Supporting responsible machine learning in heliophysics

Ayris Narock[1,2]*, Christopher Bard[1], Barbara J. Thompson[1], Alexa J. Halford[1], Ryan M. McGranaghan[1,3], Daniel da Silva[1,4], Burcu Kosar[1,5] and Mykhaylo Shumko[1,6]

[1]NASA Goddard Space Flight Center, Greenbelt, MD, United States, [2]ADNET Systems Inc, Bethesda, MD, United States, [3]Orion Space Solutions, Louisville, CO, United States, [4]University of MD, Baltimore County, Baltimore, MD, United States, [5]The Catholic University of America, Washington D.C, DC, United States, [6]University of Maryland at College Park, College Park, MD, United States

Over the last decade, Heliophysics researchers have increasingly adopted a variety of machine learning methods such as artificial neural networks, decision trees, and clustering algorithms into their workflow. Adoption of these advanced data science methods had quickly outpaced institutional response, but many professional organizations such as the European Commission, the National Aeronautics and Space Administration (NASA), and the American Geophysical Union have now issued (or will soon issue) standards for artificial intelligence and machine learning that will impact scientific research. These standards add further (necessary) burdens on the individual researcher who must now prepare the public release of data and code in addition to traditional paper writing. Support for these is not reflected in the current state of institutional support, community practices, or governance systems. We examine here some of these principles and how our institutions and community can promote their successful adoption within the Heliophysics discipline.

1 Introduction

The use of machine learning (ML) and other advanced analytics methods in Heliophysics has grown steadily in recent years and will continue to do so. As their use in both research and operations become more prevalent, it is imperative that the community adopt a conscious effort to use these, often black-box, methods in a manner that allows confidence in the interpretation of the results and meets established criteria for use. Many professional organizations are developing and promoting community standards specific to ML and artificial intelligence (AI). Such documents include values around the design, implementation, and use of AI systems and commonly title this the "ethics" of AI (European Commission, 2019; Mclarney et al., 2021; AGU - NASA Ethics in AI/ML, 2022; U.S. Department of Defense, 2020; The U.S. Intelligence Community, 2020b; The U.S. Intelligence Community, 2020a) although the scope is

broader than what is typically considered ethics within science. While they have yet to be widely adopted, once they are fully codified and required, the responsible use of machine learning will mean meeting these guidelines and will add a burden to the researcher. These are not reflected in the current state of institutional support, especially in funded research grants. Furthermore, community awareness of these policies needs to be raised and best practices set. Here we examine some of the implications of these emerging mandates and how our institutions and organizations can support successful adoption of responsible ML practices within the Heliophysics community. A version of this paper (Narock et al., 2022) was originally written in response to the 2024 Heliophysics decadal survey run by the National Academies in the United States. Previous decadal surveys and other National Academy reports have been used by many space agencies within the United States to help form their strategic plans for the coming decade (National Research Council, 2013; NASA, 2015; National Academies of Sciences, Engineering, and Medicine, 2017). Thus, this paper has a more United States centric view with recommendations intended to call attention to areas on which United States institutions might want to focus to help relieve burdens associated with responsible AI. However, we feel that many of these recommendations are applicable to other institutions as well. Additionally, recommendations not targeted specifically at institutions are broadly appliable to the entire Heliophysics community.

2 Machine learning in heliophysics

While traditional analysis for research and operations has long been reliant on heavy use of computing, in recent years there has been a striking increase towards using ML methods. As generally defined, AI is the effort to automate intellectual tasks normally performed by humans (Chollet, 2018). Any method which allows a computer to learn without explicitly being programmed is called ML, and is a subfield of AI. With these methods, the relative importance of aspects of the data are not specified by any human directly but inferred and codified by the analytical learning process itself. A recent survey of Heliophysics literature (Camporeale, 2019) shows that the percent of published works in the field that discuss ML has been growing with increasing frequency. Azari et al. (2021) extends the surveyed time range and finds an exponential increase within Heliophysics and across the wider range of related physical science domains. That a sustained sub-discipline is beginning to be carved out is further evidenced by the creation of the international "Machine Learning in Heliophysics" conference (Camporeale, 2022) as well as dedicated sessions within other long-standing Space Physics conferences such as CEDAR (Romick et al., 1987; McGranaghan et al., 2022, 2017),

GEM (Roederer, 1988; GEM, 2022), SHINE (Crooker, 2002; SHINE, 2022), AMS, (2022) and AGU (Camporeale et al., 2022).

This widespread adoption of ML is not specific to Heliophysics. Because of the unique characteristics involved, many institutions are beginning to delineate specific and separate expectations for the use and development of machine learning applications. The European Commission has released the high-level "Ethics guidelines for trustworthy AI" in 2018 (European Commission, 2019) and the more actionable "Assessment List for Trustworthy Artificial Intelligence (ALTAI)" in 2020 (European Commision, 2021), while in the United States the "Blueprint for an AI Bill of Rights" was released in 2022 (White House Office of Science and Technology Policy, 2022). NASA formulated its early expectations in 2021 with the technical memorandum, "NASA Framework for the Ethical Use of Artificial Intelligence (AI)", describing the principles to guide AI work and how each applies to NASA work (Mclarney et al., 2021). The NASA Science Mission Directorate has its own focused AI Initiative (Bolles, 2022a) and in partnership with AGU is in the process of developing a community-derived set of AI/ML standards (AGU—NASA Ethics in AI/ML, 2022).

Furthermore, federal agencies such as NASA, NSF and NOAA fall under the 2020 "Executive Order on Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government" (Trump, 2020); the policies and frameworks being issued by United States funding agencies and governmental research programs will continue to evolve in order to align with required directives. Our community needs to remain an engaged contributor in the unfolding conversation about AI regulation. However, an examination of infrastructure and support to enable compliance with current trends is prudent.

3 Principles of ethical artificial intelligence

An ethical AI system is one that acts as intended and adheres to defined ethical guidelines that reflect the fundamental values of an organization or entity. Responsible AI is complementary; the steps taken to actually enable the design and implementation of an ethical AI system. The guidance documents cited in Section 2 are examples outlining the fundamental values of each entity, with respect to AI. The Heliophysics community, in order to responsibly deploy ML methods, will therefore need to be aware of and carry out effective practices that produces work in alignment with the desired principles.

While the various guidance documents have subtle differences, they tend to focus on similar aspects. As an example, Table 1 summarizes the six high-level principles outlined by NASA (Mclarney et al., 2021): Fair, Explainable and Transparent, Accountable, Secure and Safe, Human-Centric and Societally

TABLE 1   Principles from *NASA Framework for the Ethical Use of Artificial Intelligence (AI)*.

Fair	AI systems must include considerations of how to treat people, including scrubbing solutions to mitigate discrimination and bias, preventing covert manipulation, and supporting diversity and inclusion.
Explainable & Transparent	Solutions must clearly state if, when, and how an AI system is involved, and AI logic and decisions must be explainable. AI solutions must protect intellectual property and include risk management in their construction and use. AI systems must be documented.
Accountable	Organizations and individuals must be accountable for the systems they create, and organizations must implement AI governance structures to provide oversight.
Secure & Safe	AI systems must respect privacy and do no harm. Humans must monitor and guide machine learning processes. AI system risk tradeoffs must be considered when determining benefit of use.
Human-Centric & Societally Beneficial	AI systems must obey human legal systems and must provide benefits to society. At the current state of AI humans must remain in charge, though future advancements may cause reconsideration of this requirement.
Scientifically & Technically Robust	AI systems must adhere to the scientific method NASA applies to all problems, be informed by scientific theory and data, robustly tested in implementation, well-documented, and peer reviewed in the scientific community.

Beneficial, and Scientifically and Technically Robust. Applied to the discipline of Heliophysics, some of these principles hold more relevance than others.

For instance, "Fair" would have high impact if developing a recommender system for selecting potential employees but seem to hold little significance when ML training data is not human-based. However, this principle includes the concept of mitigating data bias, which is an important aspect for all ML projects. A system to predict solar flares would have limited usefulness if a skewed data set led it to always predict a quiet Sun. There are also efforts underway to automate knowledge capture and inference within the field (connecting data, bodies of work, and researchers) that will need to be carefully engineered to avoid propagating existing systemic biases to a new medium.

The principle of "Explainable and Transparent" is also broadly applicable to most ML problems. While there is no definitive answer for what it means for an AI application to be explainable and transparent, the authors feel that this must include the ability to describe at least some aspects of why a model returns the results it does (for example, which features of the data are most impactful to a classification.) It is an easy pitfall for a researcher to apply an ML method and report results, with little to no effort to illuminate the AI logic and decisions.

As our missions become more advanced, with longer data latency or blackout periods, onboard decision making may take advantage of the currently maturing research software trained from historical data. Without well-defined oversight mechanisms

throughout the lifetime of this software it will be difficult for the process to be "Accountable" at maturity.

The fourth principle of "Secure and Safe" seems the most removed from direct implications to the Heliophysics community. Ensuring that AI systems do no harm and considering risk tradeoffs are considerations more geared towards developing technologies such as autonomous vehicles.

The ideas behind "Human-Centric and Societally Beneficial" are relevant to much work that focuses on forecasting space weather events. As ML methods mature and are considered for adoption into decision systems governing spacecraft or astronaut health, it is important to be able to trust and validate these methods.

Finally, to be "Scientifically and Technically Robust" is another broad-reaching principle. An ML approach could run afoul of this guideline in several ways. First, by blindly applying a method without full understanding of its underlying theory; many ML methods are meaningful only when certain statistical conditions within their training data are met. Second, by not meeting the criteria for robust testing due to limitations preventing others from reproducing the work. Finally, by not having the appropriate resources in place to have the method adequately peer reviewed.

Of the five principles above that are most relevant to Heliophysics, one centers on establishing governance. The other four relate to ideas surrounding to the ability to trust, validate, understand, and replicate methods; achievable only by having direct access to the trained model.

These principles, on the surface, express the same core values desired in most general codes of ethics related to scientific work and software development Gotterbarn et al. (1997). The leadership for these entities, however, have deemed that the distinction between traditional software methods and AI is important enough to define ethics within this context beyond what is already in place. Compared to traditional methods, ML methods are inherently more black-box and pose novel challenges to meeting these goals. In traditional computational algorithms the step-by-step internal logic and decision making can be concretely described, examined, and in most cases evaluated to ensure that the system is meeting specifications (Vogelsang and Borg, 2019; Hutchinson et al., 2021; Serban et al., 2021). When decision points are created by an algorithm iterating over large amounts of input information, as with ML, the ability to understand the resulting model is obscured.

4 Requirements for responsible use

Deliberately building in components to a model to make it more interpretable by design and performing post-hoc analysis on a trained model to explain its decisions, are useful approaches to clarify understanding of ML models. While individual developers can and should do their best to investigate their own methods, to make this robust the community must be able to perform their own, independent inquiries. Therefore, adhering to the principles of the emerging Open Science initiatives, like NASA's Transform to Open Science (TOPS) (Bolles, 2022b), will help address many aspects of the emerging AI guidelines by enabling outside verification and evaluation. This move towards open science is in no way specific only to ML, but we are focusing only on its implications to ML because of the AI/ML specific rules now in development. Exercises to explain or quantify predictive behavior can be undertaken when the trained machine learning model is publicly released. Providing open-source code and sharing the model training datasets can build trust in the method and allow confirmation that the data is unbiased or that appropriate efforts to mitigate bias have been undertaken. Providing all the above in a containerized environment creates a strong basis for reproducible work (The Turing Way Community, 2022).

Undertaking this approach to ensure transparency to the community adds a necessary increase in both researcher workload and infrastructural needs. No longer are high level descriptions of methodology sufficient, and a working version of research code is rarely releasable without specific intention to do so. It takes time and effort to make code and datasets releasable if we intend for them to actually be useful to an independent investigator. Clear documentation and usage instructions need to be in place, as well as explanations

of working assumptions and known and unknown biases (Lee, 2018; Serban et al., 2021). Depending on the products used and one's experience, containerization could encompass a steep learning curve. Ensuring replicable work across different computing platforms requires researchers and independent investigators have access to hardware of various architectures and performance levels. While not exhaustive, this list provides a flavor for the type of efforts and systems required for more transparency. As an illustrative example of the distance we have ahead of us to meet transparency goals, a search of the *Journal of Geophysical Research: Space Physics* between October 2021 and September 2022 returns 22 articles that mention "machine learning" (American Geophysical Union, 2022). One of these articles was a review type and wouldn't be expected to provide software. Of the remaining 21 articles, only five include public access to software or source code used to generate the paper results.

For this to be accomplished, sustained and coordinated support from affiliated agencies is needed. The Heliophysics community needs clear and straightforward pathways in place for all members. Currently, there is a high discrepancy between allowed and accessible code sharing methods between institutions. While many adopt industry-wide standard platforms, they are often deployed in a protected and non-public environment and may incur additional costs. These barriers make software collaboration across institutions difficult. In some institutions, the default posture is that all source code is protected, and special approval must be granted to share it publicly or even with collaborators outside the organization. In many cases, such approval processes are obscure or unreasonably burdensome when applied to scientific software.

Additionally, the existing level of review and governance of AI/ML work within Heliophysics is ambiguous. There are no generally accepted standards for how to go about ensuring adherence to best practices. Journals, already strapped for reviewers, find it difficult to cover the burgeoning number of ML papers with someone knowledgeable enough in both the science and computer science to properly vet the work. Once work has been done and published, possibly with the software released, the expectations for maintenance are unclear.

5 Discussion and suggested recommendations

Mandating guidelines is not sufficient for creating an open and responsible research community. Environments and culture must be cultivated to make adoption of such mandates feasible. As a perspective on a path forward, here we offer several recommendations for institutions and the community at large to promote and encourage the adoption of responsible machine learning research practices. These are by design general and not

meant to be proscriptive but instead to be touchpoints in strategic planning discussions.

5.1 Streamlining open science and establishing preferred platforms

As many institutions are now implementing open science requirements and will soon enforce open release of data and source codes (Burrell et al., 2018; Chen et al., 2018; National Academies of Sciences, Engineering, and Medicine, 2018; Barnes et al., 2020; Bolles, 2022b; European Commission, 2022), there must be a recognition that institutions should also facilitate and streamline this process. This includes funding for the time and effort to make them suitable for the general community to use, providing or funding platforms for such release, and streamlining and supporting the process for open release without difficult bureaucracy. Although some data and codes must necessarily be limited due to trafficking and export control restrictions, this should not be the assumption by default for scientific analysis work. In fact, the default should be that data and codes are open, with government and other entities working to ensure code development and data/data processing are open processes for all. Scientific codes should have a streamlined process separate from engineering-related or operations codes with minimal bureaucratic overhead.

Consideration for the initial public release effort and an early maintenance period should also be included in grant funding for scientific research. In particular, costs and permissions associated with hosting platforms should be an expected part of doing business, similar to institution-provided phone or internet service. Institutions should provide and pay for general access to a variety of computing systems and hosting services, committing to and streamlining public release of scientific software used in published research.

To aid institutions in formulating the most effective support to fund these open practices, a limited number of platforms should be identified by the community as preferred. These platforms can be used for code sharing and collaboration as well as host citable archives for released code and ML-ready datasets. Currently, popular platforms include GitHub and GitLab (for sharing and collaboration), and Zenodo (for data hosting). Care should be taken to ensure that archives be hosted and stored on public servers and not with privately-owned companies and that access is free to the user and contributors.

As part of open release of code, there should be processes to aid in the reproduction of results. Since many libraries in data science and ML are open source, they are frequently updated; therefore, separate platforms may have different combinations of library versions even if running the same code. One way to standardize this is to use containers, such as Docker or Podman. Compatibility, portability, and supportability should all be considered when recommending preferred containers (Anand, 2021).

Once recommendations for platforms are established, the community and institutions can facilitate training that is targeted to these tools and within the typical workflow of this community. Support for these platforms should be managed as a public utility and supported at the institution level, not from the individual projects.

5.2 Supporting methodological work in addition to novel research

While novel scientific results published in a peer reviewed journal is certainly an attainable goal of many investigations using these modern methods, equally worthwhile are works centered on examining new and varied methods applied to previously examined science questions. As the Heliophysics community works to build trust in and understanding of how and where ML methods are most appropriately used, it is imperative that these sorts of applications happen frequently. Additionally, when ML methods are applied to a problem and do not perform well or when we determine whether we are able to replicate results in a different environment or with a similar but distinct dataset, this is valuable information that can be leveraged for future breakthrough science. See Open Science Collaboration (2022) for a deeper discussion on replicability.

At present, much of this work is dubiously labeled original research when published, or it is rejected entirely due to lack of "novelty". Thus, we need to build value around this distinct class of paper, perhaps called a technical report, that adds depth and breadth to our understanding of advanced analytics as applied to Heliophysics without necessarily advancing the science of Heliophysics itself. This may include negative results or failures to reproduce previously-published results. Building this value will require a cultural change away from the normative behavior in science that raises personal recognition above all else.

Technical reports also serve an important role in establishing trust and dependability for machine learning methods. This is especially important as new data science techniques, structures, and algorithms are developed and iterated upon. If multiple ML methods applied to identical data produce (nearly) identical results, then this establishes the trustworthiness and robustness of the original result. There are multitudinous combinations of possible ML architectures and methodologies which can be applied to Heliophysics problems, but there are not nearly as many reviewers. Technical reports provide an outlet for this work.

Scientific journals should therefore support peer-reviewed "Methods" or "Technical Reports" articles if they do not already. Furthermore, researchers should publish their work, including the publication of negative results, under its most

appropriate article type. To promote the cultural acceptance and adoption of technical reports, funding agencies should offer opportunities to perform replicability studies and/or "re-testing" studies in which alternative methods are employed for previously published problems. Finally, some results may require significant computing time for recalculation. Institutions may wish to provide computational resources to aid in this.

5.3 Envisioning and implementing governance

Most science communities overlook the importance of governance, so education around the subject is needed. To develop a literacy for governance and the role it will play in science, the Heliophysics community should interface with existing AI/ML associations outside of Heliophysics to learn about AI/ML governance strategies. Communities such as the Stanford Artificial Intelligence Laboratory (2022), the Institute for the Future (2022), and The Alan Turing Institute (2022) have already made significant progress along this path (Cline et al., 2022) and we should learn from their lead about how to use these technologies responsibly. This may include lessons on expanding the covered domains of research teams and educational programs; how models are best maintained; setting up a framework for community evaluation of models; as well as best practices for journals that encourage responsible use and provide standards for robust peer review of publications that use AI/ML.

As a first step, Heliophysics researchers can join these communities. Heliophysics will not solve these problems on our own, but we can better connect to the groups creating the solutions and develop relationships whereby we can tailor those general solutions to the Heliophysics context. The Heliophysics community should also examine sets of principles for governance, including "Generalizing the Core Design Principles for the Efficacy of Groups" (Wilson et al., 2013), the "CSCCE Community Participation Model" (Center for Scientific Collaboration and Community Engagement, 2022), and pages 21–29 of "The Reproducibility Project" (Open Science Collaboration, 2022).

Only after we have done the work of connecting and examining shall we come to the question of implementation. This paper will not provide an answer on that front, but strongly urges a community driven approach to determining governance goals and implementation strategies.

6 Conclusion

Machine learning and other advanced data analysis techniques promise great advancements in scientific understanding. However, it is essential that one works to explain how and why and does not simply trust a black-box model output. Additionally, the global community has realized that such methods can have significant effects, possibly deleterious, on broader society; they must be responsibly used.

Towards these ends, there has been a great push towards making both data and scientific code FAIR [findable, accessible, interoperable, reusable; Wilkinson et al. (2016)] and open source, as well as defining specific expectations for the use of AI and ML [e.g. Mclarney et al. (2021)]. Although the societal impacts from research code are much less than from other sectors employing ML, the scientific community nonetheless is subject to the expanding sets of regulations surrounding its use. However, simply mandating that researchers be "ethical" and "open" is not sufficient. A public community built around responsible use of ML needs institutional and structural support beyond simply receiving research funds.

We have detailed here several additional burdens that researchers face when adhering to new ethical principles:

- Scientific products require additional time and effort to be made publicly accessible and useable.
- Bureaucratic issues and other lacks of institutional support impede, delay, and disincentivize public release.
- The emphasis on "novelty" in scientific journals disincentivizes ensuring reproducibility in scientific results.
- This emphasis also hinders experimentation with using new architectures to investigate previously-answered problems, making it difficult to establish the trustworthiness and robustness of ML methods.

These burdens are not insurmountable, but research-funding institutions must be cognizant that grants should also provide support for time and effort to satisfy ethical and open science requirements. In addition, institutions must also provide structural and policy support to streamline these efforts and encourage adoption. Some of our recommendations to accomplish these goals include:

- Establish community-wide, preferred platforms for sharing data and scientific code. Fund them as basic utilities and give scientists easy access.
- Add additional funding to grants to cover the effort required to follow newly-mandated open science requirements.
- Provide computing resources and fund projects to aid in reproducing results from published papers, possibly using alternative architectures and methods. The resulting "technical reports" will help to establish reliability and trust of ML-derived scientific results.
- Learn governance strategies and best practices from existing AI/ML communities; tweak as needed for Heliophysics applications.

The Future of Space Physics will rely on these advanced analytical methods and will need to live in the space constrained by evolving standards. At all levels we need to remain aware and forward thinking to put into place the infrastructure and culture needed for success.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

AN conceptualized the idea and led writing efforts. CB, AH, BT, and RM contributed significant writing, editing, and revising. All authors revised the manuscript before submission.

Funding

Conflict of interest

Author AN was employed by ADNET Systems Inc. RM was employed by Orion Space Solutions.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

References

AGU - NASA Ethics in AI/ML (2022). Agu - NASA ethics in AI/M. Available at: https://data.agu.org/ethics-ai-ml/ (Accessed Aug 19, 2022).

American Geophysical Union (2022). Search results from journal of geophysical research: Space Physics. Available at: https://agupubs.onlinelibrary.wiley.com/action/doSearch?AfterMonth=10&AfterYear=2021&BeforeMonth=9&BeforeYear=2022&Ppub=&field1=AllField&publication%5B%5D=21699402&rel=nofollow&sortBy=Earliest&startPage=0&text1=%22machine+learning%22&utm_campaign=R3MR425&utm_content=EarthSpaceEnvirSci&utm_medium=paidsearch&utm_source=google&pageSize=20 (Accessed Nov 02, 2022).

AMS (2022). AMS 19th conference on space weather. Available at: https://annual.ametsoc.org/index.cfm/2022/program-events/conferences-and-symposia/19th-conference-on-space-weather/ (Accessed Aug 22, 2022).

Anand, V. (2021). Containers interoperability: How compatible is portable? IBM developer. Available at: https://developer.ibm.com/articles/containers-interoperability/ (Accessed Aug 19, 2022).

Azari, A., Biersteker, J. B., Dewey, R. M., Doran, G., Forsberg, E. J., Harris, C. D. K., et al. (2021). Integrating machine learning for planetary science: Perspectives for the next decade. *Bull. AAS* 53. doi:10.3847/25c2cfeb.aa328727

Barnes, W. T., Bobra, M. G., Christe, S. D., Freij, N., Hayes, L. A., Ireland, J., et al. (2020). The SunPy project: Open source development and status of the version 1.0 core package. *Astrophys. J.* 890, 68. doi:10.3847/1538-4357/ab4f7a

Bolles, D. (2022a). ScienceMissionDirective artificial intelligence (AI) initiative. Available at: https://science.nasa.gov/open-science-overview/smd-ai-initiative (Accessed Aug 19, 2022).

Bolles, D. (2022b). Transform to open science (TOPS). Available at: https://science.nasa.gov/open-science/transform-to-open-science (Accessed Aug 19, 2022).

Burrell, A. G., Halford, A., Klenzing, J., Stoneback, R. A., Morley, S. K., Annex, A. M., et al. (2018). Snakes on a spaceship—An overview of python in heliophysics. *J. Geophys. Res. Space Phys.* 123 (10), 384–402. doi:10.1029/2018JA025877

Camporeale, E., Bortnik, J., Matsuo, T., and McGranaghan, R. M. (2022). Session proposal: Machine learning in space weather (fall meeting 2022). Available at: https://agu.confex.com/agu/fm22/webprogrampreliminary/Session159027.html (Accessed Aug 19, 2022).

Camporeale, E. (2022). ML-helio — machine learning in Heliophysics. Available at: https://ml-helio.github.io/ (Accessed Aug 19, 2022).

Camporeale, E. (2019). The challenge of machine learning in space weather: Nowcasting and forecasting. *Space Weather* 17, 1166–1207. doi:10.1029/2018SW002061

Center for Scientific CollaborationCommunity Engagement (2022). Cscce community participation model. Available at: https://www.cscce.org/resources/cpm/ (Accessed Sep 07, 2022).

Chen, X., Dallmeier-Tiessen, S., Dasler, R., Feger, S., Fokianos, P., Gonzalez, J. B., et al. (2018). Open is not enough. *Nat. Phys.* 15, 113–119. doi:10.1038/s41567-018-0342-2

Chollet, F. (2018). *Deep learning with Python*. Shelter Island, NY: Manning Publications Co.

Cline, R., Hernandez Ortiz, T., and Dunks, R. (2022). *Governance in nonprofit organizations: A literature review*. Switzerland: Zenodo. doi:10.5281/zenodo.7025611

Crooker, N. (2002). Solar-heliospheric group "shine" sheds light on murky problems. *Eos Trans. AGU.* 83, 24. doi:10.1029/2002EO000018

European Commision (2021). Assessment list for trustworthy artificial intelligence (ALTAI) for self-assessment. Available at: https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment (Accessed Aug 19, 2022).

European Commission (2019). Ethics guidelines for trustworthy AI. Available at: https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai (Accessed Aug 19, 2022).

European Commission (2022). Open science. Available at: https://research-and-innovation.ec.europa.eu/strategy/strategy-2020-2024/our-digital-future/open-science_en (Accessed Sep 07, 2022).

GEM (2022). Gem. Available at: https://gemworkshop.org/pages/gem2022/gem2022schedule.php (Accessed Aug 19, 2022).

Gotterbarn, D., Miller, K., and Rogerson, S. (1997). Software engineering code of ethics. *Commun. ACM* 40, 110–118. doi:10.1145/265684.265699

Hutchinson, B., Smart, A., Hanna, A., Denton, E., Greer, C., Kjartansson, O., et al. (2021). "Towards accountability for machine learning datasets," in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (ACM), Virtual Event Canada, March 3–10, 2021. doi:10.1145/3442188.3445918

Institute for the Future (2022). Institute for the future (IFTF). Available at: https://iftf.org/our-work/people-technology/ (Accessed Sep 01, 2022).

Lee, B. D. (2018). Ten simple rules for documenting scientific software. *PLoS Comput. Biol.* 14, e1006561. doi:10.1371/journal.pcbi.1006561

McGranaghan, R., Bhatt, A., Ozturk, D., and Kunduri, B. (2022). 2022 workshop: Data science and open science in cedar. Available at: https://cedarscience.org/workshop/2022-workshop-data-science-and-open-science-cedar (Accessed Sep 07, 2022).

McGranaghan, R. M., Bhatt, A., Matsuo, T., Mannucci, A. J., Semeter, J. L., and Datta-Barua, S. (2017). Ushering in a new frontier in geospace through data science. *JGR. Space Phys.* 122 (12), 586590. doi:10.1002/2017JA 024835

Mclarney, E., Gawdiak, Y., Oza, N., Mattmann, C., Garcia, M., Maskey, M., et al. (2021). NASA framework for the ethical use of artificial intelligence (AI) - NASA technical reports server (NTRS). Available at: https://ntrs.nasa.gov/citations/20210012886 (Accessed Aug 19, 2022).

Narock, A., Bard, C., Thompson, B. J., Halford, A., McGranaghan, R., da Silva, D., et al. (2022). Heliophysics decadal survey 2022 white paper: Responsible machine learning in Heliophysics. Available at: http://surveygizmoresponseuploads.s3.amazonaws.com/fileuploads/623127/6920789/44-54a052e28c1e3bf27580e1d768044670_NarockAyrisA.pdf (Accessed Oct 06, 2022).

NASA (2015). Heliophysics living with a star program, 10-year vision beyond 2015. Available at: https://tinyurl.com/LWS-10-year-plan.

National Academies of Sciences, Engineering, and Medicine (2018). *Open science by design*. Washington, DC: National Academies Press. doi:10.17226/25116

National Academies of Sciences, Engineering, and Medicine (2017). *Report series: Committee on solar and space Physics: Heliophysics science centers*. Washington, DC: The National Academies Press. doi:10.17226/24803

National Research Council (2013). *Solar and space Physics: A science for a technological society*. Washington, DC: The National Academies Press. doi:10.17226/13060

Open Science Collaboration (2022). The reproducibility project: A model of large-scale collaboration for empirical research on reproducibility. Available at: http://eprints.lse.ac.uk/65169/1/Kappes_Reproducibility%20project_protected.pdf (Accessed Sep 07, 2022).

Roederer, J. G. (1988). Gem: Geospace environment modeling. *Eos Trans. AGU.* 69, 786–787. doi:10.1029/88EO01064

Romick, G. J., Killeen, T. L., Torr, D. G., Tinsley, B. A., and Behnke, R. A. (1987). Cedar: An aeronomy initiative. *Eos Trans. AGU.* 68, 19–21. doi:10.1029/EO068i002p00019

Serban, A., van der Blom, K., Hoos, H., and Visser, J. (2021). "Practices for engineering trustworthy machine learning applications," in 2021 IEEE/ACM 1st Workshop on AI Engineering - Software Engineering for AI (WAIN) (IEEE), Madrid, Spain, 30-31 May 2021. doi:10.1109/wain52551.2021.00021

SHINE (2022). SHINE 2022 workshop. Available at: https://helioshine.org/?mec-events=shine-2022 (Accessed Aug 19, 2022).

Stanford Artificial Intelligence Laboratory (2022). The Stanford artificial intelligence laboratory (SAIL). Available at: https://ai.stanford.edu/ (Accessed Sep 01, 2022).

The Alan Turing Institute (2022). The alan turing Institute. Available at: https://turing.ac.uk/ (Accessed Sep 01, 2022).

The Turing Way Community (2022). *The turing way: A handbook for reproducible, ethical and collaborative research*. Switzerland: Zenodo. doi:10.5281/zenodo.6909298

The U.S. Intelligence Community (IC) (2020a). Artificial intelligence ethics framework for the intelligence community. Available at: https://www.intelligence.gov/images/AI/AI_Ethics_Framework_for_the_Intelligence_Community_1.0.pdf (Accessed Nov 02, 2022).

The U.S. Intelligence Community (IC) (2020b). Principles of artificial intelligence ethics for the intelligence community. Available at: https://www.intelligence.gov/images/AI/Principles_of_AI_Ethics_for_the_Intelligence_Community.pdf (Accessed Nov 02, 2022).

Trump, D. J. (2020). Executive order on promoting the use of trustworthy artificial intelligence in the federal government. Available at: https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-promoting-use-trustworthy-artificial-intelligence-federal-government/ (Accessed Aug 19, 2022).

U.S. Department of Defense (2020). Ethical principles for artificial intelligence. Available at: https://www.ai.mil/docs/Ethical_Principles_for_Artificial_Intelligence.pdf (Accessed Nov 02, 2022).

Vogelsang, A., and Borg, M. (2019). "Requirements engineering for machine learning: Perspectives from data scientists," in 2019 IEEE 27th International Requirements Engineering Conference Workshops (REW) (IEEE), Jeju, Korea (South), 23-27 September 2019. doi:10.1109/rew.2019.00050

White House Office of Science and Technology Policy (2022). Blueprint for an AI Bill of Rights: Making automated systems work for the American people. Available at: https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf (Accessed Sep 06, 2022).

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Sci. Data* 3, 160018. doi:10.1038/sdata.2016.18

Wilson, D. S., Ostrom, E., and Cox, M. E. (2013). Generalizing the core design principles for the efficacy of groups. *J. Econ. Behav. Organ.* 90, S21–S32. doi:10.1016/j.jebo.2012.12.010