibly validating the A²I Roadmap [1] developed by ESOC and industry.

Intelligent assistants have the potential to improve current processes by streamlining the manual and tedious tasks performed to find answers but also to discover new, hidden relations in the data. It is our hope that the work presented here will stimulate further and future developments to improve the current application and also expand the domain of application. ESOC is currently working on this already.

Acknowledgements

The authors would like to thank the team from McKinsey and Co., QuantumBlack and Airbus Defence & Space for their work and the FCTs who provided valuable feedback and comments during the agile development.

References

- [1] G. De Canio, J. Eggleston, J. Fauste, A. M. Palowski, M. Spada, Development of an actionable AI roadmap for automating mission operations, SpaceOps ID # 303, 17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023.
- [2] G. De Canio, J. Eggleston, A. M. Palowski, M. Unal, ESTIM: the ETrack Investigation and Monitoring tool for analysis of ground station passes, SpaceOps ID # 307, 17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023.
- [3] S. Alam, N. L. Chan, G. Comym, et al., Kedro (Version 0.18.4) [Computer software], <https://github.com/kedro-org/kedro>, 2022.
- [4] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient Estimation of Word Representations in Vector Space, ICLR (Workshop Poster) 2013.
- [5] Q. V. Le, T. Mikolov, Distributed Representations of Sentences and Documents, CoRR abs/1405.4053, 2014.
- [6] ESA, GAIA, https://www.esa.int/Science_Exploration/Space_Science/Gaia (accessed 30/01/23)
- [7] ESA, Sentinel-5p, https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-5P (accessed 30/01/23)
- [8] ESA, Cluster, https://www.esa.int/Science_Exploration/Space_Science/Cluster (accessed 30/01/23)