17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023.

SpaceOps-2023, ID # 303

Development of an actionable AI roadmap for automating mission operations

Gabriele De Canio[1]*, James Eggleston[1], Jorge Fauste[2], Artur M. Palowski[1], Mariella Spada[1]

[1] *European Space Operations Centre, European Space Agency, Robert-Bosch-Str. 5, 64293 Darmstadt, Germany*
[2] *European Space Astronomy Center, European Space Agency, Camino bajo del Castillo, E-28692 Madrid, Spain*
* Corresponding Author, gabriele.decanio@esa.int

Abstract

Over 14 weeks in 2021, the European Space Operations Centre of the European Space Agency in collaboration with industry developed the Artificial Intelligence for Automation (A^{2}I) Roadmap. It consists of use cases spanning across the whole spectrum of mission operations, prioritized based on foreseen impact and technological feasibility. Implementing the A^{2}I Roadmap entails the investigation and development of AI-based applications for automating, among others, predictive infrastructure maintenance, mission planning, operational simulation, mission operations systems testing and validation, spacecraft health monitoring, problem diagnosis, and the recommendation of actions to spacecraft operators. In this paper, we provide an overview of the method used to develop the roadmap, the associated use cases, and its benefits, ranging from increasing efficiency to refocusing intellectual and financial resources. We also discuss how to bring AI solutions into the risk-averse realm of mission operations and its inherent challenges.

Keywords: AI, automation, roadmap, mission operations, innovation

1. Introduction

Artificial intelligence (AI) is typically defined as the ability of a machine to perform cognitive functions that are associated with human-like tasks, such as perceiving, reasoning, learning, interacting with the environment, problem solving, and even exercising creativity. The significant increase in computing power and storage, data proliferation and its accessibility, and algorithmic advancements have yielded to an unprecedented focus on AI over the past years, which ultimately resulted in breakthroughs in multiple applications. For instance, AI is proving its value in robotics, autonomous vehicles, computer vision, natural language, virtual agents, in industries ranging from healthcare to automotive and finance. The adoption of AI continues to grow, and the technology is generating returns. Recent findings showed a nearly 25 per cent year-on-year increase from 2018 to 2019 in the use of AI in standard business processes as well as across multiple areas of their business [1]. These advancements in AI can benefit the space industry as well, in particular for the automation of mission operations, where data is abundant and typically well-structured. Greater automation leveraged by AI can bring substantial benefits as it offers the opportunity to reduce mundane and repetitive tasks and free up the time of scientists and engineers.

The European Space Operations Centre (ESOC) at the European Space Agency has for years been at the forefront of exploring the use of AI and developing AI-based applications for mission operations. Numerous rising trends in the space industry are exploiting opportunities driven by AI as well. These include, for example, mega-constellations, in-orbit servicing, and space debris management, which entail complex mission operations tasks. In 2021, to provide a coherent and strategic direction to the future developments of AI applications for automating mission operations, ESOC, in close collaboration with industry, developed the Artificial Intelligence for Automation (A^{2}I) Roadmap [2]. Considering its unique position as it covers all domains of complex mission operations, ESOC was used as a benchmark to assess the foreseen impact and technological feasibility.

This paper describes how the roadmap was conceived, developed and shaped in Sec. 2, presents the main results in Sec. 3 and discusses its implementation challenges in Sec. 4. Final considerations are summarised in Sec. 5.

2. Approach

Over 14 weeks, a cross-functional ESA-industry team that combined expertise in AI, data analytics, automation and mission operations, developed a first version of the A^{2}I Roadmap. The team, including the authors of the present paper as well as professionals from McKinsey & Co., QuantumBlack, and Airbus Defence & Space, used a first-principle approach in which mission operations were mapped onto a typical mission lifecycle, from Phase 0 (mission analysis/needs analysis) to Phase F (Disposal). Relevant stakeholders involved in the various Phases were interviewed to capture existing pain points, limitations, and areas of improvement, using activities at ESOC as a point of reference. It is worth mentioning that the work presented here was carried during the early stages of the COVID-19 pandemic, hence forcing a remote work execution, with all the difficulties that this entailed.

After consolidating the collected inputs, the team organised dedicated workshops with end-users on thematic domains: mission analysis, flight dynamics, mission preparation, routine operations, space debris, space weather, just to name a few. Each workshop focused on exploring ways in which AI could be leveraged to bring more automation to those repetitive and mundane tasks previously identified and solve, partially or entirely, the recurrent problems that affected the daily job of end-users. Ideas of possible use cases were then conceived and further elaborated, taking into account prior work and ongoing activities both at ESOC and within both the tech and space industries.

The following step entailed prioritising the use cases based on technological feasibility and expected impact in collaboration with additional AI and mission operations experts. Next, we assessed the potential for synergies between cases to maximise their cross-development, reduce deployment time and ease scaling. This led to the definition of a first draft of the roadmap.

The findings were then presented at multiple events to further validate the roadmap with an outside-in perspective and collect additional inputs from industry. This yielded to the A^{2}I Roadmap [2]. Importantly, it is not meant to be a static set of use cases, but rather a living and dynamic process that will evolve with new findings, feedback from industry and future needs. Further details and analyses that describe our approach in greater detail could not be shared and are therefore omitted in this paper.

An additional step to validate the impact of the roadmap was to develop two prototypes from the shortlisted use cases, OCAI, a virtual assistant to support flight control teams, and ESTIM, an investigation tool for ground stations passes. The prototypes were developed by cross-functional teams in only 8 and 10 weeks, respectively. Details on their implementation and technical solutions are presented in these proceedings [3, 4].

3. Results

The A^{2}I roadmap comprises 14 highly promising use cases within 5 domains, grouped as pre- and post-launch applications, as depicted in Fig. 1. The domains are:

Pre-launch applications

- Operational simulation and operations preparation
- Systems testing and validations

Post-launch applications

- Ground system maintenance and ground operations
- Flight control team and ground stations
- Satellite health monitoring and data processing

The use cases address the work of flight control teams in operations preparation and during operations via automation, AI-based diagnostics and prognostics as well as AI-based assistance. Other use cases relate to mission planning, predictive maintenance and intelligent incident classification, and intelligence and automation in systems testing and validations. Detailed descriptions can be found in Appendix A. All in all, our analyses revealed that AI can significantly support the automation of mission operations by increasing efficiency and reducing overall workload, freeing up mission operators from mundane and repetitive tasks.

Operational simulation and operations preparation	Systems testing and validations	Ground system maintenance and ground operations	Flight control team and ground stations	Satellite health monitoring and data processing
Intelligent assistance for automatic procedure generation and validation	Deviation analysis in validation and regression testing outputs	Machine learning-based incident classification and root-cause analysis assistance	AI decision support tools for controllers	Intelligent root-cause investigation and AI-assisted handling
AI assistance to support simulation validation and failure identification	AI-supported test case creation	AI-based predictive maintenance	Intelligent digital planning platform	AI-based long-term satellite health forecasting
	Automated test report generation		Short- to medium- term operations planning	AI-based short-term satellite health forecasting
				Intelligent telemetry data anomaly detection

Figure 1. The A^2I Roadmap consists of 14 use cases within 5 domains.

Rather than focusing on the actual use cases, here we present the archetypes that emerge when analysing the pertaining and underlying AI types.

1. **AI-enhanced modelling**. A few use cases address the improvement of classical models through AI-based models, e.g. reinforcement learning or deep neural surrogates, in order to improve mission-critical analyses, such as simulations and optimization calculations in mission analysis or operational simulation. Applying AI-based models to these issues may increase the accuracy and precision of analysis outcomes.

2. **Automated content generation**. Across domains, typical routine tasks consist of the repetitive generation of content such as reports (e.g. test reports, anomaly handling reports, incidence reports), the translation of text into technical language and telemetry, and the case-specific compilation of building blocks, e.g. generation of test cases. These tasks can be automated across domains based on artificial intelligence, by deploying deep learning algorithms on historic structured and unstructured data in order to provide customized output based on the input. AI technologies such as natural langue processing including transfer learning (pre-trained text algorithm) can fulfil text generation and translation tasks on an automated basis to aid the creation of structured reports and documents, as well as the automated analysis of text.

3. **Automated diagnostics**. Improving current, often manual or time intensive diagnostics of anomaly detection in telemetry, incidents and failures in ground systems or data quality checks on space weather data are promising use cases for AI. Through the implementation of continuous regression and discrete classification models, new patterns and correlations in data can help to detect anomalies faster and to find hidden anomalies, to investigate root-causes of anomalies and system failures, and to enable for faster solution diagnostics. This would also support the detection of false positives and detect whether an alert requires human intervention.

4. **Automated prognostics**. Similarly to the previous archetype, automated prognostics generates deeper knowledge based on data analysis – however, it aims to make prediction about potential future issues, such as short-term incidents, and long-term spacecraft health incl. the lifetime of subsystems. Prediction algorithms for time series forecasting are used to flag potential upcoming issues before they occur for operators to take preventative measures in order to reduce downtime or time spend in safe mode. Prognostics can further be utilized for planning and predictive maintenance tasks.

5. **Decision recommendation engine**. Decision recommendation engines analyse past human behaviour and decisions based on past documents, log files and reports through fuzzy techniques, transfer learning, genetic algorithms, evolutionary algorithms, neural networks and deep learning, and active learning. This enables these decision assistants to recommend best actions for anomaly handling, operations planning, manoeuvres or maintenance schedules. It is important to highlight that decisions are still made

by the operators; however, an AI-based decision assistance can help to significantly improve and speed up root-cause investigation and the subsequent derivation of (alternative) best handling steps.

6. **Visually enhanced simulation**. Augmentation of systems simulations and their interactions with the environment, such as spacecraft digital twins in mission simulations and during mission operations, can be used to generate additional data, have an updated copy of the current status and configuration of the spacecraft, and run simulations and predictions.

7. **AI enablers**. Several use cases were identified that rather act as enablers for AI solutions, for example through the generation of high quality, structured data, improvement of user interfaces and data visualization. They address more general issues of digitization, the improvement of data readiness and the establishment of required interfaces for interaction with AI-based solutions.

4. Discussion

4.1 Benefits

By using ESOC as a benchmark given its vantage point, we developed a roadmap that has its roots in existing pain points related to mission operations. This gave confidence on the solidity of its foundations. Importantly, when conceiving the AI solutions to address the pain points, focus was not on using AI for the sake of it, but rather on ways AI could be leveraged to automate tedious and mundane processes and tasks and enhance tools and applications. The validation steps undertaken reinforced the value and quality of the A^{2}I Roadmap.

Another important aspect is commercialisation. The industry is eager to engage with ESOC and use its infrastructure to develop, test and validate solutions that can be potentially commercialised and adopted by other space agencies and satellite operators.

With the primary aim of fostering European space sector growth in the use of applied AI in mission operations, the A^{2}I Roadmap introduces a range of additional benefits:

- Reinforces the use of AI in the operations domain, responding to ESA Agenda 2025 [5] and the ESA Technology Strategy [6]
- Increases efficiency and reduce costs of space programmes
- Enhances industry competitiveness by refocusing intellectual and financial resources
- Frees up personnel from mundane tasks to focus on more ambitious and rewarding challenges

4.2 Challenges in implementing the A^{2}I Roadmap

Historically, space is a conservative industry. Mission operations is even more so. New Space has challenged the "traditional space" sector with new ways of working, business models, technology, and products both for downstream and upstream applications. Results achieved by New Space startups are remarkable and have paved the way for embracing and spinning in novel technologies more openly. AI is no exception. However, developing and deploying AI-based solutions to mission operations is a multi-faceted endeavour that entails both technical and change management challenges. When dealing with AI, ethical aspects should also not be overlooked as the associated risks can be substantial. For the sake of brevity, these are not included in our discussion but it is worth mentioning that the ESA Data Security Policy was adhered to throughout as well as in the development of the prototypes [3, 4].

4.2.1 Technical challenges

On the technical side, we can identify the following challenges.

- **Infrastructure**. Developing and deploying AI-based applications requires the appropriate infrastructure layer that allows for their productionalization, maintenance, and monitoring. Inspired by the continuous development concept of DevOps in traditional software, we believe that an MLOps platform is the fundamental backbone and enabler to all AI applications. Importantly, the platform needs to allow parallel development and deployment in a dedicated environment that mimics all systems but that does not affect live mission operations.

- **Data**. When analysing the underlying archetypes associated with the use cases of the A^{2}I Roadmap, we can clearly recognise that machine learning is the kernel for many of them. AI in general and machine learning especially require clearly, structured data that can be easily accessed in a secure fashion. For this, another technical challenge we identified relates to data and storage. A data strategy that ensures standardized and consistent data is created, accessed, stored and versioned through a solid data governance and data architecture is also fundamental. Moreover, the data pipeline to curate and preprocess data should additionally consider potential scarcity of labelled data sets as well as highly unbalanced data sets. To address these first two technical challenges, infrastructure and data that are highly intertwined, ESOC is working with industry to develop a platform called AInabler [7, 8].

- **Explainability**. Another technical challenge relates to developing explainable and interpretable AI-based solutions. Being able to navigate and understand the decision-making process the AI took to produce specific results from a given set of inputs is paramount to ensure adoption. End-users need to trust the solution they will be using daily, and for this it must be "transparent". Developers should prioritize AI models that can accommodate this need [9 - 13].

- **Data privacy**. Preserving data privacy in a federated and distributed setting is also of relevance when scaling the AI solution across multiple spacecraft and missions as sharing data across may be forbidden as per the specific export controls regulations. Also, data that contains personally identifiable information (PII) shall be treated with care. Crucially, AI models and the underlying infrastructure need to take into account these aspects as well. Examples of relevant techniques are federated learning and differential privacy [14 - 16].

- **Continuous improvement**. In order to improve accuracy of the solutions, these should be designed to allow for retraining via feedback from end- users / domain experts and via up-to-date stream of data. Techniques such as active learning and continual learning can be used [17 - 19]. Considering the pace at which AI evolves, it is expected that novel AI models and frameworks will be continuously developed. Hence, the applications will require to be updated or replaced. It is therefore fundamental to ensure modularity in the solutions.

4.2.2 Change management

In addition to technical aspects, implementing and deploying sustainable AI solutions requires to tackle other challenges that are here considered part of the change management effort.

- **Agile way of working**. Rather than having all the system or application requirements defined upfront, by working in cross-functional, agile teams consisting of technical and strategic roles, requirements are defined and continuously fine-tuned throughout the process. For this, close interactions with the end-users are paramount. This different way of working ensures adoption as end-users are integral part of the process and can really steer the development of the applications.

- **Performance metrics**. Despite the large potential and success that AI has had in other disciplines and sectors, being objective on the actual impact of the AI applications in the business is critical. The benefits of each of the developed AI applications shall be assessed through quantitative and qualitative performance metrics or KPIs.

- **Scaling**. We foresee that each use case will be initially implemented by creating an initial proof of concept (PoC) to demonstrate key features on a specific mission. Then, the PoC will be further enhanced with additional features and validated and tested in a real operations environment. Next, such am application will be extended to other missions as well. Having in mind scalability from early design of the PoC is instrumental to ensure a successful development and deployment of applications at scale.

5. Conclusions

In this paper, we presented not only the A^{2}I Roadmap but also the expected challenges in implementing it. We believe that these are applicable for the development of AI-based solutions well beyond the current set of use cases for mission operations. We hope that our work will stimulate other space agencies to publish their roadmaps and that our approach, methodology and challenges can be used as a reference when defining other technology roadmaps.

Acknowledgements

The authors would like to thank the more than hundred people involved in defining, developing and shaping the roadmap, both from the European Space Agency and from Industry, especially McKinsey & Co., QuantumBlack, and Airbus Defence and Space for their work.

Appendix A – Artificial intelligence for Automation Roadmap: Use cases description

In the following, we report more the descriptions of the 14 use cases within 5 domains grouped as pre- and post-launch applications.

Pre-launch applications

<u>Operational simulation and operations preparation</u>

1. Intelligent assistance for automatic procedure generation and validation

Today, the procedure writing is a manual and time-intensive process. Furthermore, the validation of the system also requires high manual effort. Currently these processes are also missing an autonomous consistent flight operations plan update. From procedure writing to procedure validation, the process requires several steps taking several days, and needs to be fully revisited to make changes, which makes iterations very slow. This use case aims at developing an AI system that will support the operator through automatically suggesting elements to be included in the procedures (e.g. pre-condition or post-condition checks for every step). The system will translate the high-level context language into specific and technical actions. Furthermore, the AI system could provide suggestions for changing the procedure in case some elements of the spacecraft change (e.g. database or software parameters). Lastly, the AI could help optimize and fine tune procedure parameters (e.g. delays) based on the outcomes of the simulator.

2. AI assistance to support simulation validation and failure identification

A very high manual effort is required to generate breakpoints as well as prepare and validate the simulators. Despite simulations can be run quickly, the underlying physics models feeding into the simulations are computationally intensive and can require long times to be executed. Hence, it is difficult to use these simulations to support time sensitive decision making. The foreseen AI solution will leverage heuristics learned during training to reduce the computational time and reduce the number of calculations to obtain the same results, i.e. generate fast simulators. The AI solution will also allow for automated validation and failure identification of simulator deliveries, on-board software, databases as well as enable time sensitive decision making.

<u>Systems testing and validation</u>

1. Deviation analysis in validation and regression testing outputs

Currently, investigating system failures based on simulator data is very-time consuming and it requires great expertise and research in past logs and supplier documents to find root-causes. The complexity is increased by the system of systems architecture. In the future, if an issue occurs, a framework should be available to efficiently examine aggregated logging information to help detect similar issues that occurred in the past as well as potential solutions. Data mining can be used to analyse a large number of logs and identify line items, which are most likely going to explain root-cause of the issue. This will decrease the amount of time needed for the root-cause analysis of the problem. Furthermore, the isolation of an error in a system of systems becomes more efficient.

2. AI-supported test case creation

Adapting to new systems requires a high manual effort. As systems are complex, numerous tests are required to validate the correct functioning. Especially regression tests require a high manual effort and are thus cost intensive. The AI solution will support experts by automating the creation of test cases for ground segment software applications. The AI solution will recognise patterns contained in sets of high-quality test cases to automatically generate a large amount of reliable test cases that are customised to the system and satisfies constraints. These test cases can then be used for validating the viability of the solution.

3. Automated test report generation

High manual effort is required during the systems testing and validations phase to generate test reports and analyse them to decide on next actions to be performed, as test reports are typically written manually. The use case aims at developing an AI-based system to change the process by aggregating all relevant data and use Natural Language Processing (NLP) techniques with transfer learning to automatically generate the content of the test reports.

Post-launch applications

<u>Ground system maintenance and ground operations</u>

1. Machine learning based incident classification and root-cause analysis assistance

The lack of a centralized/standardised way of aggregating and analysing problem reports, events and incidents across systems and missions creates a high manual effort to assess incidents and check similarity to previous reports, thus implying an increased time to resolution.

The use case aims at developing an intelligent AI assistant that can support the decision making of engineers by accessing records from all different missions, showing clusters, system versions, etc. and linking entries with their context and the respective resolutions (e.g. using a knowledge graph). This will enable a fast investigation of correlations between problems and system configurations. Furthermore, it will detect historical patterns and make predictions on what type of activity may cause an increased number of incidents (e.g. a certain software rollout). This also serves as the basis for predictive maintenance.

2. AI-based predictive maintenance

Currently, there is no systematically established predictive maintenance of ground system software and systems on a subsystem level, which means that down-times can be prevented through data processing, trends identification and recommendations suggestion. By doing so, alerts/reports for maintenance activities can be sent to responsible entities before the actual down-time occurs (this includes downtimes of local hardware like frequency timing system, cryo-feeds, power amplifiers and ground segments). Given the broad scope of this use case, multiple AI types will be required, including algorithms for data mining to process system data, natural language processing for report analyses and generation, continuous time-series models for regression to predict future down-times and a recommendation engine to propose concrete actions.

<u>Flight control team and ground stations</u>

1. AI decision support tools for controllers

In case of spacecraft contingency, controllers currently need to gather information from multiple sources (e.g. flight operations procedures, live information and historical data from the control system) to analyse the anomaly and find the root causes. This activity is very time consuming and requires a lot of expert knowledge. Furthermore, it is challenging for one operator to control several spacecrafts in parallel. This use case aims at developing an AI system to automatically process the data from the sources to find correlations and potential root causes. The solution would have an underlying smart search engine for procedures, to find the correct contingency procedures to be applied in specific events. It will also recommend actions to be undertaken for anomaly handling based on past resolutions of such events and show pre-defined recovery options to resolve contingencies across all mission phases.

2. Intelligent digital planning platform

Nowadays, many iterations with flight control teams and ground stations planner are needed in order to resolve capacity conflicts in ground station pass planning. The AI solution will build an integrated platform that all stakeholders of the end-to-end planning process could use. The transfer of information between stakeholders will become seamless, more efficient and standardised. The platform will serve as the basis for descriptive, predictive and prescriptive AI applications. Data exploration, mining and analytics will enable additional real-time insights that are available across different teams. Furthermore, automated recommendations for optimized handling steps based on AI algorithms will be delivered.

3. Short- to medium- term operations planning

Planning and replanning processes of mission operations are very time consuming. The developed AI system will learn from past planning files to understand human decisions that were taken and apply these to new planning processes. The current tacit knowledge from operators and flight dynamics engineers that went into the planning process would thus be mimicked by the model.

Satellite health monitoring and data processing

1. Intelligent telemetry data anomaly detection

Current tools such as DrMUST [20] and Novelty Detection [21, 22] provide good solutions for investigating anomalies. However, the manual configuration of the nominal behaviour of parameters to use is very time intensive and difficult. It requires an in-depth understanding of the systems and the nominal behaviour of parameters. The use case aims at developing a user-friendly application that would build on the existing solutions, linking them to each other, and improve the underlying AI techniques for outlier detection. Furthermore, simulators data (or data from suppliers) could be fed into the system before the start of the mission, in order to train the algorithm. This will then be further enriched with flying data.

2. Intelligent root-cause investigation and AI-assisted handling

Root causes of telemetry data anomalies are assessed manually and require long time. The manual process entails plotting parameters over a long time period, analysing the resulting graphs, and looking for relations in the data. This effort requires knowledge of mission-specific details on activities potentially having impact on the parameters to allow drawing the right conclusions. The AI solution will analyse behaviour of parameters in historical telemetry data of the respective mission, detect of similarity in trends of other parameters at the time of the incident and compare of telemetry signature with comparable activities in the past. A list of potential root-causes will be generated based on the automatic, AI driven analysis with automatically generated recommendations for anomaly handling based on historic data.

3. AI-based long-term satellite health forecasting

For interplanetary missions, a 2-year forecast of the evolution of the mission is required (e.g. degradation, fuel, lasers, switch functionality). Forecasts are also required for the bi-yearly science operations coordination. The use case aims at using AI for long-term forecasts of the evolution of subsystems of the spacecraft. Over time, the model would learn to provide insights on the aggregated effect of the mission on subsystems. Correlation analysis with environmental conditions will also be used to forecast their impact on the spacecraft, eventually supporting the associate mission planning.

4. AI-based short-term satellite health forecasting

Monitoring spacecraft and addressing occurring anomalies are time-consuming tasks. The AI solution will use historic telemetry data to forecast the behaviours (i.e. timeseries) of the spacecraft systems and subsytems to predict onset of unwanted events in the short-term thus enabling operators to identify issues before they occur and address the root causes in a timely manner.

References

[1] Global AI Survey: AI proves its worth, but few scale impact, McKinsey 2019, https://www.mckinsey.com/featured-insights/artificial-intelligence/global-ai-survey-ai-proves-its-worth-but-few-scale-impact, (accessed 22.01.23)

[2] A²I Roadmap, https://esoc.esa.int/a2i-roadmap, (accessed 22.01.23)

[3] G. De Canio, J. Eggleston, A. M. Palowski, M. Schmidt, P. Collins, B. Sousa, D. Mesples, R. Santos, E. Ntagiou, OCAI: the Operations CompAnIon to support decision making of flight control teams, SpaceOps ID # 305, 17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023
[4] G. De Canio, J. Eggleston, A. M. Palowski, M. Unal, ESTIM: the ESTrack 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
[5] ESA Agenda 2025, March 2021
[6] ESA's Technology Strategy, Version 1.2, August 2022
[7] ESA, Ground Segment Operations Automation Using Artificial Intelligence
[8] ESA, Machine learning platform of spacecraft operational datasets
[9] Y. Gal and Z. Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. Proceedings of the 33rd International Conference on Machine Learning, 48:1050–1059, 2016
[10] F. K. Došilović et al., Explainable artificial intelligence: A survey. 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2018
[11] A. F. Markus et al., The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies. Journal of Biomedical Informatics, 2021
[12] A. G. Baydin et al., Etalumis: Bringing Probabilistic Programming to Scientific Simulators at Scale. Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 29, 2019
[13] E. Zelikman et al., Short-Term Solar Irradiance Forecasting Using Calibrated Probabilistic Models. arXiv, 2020
[14] J. Konečný et al., Federated Optimization: Distributed Machine Learning for On-Device Intelligence, 2016
[15] P. Kairouz et al., Advances and Open Problems in Federated Learning, 2021
[16] R. Shokri and V. Shmatikov, Privacy-preserving deep learning. 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2015
[17] Y. Gal et al., Deep Bayesian Active Learning with Image Data. Proceedings of the 34th International Conference on Machine Learning, PMLR 70:1183-1192, 2017
[18] S. Farquhar and Y. Gal, Towards Robust Evaluations of Continual Learning. Privacy in Machine Learning and Artificial Intelligence workshop, ICML, 2018
[19] G. I. Parisi et al., Continual lifelong learning with neural networks: A review. Neural Networks, 2019
[20] J. Martínez-Heras et al., DrMUST - A data mining approach for anomaly investigation, 2012 SpaceOps Conference, 2012
[21] J. Martinez et al., New Telemetry Monitoring Paradigm with Novelty Detection, 2012 SpaceOps Conference, 2012
[22] J. Martínez-Heras et al, Enhanced Telemetry Monitoring with Novelty Detection, AI Magazine 35 (4), 37-46, 2014