

Next-Gen Machine Learning Supported Diagnostic Systems for Spacecraft

Athanasios Vlontzos^{1*}, Gabriel Sutherland², Siddha Ganju³ and Frank Soboczinski⁴

¹Imperial College London, ²Oregon State University, ³NVIDIA Corp, ⁴King's College London
athanasios.vlontzos14@ic.ac.uk, sutherlg@oregonstate.edu, sganju@nvidia.com,
frank.soboczinski@kcl.ac.uk

Abstract

Future short or long-term space missions require a new generation of monitoring and diagnostic systems due to communication impasses as well as limitations in specialized crew and equipment. Machine learning supported diagnostic systems present a viable solution for medical and technical applications. We discuss challenges and applicability of such systems in light of upcoming missions and outline an example use case for a next-generation medical diagnostic system for future space operations. Additionally, we present approach recommendations and constraints for the successful generation and use of machine learning models aboard a spacecraft.

1 Introduction

Space agencies have a renewed drive to take human space exploration beyond Low Earth Orbit (LEO) and into deep space. NASA's Artemis program outlines a clear path to return to the Moon and to go beyond to Mars [NASA, 2017]. Additionally, recent successes in the commercial space sector by major players such as SpaceX and Blue Origin make human spaceflight more accessible, affordable and future long term-missions a reality. Yet future long duration spaceflight require systems that are independent of LEO operations such as constant communication, the ability to transfer large amounts of data via multiple systems in a relatively short time-frame or the ability to request and exchange crew if needed. On Earth, machine learning (ML) and machine automation is already driving the next industrial revolution and resulted in fully autonomous industrial processes in domains such as agriculture as well as manufacturing [Ayaz *et al.*, 2019; Yang *et al.*, 2019]. Spaceflight itself, however, is far behind such advances. Here we discuss challenges ML supported systems face in the space domain as well as the applicability and advantages of ML systems on a spacecraft. We highlight aforementioned items via an example of an autonomous medical system and describe an infrastructure for the successful development of such systems.

*Contact Author

2 Challenges for ML aboard Spacecraft

Space is hard and manned space exploration is dangerous and unforgiving. Moreover, the long term effects of human presence in microgravity are still not fully known. Current space missions do rarely include a medical officer among the crew and rely on specialists who are also trained in emergency medicine. While this is sufficient for minor and trained emergency cases, it does not allow for more serious and complex medical treatments. Health related checks and emergencies are handled in a tele-medicine regime where instructions are communicated to the astronauts via ground to spacecraft channels. In extreme emergencies the astronauts can always make the journey back on Earth. As missions veer further away from Earth, returning for medical treatment and relying on simultaneous communications becomes infeasible as both distance and communication latency increases exponentially. For a deployed ML system aboard a spacecraft there are several challenges to consider: a) limited live testing abilities available, and testing in the form of payloads on missions are expensive; b) systems that are deployed need to be at a high technology readiness level (TRL) [Lavin *et al.*, 2021]; c) environmental effects may influence deployed systems. For example, how ionizing radiation can affect deployed space capable hardware [Page and Benedetto, 2005], affecting the consistency of sensor behavior such as early or late sensor fusion; leading to potentially corrupted data to be processed e) payloads are constrained by weight, so shipping large compute infrastructures is infeasible; f) fully autonomous applications can be found mainly on controlled environments that assume almost complete access to information and environmental parameters; g) and more importantly the lack of labeled data for each task along with limited interpretability and explainability of current ML systems add to the complexity. Contrary to that, any space-faring vehicle is faced with extreme environmental conditions that not only are hard to control or predict, but in some cases are challenging to human scientific comprehension. Hence, ML systems must be as robust as possible to changes of their operating environments.

3 Applicability & Advantages

Incorporating autonomous processes in a spacecraft is a complex task as technology and the associated needs constantly

develop. Some key points, however, are: (1) Reduce Latency and Earth Reliance - A significant amount of operations aboard modern day spacecrafts and the ISS, require the constant communication with mission control on Earth [Dempsey, 2017]. As space exploration expands further away from LEO and the Moon, the delay in communication represent an insurmountable obstacle for remote guidance, control and communications. Characteristically, we note that the round-trip time for current communications with Mars ranges between 5 and 20 minutes depending on the state of the two planets orbits [NASA, 2020]. ML, thus, can resolve the dependency on communications and perform the mission critical information processing aboard. (2) Adapting Maintenance - Modern spacecrafts both manned and unmanned, constitute extremely complex systems that even with the use of automated checks are still prone to faults, especially when faced with extraordinary circumstances. The crew might not be able to solve the issues by themselves. To tackle this we believe that ML systems can perform functions that transcend anomaly detection and fault prediction. An ideal deployed ML solution needs to account for automated maintenance and resolution of faults both in hardware and in the software of the spacecraft. (3) Reduce latency for functions that don't require manual intervention and can be conducted at scale. For example, automated checksums for data or models (4) Recent advances in ML like bit quantization, pruning, and hardware approximations [Wang *et al.*, 2019] enable inference on resource constraint edge devices which can also be similar to the target hardware in space [Hiemstra *et al.*, 2020].

4 Medical Use Case

In the previous sections we have set out some challenges as well as advantages for the use of ML on spacecrafts. In this section we will be exploring the medical use case. As mentioned before, medical treatment in space relies on communication with the earth, a reliance we need to remove as we progress into deep space missions. A major part of medical treatment is the ability to inspect bodily functions in a non invasive manner through the use of medical imaging devices. Magnetic Resonance (MRI), CT-Scans and X-Rays are widely used on earth but are not ideal for space use, as we explain below. To this effect, we propose the use of Point of Care Ultrasounds (POCUS) as a lightweight non ionizing imaging solution.

ML Enabled Imaging: Medical imaging devices are often slow, requiring significant resources to operate and store and depend on the use of ionizing radiation (e.g. X-Rays/CT Scans) that has negative effects on patients and doctors alike [FDA, 2021]. One notable exception to the above constraints is Point of Care Ultrasounds (POCUS). Mobile device enabled Ultrasound (US) probes are safe for patients and doctors, require minimal resources to operate — weighting less than 1kg— and offer real time imaging capabilities [GEHealthcare, 2021]. However, acquiring medically significant images with a POCUS probe is non trivial and requires expert operators.

As expert users might not always be available in deep space missions we envision the use of ML enabled POCUS probes

that guide the user to medically relevant areas.

Approaching the task of navigation one can identify straight away the need to be able to determine where the intelligent agent is with respect to the patients body. Basic anatomy knowledge by both patient and other crew members is assumed for the purposes of this application. Navigating to points of medical interest, though is non trivial. Following medical practice of first identifying standard planes and anatomical landmarks; reinforcement learning (RL) solutions of disembodied agents have been proposed in [Alansary *et al.*, 2019; Vlontzos *et al.*, 2019] to perform the above tasks, providing high accuracy and low computational constraints. [Milletari *et al.*, 2019; Hase *et al.*, 2020; Li *et al.*, 2021] expand the RL methods including the degrees of freedom of the US probe into the agents action space and optimize directly on identifying landmarks while controlling a virtual probe. All the above methods, find themselves limited to the anatomies that they have been trained on, with their computational burden rising exponentially when trained to find standard planes or landmarks of multiple anatomies.

As such we believe that a hierarchical approach to the problem would help keep complexity and computation low enough for space craft applications while maintaining high enough performance. Approaches like MAX-Q [Dietterich, 2000] have withstood the test of time and have theoretical guarantees on convergence and the ability of the algorithm that given a set of subtasks it can find the global optimal policy. In short, MAX-Q constructs a hierarchical action tree that a higher level agent sets out subtasks for other agents. In this fashion one can design a fully autonomous agent that collects information in form of images through POCUS navigation, assesses biomarkers based on measurements derived from landmarks, and regresses the diagnosis.

Furthermore, [Eslami *et al.*, 2018; Dinesh and Grauman, 2016; Ramakrishnan and Grauman, 2018] have shown that incorporating modules that aid agents "imagine" how scenes would look like from different point of views, increases their performance in scene understanding. Hence, we are firm believers that incorporating such modules in approaches like [Milletari *et al.*, 2019] would increase their performance capabilities.

Finally, as medical applications are of critical importance, appropriate checks and balances should be put in place to avoid any harm to the crew. We believe that a rule based set of parameters should be developed in collaboration with physicians that would constitute a fallback system. Having no dependence on learned or inferred information we are able to guarantee a basic level of care to the astronauts in case all other systems fail. On top of the safety-net rule based system, a causal inference infrastructure can be used to assess the the probabilities of causality (Sufficiency and Necessity) [Buesing *et al.*, 2018; Kusner *et al.*, 2017; Louizos *et al.*, 2017; Vlontzos *et al.*, 2020; Budd *et al.*, 2021]. Causal counterfactual inference enable us to assess the causal links and potential outcomes of treatments providing, thus, a more informed decision process.

5 Infrastructure considerations for ML systems on Spacecraft

Developing novel ML methods to provide mission critical treatments to astronauts is a hard task. In the previous section we set out some thoughts on a potential application for medical imaging aboard a spacecraft. In this section we will be briefly exploring infrastructure considerations that are directly related to the operation and development of medical imaging ML solutions. These considerations are also apparent on earth bound ML applications but gain increased importance due to the edge cases that they are called to operate on during spaceflight.

Data Collection Medical data acquisition is faced with a series of challenges. The privacy of the patients donating their data has to be protected throughout the data acquisition and model development processes. In order to comply with all legal obligations, we propose the data to have NIHR [NIHR, 2021] approved methods in place for making patient data anonymous. While, storage of data must be in a HIPAA [HIPAA, 2021] compliant platform. Federated learning also present promising methods in regards to secure data processing [Konečný *et al.*, 2016]. Going beyond the need of anonymization and maintenance we would like to draw the readers attention to two constraints of increased importance on medical ML applications

Domain Shift: Domain shift robustness is an open topic of ML research that focuses on making existing methods robust to distributional shifts from the underlying training data domain [Long *et al.*, 2016; Zhang *et al.*, 2020]. In medical applications and more prominently in medical imaging applications domain shifts are easier to manifest and harder to overcome. The first major factor in this phenomenon comes from the equipment used. Medical equipment, when installed in hospitals and health centers is configured by the seller to the exact specifications of the attending physicians, as such, two different doctors using the same base equipment on the same patient can result in two quite different images. The standardization of medical equipment configurations in space missions would aid decrease the equipment induced domain shifts.

Another domain shift source comes from the training patient characteristics. Different populations exhibit different medical characteristics. Hereditary traits as well as phenotype derived attributes place an individual in different medical risk groups and force pathologies to manifest with different probabilities. These differences, are often picked up by ML algorithms as unwanted inductive biases, skewing the learned conclusions. In the context of space missions, astronauts have diverse backgrounds both genetically and in terms of phenotype, as such models trained on a general population not representative of the characteristics of the crew can provide skewed estimations of vital for diagnosis biomarkers. It is imperative, then, that the deployed models are calibrated to the genetic background and phenotype of the astronauts.

Anomalous phenomena: As stated in the above motivation, the effects on human well-being stemming from prolonged exposure to cosmic radiation and other space related phenomena are not fully understood. It is unknown, hence how the

human body might react to adverse conditions. ML applications cannot, then, be expected to cover the full range of scenarios, on the contrary they should be expected to fail when presented with data that constitute an anomalous effect. In order to combat potential failure cases of the medical systems on-board a spacecraft, we suggest accompanying any ML algorithm with an anomaly detection mechanism (perhaps probabilistic or out-of-distribution mechanics), that flags non standard bio markers [Winkens *et al.*, 2020], and an active learning [Settles, 2012] feedback pipeline such that new, medical phenomena are incorporated in the evolution of the medical algorithms on-board.

Computational Considerations High communication latency creates the demand for on-board computation. However, cosmic radiation and weight offers strict constraints on the available on-board compute resources. Cosmic radiation is able to produce errors in modern scale computing units, while weight restrictions exist in all parts of a mission, from lift off to grounding. ML applications should, then, be able to function without heavy computational needs, a non trivial feat especially on medical imaging algorithms. This, reinforces our case towards the use of POCUS as they have been shown to be computationally and physically lightweight. Recently [Hiemstra *et al.*, 2020] has shown that NVIDIA’s Jetson Xavier modules are able to withstand a significant level of proton based radiation, making them optimal candidates for on-board inference and fine-tuning infrastructure. NASA has also awarded an SBIR contract to Numem [Gharia, 2020] to develop a radiation hardened DNN co-processor for a wide variety of ML applications, from low power machine vision to healthcare.

6 Conclusion

ML supported medical diagnostic systems on spacecraft are necessary for long-term space mission to overcome limitations in ground-to-spacecraft communication and lack of qualified medical crew. In this extended abstract, we outlined the importance of incorporating ML enabled medical applications on spacecrafts and considered challenges that need to be overcome in terms of anomaly detection and domain adaptation. Finally, we discussed ideas on how to augment existing POCUS algorithms such that they constitute a complete diagnostic system. We gave the example of a hierarchical RL method that also includes causal inference checks and balances such that the health and safety of crew members is guaranteed. We hope that with this paper we highlighted the final frontier for ML applications: space capable ML medical systems.

7 Acknowledgments

We would like to thank NASA, the Canadian Space Agency (CSA), European Space Agency (ESA), NVIDIA, the SETI Institute and Frontier Development Lab (FDL) for their continuous support and for making this collaboration possible.

References

[Alansary *et al.*, 2019] Amir Alansary, Ozan Oktay, Yuanwei Li, Loic Le Folgoc, Benjamin Hou, Ghislain Vail-

- lant, Konstantinos Kamnitsas, Athanasios Vlontzos, Ben Glocker, Bernhard Kainz, et al. Evaluating reinforcement learning agents for anatomical landmark detection. *Medical image analysis*, 53:156–164, 2019.
- [Ayaz *et al.*, 2019] Muhammad Ayaz, Mohammad Ammad-Uddin, Zubair Sharif, Ali Mansour, and El-Hadi M. Aggoune. Internet-of-things (iot)-based smart agriculture: Toward making the fields talk. *IEEE Access*, 7:129551–129583, 2019.
- [Budd *et al.*, 2021] Samuel Budd, Arno Blaas, Adrienne Hoarfrost, Kia Khezeli, Krittika D’Silva, Frank Soboczenski, Graham Mackintosh, Nicholas Chia, and John Kalantari. Prototyping crisp: A causal relation and inference search platform applied to colorectal cancer data. In *2021 IEEE 3rd Global Conference on Life Sciences and Technologies (LifeTech)*, pages 517–521. IEEE, 2021.
- [Buesing *et al.*, 2018] Lars Buesing, Theophane Weber, Yori Zwols, Sebastien Racaniere, Arthur Guez, Jean-Baptiste Lespiau, and Nicolas Heess. Woulda, coulda, shoulda: Counterfactually-guided policy search. *arXiv preprint arXiv:1811.06272*, 2018.
- [Dempsey, 2017] Robert C. Dempsey. *The International Space Station: operating an outpost in the new frontier*. National Aeronautics and Space Administration, Lyndon B. Johnson Space Center, 2017.
- [Dietterich, 2000] Thomas G Dietterich. Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of artificial intelligence research*, 13:227–303, 2000.
- [Dinesh and Grauman, 2016] Jayaraman Dinesh and Kristen Grauman. Look-Ahead Before You Leap : End-to-End Active Recognition. *ECCV*, 2016.
- [Eslami *et al.*, 2018] S. M. Ali Eslami, Danilo Jimenez Rezende, Frederic Besse, Fabio Viola, Ari S. Morcos, Marta Garnelo, Avraham Ruderman, Andrei A. Rusu, Ivo Danihelka, Karol Gregor, David P. Reichert, Lars Buesing, Theophane Weber, Oriol Vinyals, Dan Rosenbaum, Neil Rabinowitz, Helen King, Chloe Hillier, Matt Botvinick, Daan Wierstra, Koray Kavukcuoglu, and Demis Hassabis. Neural scene representation and rendering. *Science*, 360(6394):1204–1210, 2018.
- [FDA, 2021] FDA. Medical x-ray imaging. <https://www.fda.gov/radiation-emitting-products/medical-imaging/medical-x-ray-imaging>, 2021.
- [GEHealthcare, 2021] GEHealthcare. Point of care ultrasounds. <https://www.gehealthcare.co.uk/products/ultrasound/point-of-care-ultrasound>, 2021.
- [Gharia, 2020] Nilesh Gharia. Dnn radiation hardened co-processor companion chip to nasa’s upcoming high-performance spaceflight computing processor. <https://www.sbir.gov/node/1882269>, 2020.
- [Hase *et al.*, 2020] Hannes Hase, Mohammad Farid Azampour, Maria Tirindelli, Magdalini Paschali, Walter Simson, Emad Fatemizadeh, and Nassir Navab. Ultrasound-guided robotic navigation with deep reinforcement learning, 2020.
- [Hiemstra *et al.*, 2020] D. M. Hiemstra, C. Jin, Z. Li, R. Chen, S. Shi, and L. Chen. Single event effect evaluation of the jetson agx xavier module using proton irradiation. In *2020 IEEE Radiation Effects Data Workshop (in conjunction with 2020 NSREC)*, pages 1–4, 2020.
- [HIPPA, 2021] HIPPA. Hippa data protection guidelines. <https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html>, 2021.
- [Konečný *et al.*, 2016] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- [Kusner *et al.*, 2017] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. In *Advances in Neural Information Processing Systems*, pages 4066–4076, 2017.
- [Lavin *et al.*, 2021] Alexander Lavin, Ciarán M. Gilligan-Lee, Alessya Visnjic, Siddha Ganju, Dava Newman, Sujoy Ganguly, Danny Lange, Atılım Güneş Baydin, Amit Sharma, Adam Gibson, Yarin Gal, Eric P. Xing, Chris Mattmann, and James Parr. Technology readiness levels for machine learning systems, 2021.
- [Li *et al.*, 2021] Keyu Li, Jian Wang, Yangxin Xu, Hao Qin, Dongsheng Liu, Li Liu, and Max Q. H. Meng. Autonomous navigation of an ultrasound probe towards standard scan planes with deep reinforcement learning, 2021.
- [Long *et al.*, 2016] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. *arXiv preprint arXiv:1602.04433*, 2016.
- [Louizos *et al.*, 2017] Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *Advances in Neural Information Processing Systems*, pages 6446–6456, 2017.
- [Milletari *et al.*, 2019] Fausto Milletari, Vighnesh Birodkar, and Michal Sofka. Straight to the point: reinforcement learning for user guidance in ultrasound, 2019.
- [NASA, 2017] NASA. Nasa artemis program. <https://www.nasa.gov/artemisprogram>, 2017.
- [NASA, 2020] NASA. Nasa mars 2020 mission perserverance rover. <http://shorturl.at/xzi25>, 2020.
- [NIHR, 2021] NIHR. Data sharing guidelines. <https://www.nihr.ac.uk/documents/nihr-position-on-the-sharing-of-research-data/12253>, 2021.
- [Page and Benedetto, 2005] T.E. Page and J.M. Benedetto. Extreme latchup susceptibility in modern commercial-off-the-shelf (cots) monolithic 1m and 4m cmos static random-access memory (sram) devices. In *IEEE Radiation Effects Data Workshop, 2005.*, pages 1–7, 2005.

- [Ramakrishnan and Grauman, 2018] S K Ramakrishnan and K Grauman. Sidekick Policy Learning for Active Visual Exploration. In *EECV*, 2018.
- [Settles, 2012] Burr Settles. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1):1–114, 2012.
- [Vlontzos *et al.*, 2019] Athanasios Vlontzos, Amir Alansary, Konstantinos Kamnitsas, Daniel Rueckert, and Bernhard Kainz. Multiple landmark detection using multi-agent reinforcement learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 262–270. Springer, Cham, 2019.
- [Vlontzos *et al.*, 2020] Athanasios Vlontzos, Henrique Bergallo Rocha, Daniel Rueckert, and Bernhard Kainz. Causal future prediction in a minkowski space-time. *arXiv preprint arXiv:2008.09154*, 2020.
- [Wang *et al.*, 2019] Erwei Wang, James J. Davis, Ruizhe Zhao, Ho-Cheung Ng, Xinyu Niu, Wayne Luk, Peter Y. K. Cheung, and George A. Constantinides. Deep neural network approximation for custom hardware. *ACM Computing Surveys*, 52(2):1–39, May 2019.
- [Winkens *et al.*, 2020] Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R. Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, Taylan Cemgil, S. M. Ali Eslami, and Olaf Ronneberger. Contrastive training for improved out-of-distribution detection, 2020.
- [Yang *et al.*, 2019] Hui Yang, Soundar Kumara, Satish Bukkapatnam, and Fugee Tsung. The internet of things for smart manufacturing: A review. *IIE Transactions*, pages 1–35, 01 2019.
- [Zhang *et al.*, 2020] Yifan Zhang, Ying Wei, Qingyao Wu, Peilin Zhao, Shuaicheng Niu, Junzhou Huang, and Mingkui Tan. Collaborative unsupervised domain adaptation for medical image diagnosis. *IEEE Transactions on Image Processing*, 29:7834–7844, 2020.