

Trustworthy Artificial Intelligence for Environmental Sciences: An Innovative Approach for Summer School

Amy McGovern^a, David John Gagne II^b, Christopher D. Wirz^b, Imme Ebert-Uphoff^c, Ann Bostrom^d, Yuhan Rao^e, Andrea Schumacher^c, Montgomery Flora^e, Randy Chase^a, Antonios Mamalakis^c, Marie McGraw^c, Ryan Lagerquist^c, Robert J. Redmon^f, Taysia Peterson^b

^a University of Oklahoma, Norman OK

^b National Center for Atmospheric Research, Boulder CO

^c Colorado State University, Fort Collins CO

^d University of Washington, Seattle WA

^e Cooperative Institute for Satellite Earth System Studies, North Carolina State University,
Asheville NC

^f NOAA Center for Artificial Intelligence (NCAI)

Corresponding author: Amy McGovern, amcgovern@ou.edu

Early Online Release: This preliminary version has been accepted for publication in *Bulletin of the American Meteorological Society*, may be fully cited, and has been assigned DOI 10.1175/BAMS-D-22-0225.1. The final typeset copyedited article will replace the EOR at the above DOI when it is published.

ABSTRACT: Many of our generation's most pressing environmental science problems are *wicked* problems, which means they cannot be cleanly isolated and solved with a single 'correct' answer (e.g., Rittel 1973; Wirz 2021). The NSF AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography (AI2ES) seeks to address such problems by developing synergistic approaches with a team of scientists from three disciplines: environmental science (including atmospheric, ocean, and other physical sciences), AI, and social science including risk communication. As part of our work, we developed a novel approach to summer school, held from June 27-30, 2022. The goal of this summer school was to teach a new generation of environmental scientists how to cross disciplines and develop approaches that integrate all three disciplinary perspectives and approaches in order to solve environmental science problems. In addition to a lecture series that focused on the synthesis of AI, environmental science, and risk communication, this year's summer school included a unique Trust-a-thon component where participants gained hands-on experience applying both risk communication and explainable AI techniques to pre-trained ML models. We had 677 participants from 63 countries register and attend online. Lecture topics included trust and trustworthiness (Day 1), explainability and interpretability (Day 2), data and workflows (Day 3), and uncertainty quantification (Day 4). For the Trust-a-thon we developed challenge problems for three different application domains: (1) severe storms, (2) tropical cyclones, and (3) space weather. Each domain had associated user persona to guide user-centered development.




FIG. 1. Countries where AI2ES summer school attendees came from with color indicating number of participants from the different regions. Graphic is from (Dula and Craven 2022).

1. Introduction

The NSF AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography (AI2ES) recently held our second annual summer school on trustworthy AI for the environmental sciences. We are an NSF funded AI institute focused on creating trustworthy AI for a wide variety of atmospheric, climate, and ocean related applications. The goal of this In Box article is to share the innovative approach we developed for the summer school in hopes that it will facilitate the development and refinement of future interdisciplinary educational efforts. To reach this goal, we describe the details of the overall structure, our main insights, the results of an external evaluation, and our plans for continued innovation of the summer school.

2. Goals and audience

A main focus of AI2ES is to improve workforce development and broaden participation, of which summer schools are a key part. Our goal with the 2022 summer school was to share our latest research in creating and understanding the nature of trust and trustworthy AI for the environmental sciences and to facilitate the use of these techniques by our participants. Our target audience was




FIG. 2. Highest level of schooling from summer school attendees. Graphic is from (Dula and Craven 2022).

primarily early career scientists (graduate students, postdoctoral researchers and scientists newer to the workforce), and targeted more toward environmental scientists with interest in AI (rather than the other way round). Figure 1 shows the locations of our registrants, with over 600 people attending from around the world. While the majority were from the US, we had participants from 63 countries representing a wide range of economic and AI development (Zhang et al. 2022). Figure 2 shows the education level of the attendees. From these numbers, the majority of our attendees were current graduate students or people in the workforce.

3. Summer school lectures on trustworthy AI

The summer school was broken into two halves: each morning we held a series of lectures focused on specific topics; each afternoon, we had a Trust-a-thon with a specific focus tied to the morning lectures. We describe the trust-a-thon in the following section. Each day of summer school focused on a specific theme, with the themes connected across the lectures and trust-a-thon activities. The lectures for each day were created in a convergent manner involving an interdisciplinary team of AI2ES researchers and collaborators from AI, environmental sciences, and risk communication. This convergent development of the materials benefited all participants by giving them a synergistic perspective across disciplines. We also placed a strong emphasis on making the lectures engaging

and as participatory as we could, given the virtual format and how many participants we had. We used Slido, an online and interactive polling platform, to host 36 different activities (polls, surveys, etc.) across the 4 days. We also used Slido for open questions all throughout the lectures, which was very active with participants asking questions and members from across AI2ES answering them, sharing resources, and facilitating conversations.

Day 1's theme was trust and trustworthiness in AI for environmental sciences. We first approached this from the risk communication perspective, where we shared information from the AI2ES team's interviews with forecasters about trust in AI products (Cains et al. 2022), as well as from the AI and trust literature (e.g., Hoff and Bashir (2015); Chiou and Lee (2021)). Two critical aspects of trust for the forecasters are the performance of the model and the ability to peer "inside the black box" (McGovern et al. 2019). As such, we covered common evaluation metrics as well as eXplainable AI (XAI, e.g., see Mueller et al. (2019); Biran and Cotton (2017)) methods for traditional ML methods. As part of this theme we also focused on strategies for effective interdisciplinary work (Peek and Guikema 2021; Morss et al. 2021).

Day 2's theme focused in more depth on the explainability and interpretability of ML models and how these affect trust. Again, we first presented this interdisciplinary topic from a social science perspective making use of data from our forecaster interviews. We then jumped into technical details of XAI methods for deep learning, including an introduction to XAI methods of particular interest for environmental science tasks, and benchmarks for the environmental sciences (Mamalakis et al. 2022a,b). We highlighted the need to include multiple approaches to XAI as no one approach provides a single "right" answer (McGovern et al. 2019).

On the third day, we focused on the importance of trustworthy data and workflows and how these influence trust in the final models learned by the AI methods. We emphasized the importance of considering ethics and the potential for bias all throughout the development process (McGovern et al. 2022a). Our overview of trustworthy data and workflows demonstrated how these factors can influence trust and users' perceptions, as well as help address issues with reproducibility and replicability (NASEM 2019). We also examined the use of case studies in detail, highlighting their role in developing trust.

The final day focused on uncertainty throughout the lifecycle of the AI system development. This included learning about common methods for uncertainty quantification of AI methods

(from Quantile regression to Bayesian Neural Networks), metrics to evaluate uncertainty estimates obtained by those methods (from spread-skill plots to discard tests), and risk communication strategies for communicating uncertainty to diverse audiences (Van der Bles et al. 2019; Millet et al. 2020; Morgan 2009; Padilla et al. 2020). Finally, the role of uncertainty in trust in AI was discussed.

The lectures were accompanied by detailed Jupyter notebooks that illustrate the use of the various concepts, especially the XAI methods, for environmental science applications. These notebooks provided a set of tools for participants to use in the Trust-a-thon component of the summer school, as well as in their future research projects.

4. Trust-a-thon

The trust-a-thon was envisioned as a twist on the traditional machine learning hackathon. The innovative idea for the trust-a-thon came from second author (Gagne). To the best of our knowledge this is the first event of its kind in the field. In traditional machine learning hackathons, participants train machine learning models given data and a specific task. In contrast, for the trust-a-thon the participants were given the following elements:

1. Description of an environmental science application, a task, and an accompanying data set. Development of these was led by AI2ES environmental science experts and collaborators.
2. Accompanying code that implements one or more simple machine learning models for the desired task. Development led by AI2ES AI experts.
3. Fictional user personas that span different types of potential end users, such as emergency managers, forecasters, and more. Development led by AI2ES risk communication and social science experts.

Note that even though experts from specific disciplines led the development of each element above, the overall development of these elements was a convergent effort where experts from all three disciplines contributed to all elements.

The goal of the trust-a-thon was for each team to develop a more nuanced understanding of the trustworthiness of the provided machine learning models for a given task from the perspective

of a potential end user of the system, thus *integrating insights and tools from the disciplines of environmental science, AI, and risk communication/social science*.

Challenge problems were developed for three different application domains: (1) severe storms, (2) tropical cyclones, and (3) space weather. Each challenge problem featured Jupyter notebooks outlining the problem and how to train a model and apply to that problem the evaluation, explanation, case study, and uncertainty quantification methods discussed in the lectures. Each team also received a problem guide with fictional user personas to give them a better sense of what aspects of the problem users might care about the most. These are discussed separately in the next section. Each day of the summer school concluded with a small group discussion of several reflection questions, which each group then answered and posted on the summer school blog.

5. Use of Persona to Encourage Participants to teach user-centered development

As mentioned above, we added personas to each of the trust-a-thon challenge problems in order to encourage participants to focus their development on the intended users' needs. The use of personas, or user profiles that capture an archetypal user, is a common strategy in user experience work that we adapted for the summer school. Our emphasis on user-centered design is a key element differentiating the trust-a-thon from a standard hackathon, which typically focuses on performance metrics, because it focuses on developing, refining, and communicating to better serve a specific audience. Our goal is that teaching this approach will not only generate AI/ML models that are more useful to and used by the targeted end-users, but that it will also help foster trust among research teams and those who they aim to serve.

The idea to create and include personas for the trust-a-thon originated from the third author (Wirz). Personas were developed in close collaboration with coauthors who were experts in each of the three domain areas (severe weather, space weather, and tropical cyclones). The user personas spanned different types of potential end-users such as emergency managers, forecasters, and more, which were customized to each application. A sample persona is provided in Figure 3, while a list of all personas is included in the supplemental materials.

From this example you can see how we gave context about the decision space the user is operating within and their specific needs (long lead-times and event duration). We tailored these needs to

Flight Planner for a Large Airline Company

- *Background:* Geomagnetic storms and disturbances can interfere with radio communication systems on certain routes. These events can result in radio communication not functioning for several days in certain areas. During these events, flights must be redirected to areas and latitudes where they can use other forms of communication.
- *User needs:* The flight planner needs to know about when events are going to start at least 7 days beforehand so they can reroute flights and notify their customers. The more time they have, the better. Their customers do not like changed itineraries and it can cost the company a lot of money. They also need to know how long the events will last, so they know just how many flights they need to change and what flights will not need to be changed. This information will help them better manage customer relations.

FIG. 3. Example persona for the user-centered development in the trust-a-thon.

be compatible with the dataset we provided and the associated AI/ML related activities. We then encouraged participants to focus on meeting these specific needs.

In sum, the trust-a-thon and user persona approach to teaching AI development is innovative in the ways that it emphasizes users' needs and perceptions over a singular focus on maximizing performance metrics. This reflexive and contextual approach is essential for developing AI that is both used and trusted by users.

6. Insights from the Summer School

Many of our generation's most pressing environmental science problems are *wicked* problems, which means they cannot be cleanly isolated and solved with a single 'correct' answer (e.g., Rittel 1973; Wirz 2021). Given their complexity and importance, addressing these problems often requires convergent approaches. Convergent research has two parts: it is use-inspired and driven by compelling problems, which weather and climate provide, and it requires deep integration across multiple disciplines that come together to create something that could not be done individually (National Science Foundation 2018). In our work, we work synergistically across three disciplines, environmental science, AI, and social science, including risk communication. This fact is the motivation behind the approaches developed for AI2ES in general, and for this summer school in particular.

For the summer school we made it our central goal to teach a new generation of environmental scientists how to work across disciplines and develop approaches that integrate all three disciplinary perspectives and approaches in order to solve environmental science problems. We did this through

two unique aspects of the summer school. First, we emphasized *synergistic co-development of the entire curriculum*, including both lectures and trust-a-thon, by a team of experts spanning the three disciplines of environmental science, AI, and social science. While prior AI2ES summer schools and courses have also sought to teach these three elements together, this year we took the co-development and integration of these three topics to an entirely new level, which we believe was possible only now thanks to our two years of extensive experience of integrating these topics in our AI institute activities (founded in 2020).

The second aspect is the development of an innovative trust-a-thon component that focused on *evaluating* machine learning models from a user-centered perspective while integrating tools and perspectives from the three disciplines. In contrast, most ML hackathons focus on *developing* and tuning ML models against a single performance metric that may not align with operational usefulness. The trust-a-thon activities encouraged participants to embrace a more holistic view of model quality and how well the ML models presented were fit for purpose. Furthermore, the user persona component helped us demonstrate the importance of user-centered development and give students the opportunity to put the ideas into practice.

Finally, it should be mentioned that in our experience the development of such a synergistic curriculum tends to be much more time consuming than developing a more segmented curriculum. That is because experts from all disciplines need to be involved in the planning of all lectures and trust-a-thon activities. Thus, the development of slides and materials for different days could not just be split and delegated to subgroups early on. Instead, the team had countless meetings among the entire group to coordinate development, followed by subgroup development and large group review. Overall, the preparation required much more time from a larger team than we had anticipated, but we are delighted by the feedback we received from participants about the unique perspectives they learned in this unique setting, which made our effort worthwhile.

We offer to the readers a few lessons learned about things that we tried that also did not work over the course of several years of summer school. The biggest lessons center on maintaining interest and attendance throughout the week. While the online mode is successful in bringing people from around the world, it also means that people often drop off quickly in participation as the week goes on. This becomes especially obvious if the sessions take too many hours. We are not suggesting that in-person interaction is required, as online provides much greater inclusivity, but instead that

designers of such summer schools ensure that there is a strong interactive component to maintain interest and that the sessions do not stretch so long that they lead to meeting fatigue. Another aspect to reducing the dropoff is to ensure that participants understand the expectations and prerequisite knowledge that is expected of them. Our first offering of the summer school in 2020 focused on a more introductory approach and our later offerings assumed that participants had much of this introductory background. While this worked for the majority of participants, some participants expected to be focused on introductory material again. Ensuring that everyone understands the expectations ahead of time will help address this.

Additional lessons learned focus on the global nature of an online summer school. If a hack-a-thon or trust-a-thon type of approach is to be included, the world-wide nature of the audience makes it critical that teams are formed in local time zones and that teams have people from around the world who can answer questions at their local times. With participants from around the world but all of the instructors residing in the US, we did not meet the needs of the participants in opposite time-zones in a timely manner as we were often asleep when they had questions. Embedded in this is a critical additional lesson, timely feedback is crucial to maintaining interest and engagement in an online course. For example, one of the facilitators for the tropical use case engaged with the team blog posts each day with constructive feedback, resulting in higher engagement for that use case throughout the week.

Finally, it is important for all of the people creating the course to realize up-front that doing a course such as this is a significant time investment and should be recognized as such on annual evaluations. For example, a summer course such as this does not typically count as a course load for a faculty member but it should be recognized either as an overload course or as a service offering to the community.

7. Summer School Evaluation

As part of their role in evaluating AI2ES, Horizon Research Inc (HRI) evaluated the summer school, including surveys and interviews with participants. We provide highlights from the full evaluation and full details can be found at (Dula and Craven 2022). HRI evaluated the individual components: lectures and trust-a-thon, as well as the summer school as a whole. Overall the evaluation of the lectures was very positive, with the vast majority of participants agreeing that the

pacing was appropriate. When asked about their most useful topics across all of summer school, many chose the lectures and then specifically the XAI and uncertainty quantification lectures. Two of the interviews provided quotes that explain why they enjoyed the lectures particularly and why they felt they were high impact:

The entire summer school had an impact on my learning because the lectures were clear, and I think they were meticulously organized, ... and most of what I learned from the summer school, I learned from the lectures. -Attendee

It is not possible to give an in-depth learning in a short period of time, but their training is tremendous, and I'm totally fascinated by the kind of breadth they provide. ... So that is what I feel is the major contribution of summer school, to make you aware that these kinds of ideas are there -Attendee

When specifically asked about the Trust-a-thon and what worked well, one of the attendees highlighted the user personas.

I also like the idea of the personas because it gave you an example, a user that you're trying to provide this model for. To build this model for, let's say, someone working in a metrology center or something like that, and this is a person going to be using the model that you've built. So having this understanding of how to relay information to them was an amazing thing, that I think a key takeaway that we all would've gotten from the sessions altogether. -Attendee

As with the lectures, the trust-a-thon overall was rated well with a few adjustments suggested for future years (expected since this was our first year of trying a trust-a-thon instead of a hack-a-thon).

As part of the overall evaluation, HRI performed a survey and demonstrated statistically significant improvements in respondents' self-reported understanding of trustworthy AI for weather and climate applications. Figure 4 shows the results for the lectures and the trust-a-thon separately and then the composite scores overall, all of which are positive. Finally, we finish our summary of the evaluation with a quote from a participant highlighting their overall experience in an "intensive week" but that really changed their understanding and views on AI for weather and climate.

First of all, in the aspect of acquiring more knowledge about environmental science, that was highly achieved because I can say that through the Summer School, I didn't leave as

Participants Indicating Good or Strong Understanding[†] of the Lecture Topics

	Percent of Respondents (N = 249)	
	Retrospective Pre	Post
What trustworthy AI means	22	86
The importance of thinking about end users	36	83
The importance of trustworthy data and workflows	24	79
Available techniques for XAI	15	72
The difference between explainable AI (XAI) and interpretable AI	13	69
How to communicate uncertainty	22	68
Uncertainty quantification methods	20	68
How to pick case studies	17	68

[†] Includes those indicating “Good” or “Strong” understanding on a four-point scale from 1 “None” to 4 “Strong.”

Participants Indicating Good or Strong Understanding[†] of the Trust-a-Thon Topics

	Percent of Respondents (N = 95)	
	Retrospective Pre	Post
The importance of thinking about the end user when developing AI	16	84
How to assess AI trustworthiness	13	73
How to implement machine learning systems that have been developed for a range of environmental science applications	16	72
How to implement explanatory AI (XAI) methods	18	69
How to implement ethical AI methods	16	59

[†] Includes those indicating “Good” or “Strong” understanding on a four-point scale from 1 “None” to 4 “Strong.”

Participants' Understanding of Summer School Topics




FIG. 4. Summer school evaluation by HRI. Top row: Pre and post survey for lecture. Middle row: pre and post survey for Trust-a-thon. Bottom graphic: combined score showing statistically significant improvements (two-tailed paired samples t-test, $p < 0.05$, denoted by the *). Table and graphics are from (Dula and Craven 2022).

I joined. It was an intensive week. ... I covered quite a lot through that summer school. ... And then in the direction of trustworthy AI, I can say that it helped me to see how best I can use to trustworthy AI and not just in the environmental science space, but in my application of AI generally, because I've been using AI for a while, and I didn't see interpretability as a thing to always show people, as a primary thing to show. -Attendee

8. Future Summer Schools: Personalized learning journeys

Over the past three years, we have developed a large library of materials, including lectures and Jupyter notebooks from summer schools, short courses, and tutorials, designed to teach environmental, atmospheric, ocean, climate, and physical scientists about AI/ML and risk communication. For Summer 2023, we are planning to create an online course which will facilitate personalized learning journeys. These learning journeys will leverage our existing material and will be targeted across the full spectrum of environmental and earth science researchers, ranging from those who are new to AI/ML to those who are more experienced and need to learn a specific method or application. Topics will include basic AI/ML methods relevant to environmental, atmospheric, ocean, and climate science problems, as well as deep learning, XAI, trust in AI, and other topics.

Overall, we hope our approach, insights, evaluation, and future plans serve as a foundation for teaching development, both of AI and non-AI products, in ways that emphasize the importance of end users' perceptions and needs.

Acknowledgments. This material is based upon work supported by the National Science Foundation under Grant No. ICER-2019758. AI2ES summer school was a collaboration between AI2ES, the NOAA Center for AI (NCAI), the National Center for Atmospheric Research (NCAR), Radian Earth Foundation, and the Learning the Earth with Artificial Intelligence and Physics (LEAP) Science Technology Center. NCAR is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977. Computing credits for the Trust-a-thon Jupyterhub platform were provided by Google Cloud Platform.

The authors would like to acknowledge the contributions of many additional people who worked tirelessly on making this summer school a success. For the Trust-a-thon, major contributors are Rob Redmon, Manoj Nair, and LiYin Young for the space weather application; Hamed Alemohammad, Renee Pieschke, Jason Stock, Marie McGraw, Akansha Singh Bansal, Kate Musgrave and Imme

Ebert-Uphoff for the tropical cyclone application; and Randy Chase and Monte Flora for the severe weather application. For the lecture materials, additional contributors are Mariana Cains, Julie Demuth, Katherine Haynes and Philippe Tissot.

Data availability statement. All lectures, slides, and jupyter notebooks from summer school are online. The slides and GitHub are shared using Zenodo (McGovern et al. 2022b; Flora et al. 2022) and the notebooks and user personas are on the AI2ES website in the education section: <https://www.ai2es.org/products/education/>. Recordings of the lectures are available at <https://youtu.be/3ZL7U0r7n0g>.

References

Biran, O., and C. Cotton, 2017: Explanation and justification in machine learning: A survey. *IJCAI-17 workshop on explainable AI (XAI)*, Vol. 8, URL http://www.cs.columbia.edu/~orb/papers/xai_survey_paper_2017.pdf.

Cains, M. G., and Coauthors, 2022: Nws forecasters' perceptions and potential uses of trustworthy ai/ml for hazardous weather risks. *102nd American Meteorological Society Annual Meeting*, AMS.

Chiou, E. K., and J. D. Lee, 2021: Trusting automation: Designing for responsivity and resilience. *Human factors*, 00187208211009995.

Dula, J. A., and L. M. Craven, 2022: 2022 Trustworthy Artificial Intelligence for Environmental Science (TAI4ES) Summer School Feedback. Tech. rep., Horizon Research Inc. Available at https://www.ai2es.org/wp-content/uploads/2023/02/HRI_2022_Summer_School_AI2ES.pdf.

Flora, M., R. Redmon, M. McGraw, A. S. Bansal, D. J. Gagne, and J. Stock, 2022: ai2es/tai4es-trustathon-2022: Trustworthy Artificial Intelligence for Environmental Science (TAI4ES) Summer School 2022. Zenodo, URL <https://doi.org/10.5281/zenodo.6784569>, <https://doi.org/10.5281/zenodo.6784569>.

Hoff, K. A., and M. Bashir, 2015: Trust in automation: Integrating empirical evidence on factors that influence trust. *Human factors*, **57** (3), 407–434.

Mamalakis, A., E. A. Barnes, and I. Ebert-Uphoff, 2022a: Investigating the fidelity of explainable artificial intelligence methods for applications of convolutional neural networks in geoscience. *Artificial Intelligence for the Earth Systems*, 1 – 42, <https://doi.org/10.1175/AIES-D-22-0012.1>, URL <https://journals.ametsoc.org/view/journals/aiies/aop/AIES-D-22-0012.1/AIES-D-22-0012.1.xml>.

Mamalakis, A., I. Ebert-Uphoff, and E. A. Barnes, 2022b: Neural network attribution methods for problems in geoscience: A novel synthetic benchmark dataset. *Environmental Data Science*, **1**, e8, <https://doi.org/10.1017/eds.2022.7>.

McGovern, A., I. Ebert-Uphoff, D. J. Gagne, and A. Bostrom, 2022a: Why we need to focus on developing ethical, responsible, and trustworthy artificial intelligence approaches for environmental science. *Environmental Data Science*, **1**.

McGovern, A., R. Lagerquist, D. Gagne, G. Jergensen, K. Elmore, C. Homeyer, and T. Smith, 2019: Making the black box more transparent: Understanding the physical implications of machine learning. *Bulletin of the American Meteorological Society*, **100** (11), 2175–2199, URL <https://doi.org/10.1175/BAMS-D-18-0195.1>.

McGovern, A., and Coauthors, 2022b: Trustworthy Artificial Intelligence for Environmental Science (TAI4ES) Summer School 2022. Zenodo, URL <https://doi.org/10.5281/zenodo.6784187>, This material is based upon work supported by the National Science Foundation under Grant No. ICER-2019758., <https://doi.org/10.5281/zenodo.6784187>.

Millet, B., A. P. Carter, K. Broad, A. Cairo, S. D. Evans, and S. J. Majumdar, 2020: Hurricane risk communication: visualization and behavioral science concepts. *Weather, climate, and society*, **12** (2), 193–211.

Morgan, M. G., 2009: *Best practice approaches for characterizing, communicating and incorporating scientific uncertainty in climate decision making: Synthesis and assessment product 5.2 Report*, Vol. 5. US Climate Change Science Program.

Morss, R. E., H. Lazarus, and J. L. Demuth, 2021: The “inter” within interdisciplinary research: Strategies for building integration across fields. *Risk Analysis*, **41** (7), 1152–1161.

Mueller, S. T., R. R. Hoffman, W. Clancey, A. Emrey, and G. Klein, 2019: Explanation in human-ai systems: A literature meta-review synopsis of key ideas and publication and bibliography for explainable ai. *arXiv preprint arXiv:1902.01876*, URL <https://arxiv.org/abs/1902.01876>.

NASEM, 2019: Reproducibility and replicability in science.

National Science Foundation, 2018: Growing Convergence Research. URL https://www.nsf.gov/news/special_reports/big_ideas/convergent.jsp.

Padilla, L., M. Kay, and J. Hullman, 2020: Uncertainty visualization.

Peek, L., and S. Guikema, 2021: Interdisciplinary theory, methods, and approaches for hazards and disaster research: An introduction to the special issue. *Risk Analysis*, **41** (7), 1047–1058.

Rittel, W. M., H.W.J., 1973: Dilemmas in a general theory of planning. *Policy Sciences*, **4**, 155–169, <https://doi.org/10.1007/BF01405730>.

Van der Bles, A. M., S. Van Der Linden, A. L. Freeman, J. Mitchell, A. B. Galvao, L. Zaval, and D. J. Spiegelhalter, 2019: Communicating uncertainty about facts, numbers and science. *Royal Society open science*, **6** (5), 181 870.

Wirz, C., 2021: Risk perceptions for wicked issues: Toward more nuanced risk communication. Doctoral dissertation, University of Wisconsin-Madison.

Zhang, D., and Coauthors, 2022: The AI index 2022 annual report. AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University.