

AI for Anticipatory Action: Moving beyond Climate Forecasting

Benjamin Q. Huynh¹, Mathew V. Kiang²

¹ Johns Hopkins University, Department of Environmental Health and Engineering

² Stanford University, Department of Epidemiology and Population Health
bhuyh@jhu.edu, mkiang@stanford.edu

Abstract

Disaster response agencies have been shifting from a paradigm of climate forecasting towards one of anticipatory action: assessing not just what the climate will be, but how it will impact specific populations, thereby enabling proactive response and resource allocation. Machine learning models are becoming exceptionally powerful at climate forecasting, but methodological gaps remain in terms of facilitating anticipatory action. Here we provide an overview of anticipatory action, review relevant applications of machine learning, identify common challenges, and highlight areas where machine learning can uniquely contribute to advancing disaster response for populations most vulnerable to climate change.

Background

Anticipatory action – proactive measures taken in response to forecasted risks and disasters before they occur – has emerged as a crucial framework for disaster response (Figure 1) (Weingärtner and Wilkinson 2019; Thalheimer, Simperingham, and Jjemba 2022; de la Poterie et al. 2023; Pople et al. 2021). Central to anticipatory action is impact-based forecasting, an approach that focuses on predicting the specific impacts of an impending disaster, such as health outcomes, disruption of services, or economic loss (Yu et al. 2022; Merz et al. 2020; Huynh et al. 2021; Liu 2012; Gerl et al. 2016). By contrast, climate forecasting typically reports environmental conditions, such as wind speeds or rainfall. Impact-based forecasting is typically more accessible to humanitarian practitioners, but harder to predict; climate forecasting is generally more accurate, but may not adequately communicate potential hazards. However, the two approaches are not mutually exclusive – in fact, we examine here areas in which machine learning can uniquely be applied to leverage the predictive power of climate forecasting as well as the applied perspective of impact-based forecasting.

Despite the promise of increased resilience from anticipatory action, major gaps remain in terms of developing robust, trustworthy systems. Forecast-based financing, a system where financing is allocated proactively based on risk estimates for models, has become increasingly used

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

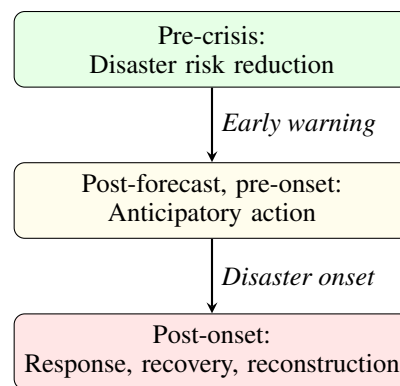


Figure 1: Typical timeline of disaster preparedness, anticipatory action, and response. *Disaster risk reduction* refers to strengthening resilience before crises occur, and *anticipatory action* refers to proactive response and resource allocation after an early warning system flags an incoming disaster.

within United Nations agencies and non-governmental organizations. However, because regions that receive higher-risk forecasts receive more aid, poor predictive performance and causal validity of such forecasts may lead to misallocation of resources (Coughlan de Perez et al. 2015; Van den Homberg, Gevaert, and Georgiadou 2020). The gap between the promise and current reality of anticipatory action may be in part aided by advances in AI research; we examine here potential areas of opportunity where AI can be used to improve anticipatory action.

Challenges and Opportunities

Figure 2 shows a typical undertaking of the forecasting portion of anticipatory action. A model is developed based on spatiotemporal environmental data and used to make population-specific impact forecasts. From there, the forecasts are used to enable proactive resource allocation. We broadly identify two strategies to forecast disaster impacts: (1) a predictive forecasting approach where the impacts of the disaster are directly modeled; and (2) a causal approach wherein the disaster is meteorologically modeled, then out-

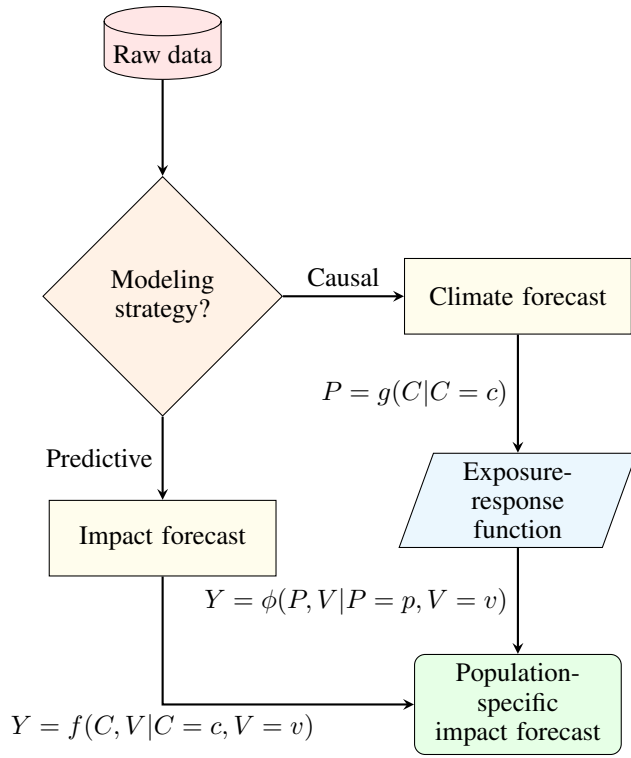


Figure 2: Typical workflow for the forecasting portion of anticipatory action. After population-specific impact forecasts are provided, targeted resource allocation can be conducted.

puts are run through an exposure-response function to ascertain the impact of the disaster. A subsequent step is taken after either approach to apply the model to a target population for a population-specific forecast. Table 1 shows examples of how impact-based forecasts can be used.

Problem Formulation

Formally, let $C = c_1, c_2, \dots, c_k$ be a set of climate-related features and let $V = v_1, v_2, \dots, v_k$ be a set of population-specific features. The predictive forecasting approach to directly model disaster impacts is represented by

$$Y_{s,t+n} = f(C, V|C = c_{s,t}, V = v_{s,t})$$

where Y is the outcome of interest, s and t are spatiotemporal indices, and n is the forecasting lead time. Machine learning approaches have been used for impact-based forecasting for outcomes related to disasters, such as modeling displacement, food insecurity, disease outbreaks, economic loss, or infrastructure assessment (Pham and Luengo-Oroz 2022; Huynh and Basu 2019; Foini et al. 2023; Wang, Chen, and Marathe 2019; Proma et al. 2022; Oshri et al. 2018; Nethery et al. 2021). Most machine learning approaches for impact-based forecasting directly model the impact without using climate forecasting as an intermediary; to the extent climate variables are used, they are typically just more predictors for the model.

In terms of climate forecasting, such models can be rep-

resented as

$$P_{s,t+n} = g(C|C = c_{s,t})$$

where P is a modeled climate event, such as a risk score or climate condition, and g is the climate forecasting function. Machine learning models for climate forecasting, while historically inferior to global simulation methods like numerical weather prediction, have since reached comparable or even superior performance to simulation-based methods for a variety of forecasting applications, ranging from tropical cyclones, heat waves, rainfall, floods, wildfires, and more (Chen, Zhang, and Wang 2020; Mosavi, Ozturk, and Chau 2018; Cifuentes et al. 2020; Jain et al. 2020; Zhang et al. 2023; Bi et al. 2023; Rolnick et al. 2022). Machine learning models are far less computationally expensive than global simulation models, and can be used to supplement them by emulating their predictions. Because such models are emulating mechanistic physical processes, predictive performance is often very high compared to other forecasting domains. Many machine learning models frame their outputs as risk scores, providing spatiotemporal risk estimates for a given type of natural disaster.

Causally forecasting disaster impacts is represented by

$$Y_{s,t+n} = f(C, V|do(C = c_{s,t}), V = v_{s,t})$$

where the *do* operator refers to causally setting the climate variables C to be specific values, rather than using a predictive relationship (Pearl 2009). In practice, causal forecasting is facilitated by an *exposure-response function*:

$$Y = \phi(P, V|P = p, V = v)$$

where ϕ represents the relationship between the climate event and the impact of interest. Exposure-response functions are typically used in environmental epidemiology, where the relationship between an exposure and a health outcome is identified, then that relationship is used to simulate health outcomes based on known exposure levels (Aunan and Pan 2004). Depending on how they are identified, exposure-response functions can vary in levels of causal validity, ranging from results from randomized controlled trials to associations identified from observational data using causal inference. In other fields, this relationship may also be known as a *damage function* (Merz et al. 2020).

Open Areas for Improvement

The decision to use a causal or predictive modeling strategy is context-specific, and both approaches have their own advantages and disadvantages. For causal forecasting, exposure-response functions can be difficult to identify, depending on the exposure. Predictive forecasting is subject to confounders, threatening causal validity (Peters, Bühlmann, and Meinshausen 2016). Both approaches face challenges of dataset shift and algorithmic fairness. Forecasting for anticipatory action thus remains an open problem with multiple potential areas for improvement, and we discuss here how AI can be used to address such issues.

Exposure-response functions Exposure-response functions are difficult to obtain, and especially so for disaster-related impacts. Disasters vary in magnitude, location, and

| Hazard | Climate Forecast | Exposure-response function(s) | Impact-based Forecast |
|----------|--|--|---|
| Cyclone | A category 4 tropical cyclone with windspeed 220km/h is anticipated to make landfall in 36-48 hours. | Relationship between tropical cyclone and building damage/displacement | 40-60% of housing in Region A will be damaged, with an anticipated 20,000-40,000 displaced individuals. |
| Wildfire | Wildfire smoke will reach City B tomorrow morning and cover the city for 48-72 hours. | Relationship between smoke exposure and hospitalizations; relationship between excess hospitalizations and drug shortages | 200-400 excess hospitalizations are expected, and extra inhalers will be necessary to prevent shortages. |
| Drought | Precipitation far below average is likely for the summer in Province C. | Relationship between precipitation and crop failure/food insecurity, relationship between food insecurity and malnourishment | Food insecurity will increase, and prices for affected crops will spike by 55-75%. Incidence of child malnourishment will increase by 20-30%. |

Table 1: Example scenarios of disasters and their potential forecasts. Cyclone refers to tropical cyclones.

population characteristics, with relatively infrequent occurrence – estimating a valid association between a disaster and an impact typically requires high-resolution data and a plausible empirical strategy (Nethery et al. 2021; Parks et al. 2022). Even air pollution, which is continuously monitored with high-resolution exposure and health outcomes data, drives contention over what the exposure-response function should be (Zigler and Dominici 2014). Machine learning can be useful in improving estimation of exposure-response relationships by producing better characterizations of disaster impacts. For example, satellite imagery in conjunction with AI has been used to assess building damage from disasters (Novikov et al. 2018; Zhao and Zhang 2020; Oshri et al. 2018); such information could be used to develop exposure-response relationships between disaster occurrence and building damage, potentially enabling forecasts of building damage.

As another example, AI can be used to improve estimates of health outcomes such as mortality or disease prevalence after a disaster. In low-resource environments, post-disaster estimates of mortality or disease prevalence are conducted by taking surveys across clusters of households. Results from such surveys can vary drastically due to sampling bias, which can be difficult to overcome through randomization due to poor access to households post-disaster. The survey approach remains the gold-standard, as official national or international statistics are generally too slow to capture the magnitude of a disaster shortly after it has occurred (Kishore et al. 2018; Gang, O’Keeffe, and Roberts 2023). AI can be used in conjunction with post-stratification, a technique typically used in the social sciences to learn from biased samples, to estimate adjusted mortality or disease prevalence rates (Gelman et al. 2016; Broniecki, Leemann, and Wüest 2022; Van der Heyden et al. 2014).

Additionally, mortality or disease prevalence surveys typically obtain population-wide estimates by extrapolating from survey results:

$$\frac{E_s}{N_s} \approx \frac{E_f}{N_f}$$

where E and N represent events and population size respectively, and s and f indices indicate observations from the survey data and the full population at large, respectively. Typically, E_f is estimated by assuming a value for N_f (as well as adjusting for any covariates), but there is often uncertainty regarding the value of N_f as populations tend to migrate in response to disasters, potentially affecting estimates of E_f . Any AI methods to estimate N_f , whether through migration or population modeling, can thus be valuable in improving identification of exposure-response functions (Doocy, Cherewick, and Kirsch 2013; Robinson, Hohman, and Dilkina 2017; Yeh et al. 2020).

Dataset Shift All forecasting approaches are subject to dataset shift, where there is a difference in distribution between training data and deployment-time testing data. This is particularly true for disaster-related forecasts: first, unprecedented events happen regularly under climate change, and models may not provide valid results for environmental conditions that have never been seen before (Dickerman and Hernan 2020); secondly, because disasters happen relatively infrequently, there will not be sufficient historical data regarding how a disaster of any given magnitude affects any given region or subpopulation of a country. Taken together, high-resolution country-wide forecasts will require targeted population-specific forecasts in a manner that addresses dataset shift.

One approach to handling dataset shift is leveraging *transportability* of causal effects (Bareinboim and Pearl 2012; Pearl and Bareinboim 2014). Transportability provides a framework to assess whether causal effects can be applied from one environment to another. As an example, suppose we have a plausible causal effect estimate saying that wildfire smoke caused an $X\%$ increase in hospitalizations per $\mu\text{g}/\text{m}^3$ of fine particulate matter in City A. We want to know how many additional hospitalizations we’d expect to see if smoke reached City B, but City B has a different age distribution than City A, and age is a confounder in that it affects both exposure levels to smoke and likelihood of hospitalization. Transportability theory determines that if we can iden-

tify both the full causal effect:

$$P(\text{hospitalization}|\text{do}(\text{smoke}))$$

and the age-specific causal effect:

$$P(\text{hospitalization}|\text{do}(\text{smoke}), \text{age})$$

then it is feasible to transport causal estimates from City A to City B. Transportability has been applied in various complex scenarios in health (Rudolph et al. 2018; Prospero et al. 2020), and could also be useful for improving causal validity of disaster forecasts.

Relatedly, another approach to handle dataset shift for impact forecasting involves *invariant prediction* (Peters, Bühlmann, and Meinshausen 2016; Pfister, Bühlmann, and Peters 2019; Arjovsky et al. 2019). Briefly, approaches leveraging invariant prediction allow for selection of causal features from a model by identifying model configurations with high prediction invariance, enabling more robust predictions less likely to be influenced by spurious correlations. Using approaches based on invariant prediction would improve the causal validity of predictive impact forecasting without requiring the intermediary step of calculating an exposure-response function, thereby improving the reliability of impact predictions in relationship to varying climate variables. Although these approaches require the strong assumption of there being no unobserved confounders present, sensitivity analyses for potentially unobserved confounders could be used to provide uncertainty intervals for predictions (Rosenbaum and Rubin 1983; Yadlowsky et al. 2022; Jung et al. 2017). Applications of invariant causal prediction have recently been gaining traction within medicine (Prospero et al. 2020; Subbaswamy, Schulam, and Saria 2019; Subbaswamy, Chen, and Saria 2022) and earth sciences (Runge et al. 2019; Sheth et al. 2022), suggesting that applications to climate disasters and anticipatory action may be feasible.

Data Scarcity and Algorithmic Bias Climate data are inequitably collected, with marginalized populations typically having poorer and lower-quality data (Rolnick et al. 2022; Grover 2022). 60% of countries lack basic water information services that can be used for flood preparedness, and roughly half of all countries do not have multi-hazard early warning systems (Cullmann et al. 2021). The World Meteorological Organization roughly estimates that countries without such early warning systems have over eight times the disaster-related mortality rates than countries with early warning systems (Honore, Kumar, and Speck 2022). Data scarcity degrades the performance of predictive models for regions that are most vulnerable to disasters in the first place, making anticipatory action more difficult within low-resource countries. Typical numerical weather prediction models rely on data collected globally to make better localized predictions, but AI efforts to ameliorate data scarcity such as localized predictions via transfer learning, or up/downscaling climate predictions may help circumvent such requirements and improve impact-based forecasts (Xie et al. 2016; Li et al. 2021; Vosper et al. 2023).

Careful consideration needs to be given to the ethical implications of AI models for impact forecasting, particularly as forecast-based financing becomes increasingly used

(Coughlan de Perez et al. 2015). In particular, practitioners should exercise caution as to how the choice of metric affects equity considerations. Treating economic loss as an absolute metric for impact assessment favors wealthier regions; considering risk and loss as relative has been shown to provide more equitable results in risk assessment (Kind, Botzen, and Aerts 2020). Additionally, considering how disasters impact regions beyond economic loss, and instead using metrics like health or wellness, may provide a more nuanced characterization of disaster impacts (Hino and Nance 2021). Algorithmic audits and assessments commonly used for composite indicators or risk scoring machine learning models should be used to assess the algorithmic fairness of impact-based forecasts (Mayson 2019; Obermeyer et al. 2021; Gevaert et al. 2021; Huynh et al. 2023; Bhagwat, Ferryman, and Gibbons 2023; Paulus and Kent 2020).

Discussion

In this paper, we surveyed how AI advances can be used to improve anticipatory action in terms of current methodological gaps and potential areas for improvement. There are several limitations to our paper. First, this is not an exhaustive review of all relevant literature: we refer to other review papers where applicable, and subfields such as AI for climate forecasting or causal prediction are sufficiently deep and nuanced that reviewing all relevant works is out of scope. Secondly, there are other crucial areas of improvement that are not mentioned here, such as model explainability or uncertainty quantification – our subjective judgment is that they represent long-standing challenges of fundamental importance that are not sufficiently specific to anticipatory action to be in scope of this survey.

Amidst intensifying climate change and increased frequency of natural disasters, the need for improved anticipatory action will only rise. Some of the AI-related challenges and opportunities we raise may appear too intractable or open-ended. We argue that since anticipatory action is already being widely used and deployed, any improvements to existing systems, however marginal, could be consequential and life-saving. Such a perspective could be used to foster a long-term vision for a future where we have better data and tools, enabling protection of the most vulnerable from disasters before they occur.

Acknowledgments

MVK is supported in part by the National Institute on Drug Abuse of the National Institutes of Health (R00DA051534).

References

- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.
- Aunan, K.; and Pan, X.-C. 2004. Exposure-response functions for health effects of ambient air pollution applicable for China—a meta-analysis. *Science of the total environment*, 329(1-3): 3–16.
- Bareinboim, E.; and Pearl, J. 2012. Transportability of causal effects: Completeness results. In *Proceedings of the*

- AAAI Conference on Artificial Intelligence, volume 26, 698–704.
- Bhagwat, A. M.; Ferryman, K. S.; and Gibbons, J. B. 2023. Mitigating algorithmic bias in opioid risk-score modeling to ensure equitable access to pain relief. *Nature medicine*, 29(4): 769–770.
- Bi, K.; Xie, L.; Zhang, H.; Chen, X.; Gu, X.; and Tian, Q. 2023. Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970): 533–538. Number: 7970 Publisher: Nature Publishing Group.
- Broniecki, P.; Leemann, L.; and Wüest, R. 2022. Improved Multilevel Regression with Poststratification through Machine Learning (autoMrP). *The Journal of Politics*, 84(1): 597–601.
- Chen, R.; Zhang, W.; and Wang, X. 2020. Machine learning in tropical cyclone forecast modeling: A review. *Atmosphere*, 11(7): 676.
- Cifuentes, J.; Marulanda, G.; Bello, A.; and Reneses, J. 2020. Air temperature forecasting using machine learning techniques: a review. *Energies*, 13(16): 4215.
- Coughlan de Perez, E.; van den Hurk, B.; Van Aalst, M.; Jongman, B.; Klose, T.; and Suarez, P. 2015. Forecast-based financing: an approach for catalyzing humanitarian action based on extreme weather and climate forecasts. *Natural Hazards and Earth System Sciences*, 15(4): 895–904.
- Cullmann, J.; Dilley, M.; Egerton, P.; Grasso, V. F.; Honoré, C.; Lúcio, F.; Luterbacher, J.; Nullis, C.; Power, M.; Rea, A.; Repnik, M.; and Stander, J. 2021. The State of Climate Services 2021: Water. Technical report, World Meteorological Organization.
- de la Poterie, A. T.; Castro Jr, E.; Rahaman, H.; Heinrich, D.; Clatworthy, Y.; and Mundorega, L. 2023. Anticipatory action to manage climate risks: Lessons from the Red Cross Red Crescent in Southern Africa, Bangladesh, and beyond. *Climate Risk Management*, 39: 100476.
- Dickerman, B. A.; and Hernán, M. A. 2020. Counterfactual prediction is not only for causal inference. *European journal of epidemiology*, 35: 615–617.
- Doocy, S.; Cherewick, M.; and Kirsch, T. 2013. Mortality following the Haitian earthquake of 2010: a stratified cluster survey. *Population health metrics*, 11(1): 1–9.
- Foini, P.; Tizzoni, M.; Martini, G.; Paolotti, D.; and Omodei, E. 2023. On the forecastability of food insecurity. *Scientific Reports*, 13(1): 2793.
- Gang, K. B. A.; O’Keeffe, J.; and Roberts, L. 2023. Cross-sectional survey in Central African Republic finds mortality 4-times higher than UN statistics: how can we not know the Central African Republic is in such an acute humanitarian crisis? *Conflict and Health*, 17(1): 21.
- Gelman, A.; Goel, S.; Rivers, D.; Rothschild, D.; et al. 2016. The mythical swing voter. *Quarterly Journal of Political Science*, 11(1): 103–130.
- Gerl, T.; Kreibich, H.; Franco, G.; Marechal, D.; and Schröter, K. 2016. A review of flood loss models as basis for harmonization and benchmarking. *PloS one*, 11(7): e0159791.
- Gevaert, C. M.; Carman, M.; Rosman, B.; Georgiadou, Y.; and Soden, R. 2021. Fairness and accountability of AI in disaster risk management: Opportunities and challenges. *Patterns*, 2(11).
- Grover, A. 2022. Rethinking Machine Learning for Climate Science: A Dataset Perspective. In *AAAI Symposium on The Role of AI in Responding to Climate Challenges*.
- Hino, M.; and Nance, E. 2021. Five ways to ensure flood-risk research helps the most vulnerable. *Nature*, 595(7865): 27–29.
- Honore, C.; Kumar, A.; and Speck, R. M. 2022. Global status of multi-hazard early warning systems: Target G. Technical report, United Nations Office for Disaster Risk Reduction.
- Huynh, B. Q.; and Basu, S. 2019. Forecasting Internally Displaced Population Migration Patterns in Syria and Yemen. *Disaster Medicine and Public Health Preparedness*, 1–6. Publisher: Cambridge University Press.
- Huynh, B. Q.; Chin, E. T.; Koenecke, A.; Ouyang, D.; Ho, D. E.; Kiang, M. V.; and Rehkopf, D. H. 2023. Potential for allocative harm in an environmental justice data tool. *arXiv preprint arXiv:2304.05603*.
- Huynh, B. Q.; Kwong, L. H.; Kiang, M. V.; Chin, E. T.; Mohareb, A. M.; Jumaan, A. O.; Basu, S.; Geldsetzer, P.; Karaki, F. M.; and Rehkopf, D. H. 2021. Public health impacts of an imminent Red Sea oil spill. *Nature sustainability*, 4(12): 1084–1091.
- Jain, P.; Coogan, S. C.; Subramanian, S. G.; Crowley, M.; Taylor, S.; and Flannigan, M. D. 2020. A review of machine learning applications in wildfire science and management. *Environmental Reviews*, 28(4): 478–505.
- Jung, J.; Concannon, C.; Shroff, R.; Goel, S.; and Goldstein, D. G. 2017. Simple rules for complex decisions. *arXiv preprint arXiv:1702.04690*.
- Kind, J.; Botzen, W. W.; and Aerts, J. C. 2020. Social vulnerability in cost-benefit analysis for flood risk management. *Environment and Development Economics*, 25(2): 115–134.
- Kishore, N.; Marqués, D.; Mahmud, A.; Kiang, M. V.; Rodriguez, I.; Fuller, A.; Ebner, P.; Sorensen, C.; Racy, F.; Lemery, J.; et al. 2018. Mortality in puerto rico after hurricane maria. *New England journal of medicine*, 379(2): 162–170.
- Li, Q.; Wang, Z.; Shangguan, W.; Li, L.; Yao, Y.; and Yu, F. 2021. Improved daily SMAP satellite soil moisture prediction over China using deep learning model with transfer learning. *Journal of Hydrology*, 600: 126698.
- Liu, J. 2012. Progresses and Challenges on QPE/QPF Utilization in Hydrology. *Tropical Cyclone Research and Review*, 1(2): 194–197.
- Mayson, S. G. 2019. Bias in, bias out. *The Yale Law Journal*, 2218–2300.
- Merz, B.; Kuhlicke, C.; Kunz, M.; Pittore, M.; Babeyko, A.; Bresch, D. N.; Domeisen, D. I.; Feser, F.; Koszalka, I.; Kreibich, H.; et al. 2020. Impact forecasting to support emergency management of natural hazards. *Reviews of Geophysics*, 58(4): e2020RG000704.

- Mosavi, A.; Ozturk, P.; and Chau, K.-w. 2018. Flood prediction using machine learning models: Literature review. *Water*, 10(11): 1536.
- Nethery, R. C.; Katz-Christy, N.; Kioumourtzoglou, M.-A.; Parks, R. M.; Schumacher, A.; and Anderson, G. B. 2021. Integrated causal-predictive machine learning models for tropical cyclone epidemiology. *Biostatistics*, kxab047.
- Novikov, G.; Trekin, A.; Potapov, G.; Ignatiev, V.; and Burnaev, E. 2018. Satellite imagery analysis for operational damage assessment in emergency situations. In *Business Information Systems: 21st International Conference, BIS 2018, Berlin, Germany, July 18-20, 2018, Proceedings 21*, 347–358. Springer.
- Obermeyer, Z.; Nissan, R.; Stern, M.; Eaneff, S.; Bembenek, E. J.; and Mullainathan, S. 2021. Algorithmic bias playbook. *Center for Applied AI at Chicago Booth*.
- Oshri, B.; Hu, A.; Adelson, P.; Chen, X.; Dupas, P.; Weinstein, J.; Burke, M.; Lobell, D.; and Ermon, S. 2018. Infrastructure quality assessment in africa using satellite imagery and deep learning. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 616–625.
- Parks, R. M.; Benavides, J.; Anderson, G. B.; Nethery, R. C.; Navas-Acien, A.; Dominici, F.; Ezzati, M.; and Kioumourtzoglou, M.-A. 2022. Association of tropical cyclones with county-level mortality in the US. *JAMA*, 327(10): 946–955.
- Paulus, J. K.; and Kent, D. M. 2020. Predictably unequal: understanding and addressing concerns that algorithmic clinical prediction may increase health disparities. *NPJ digital medicine*, 3(1): 99.
- Pearl, J. 2009. *Causality*. Cambridge university press.
- Pearl, J.; and Bareinboim, E. 2014. External Validity: From Do-Calculus to Transportability Across Populations. *Statistical Science*, 29(4): 579 – 595.
- Peters, J.; Bühlmann, P.; and Meinshausen, N. 2016. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 78(5): 947–1012.
- Pfister, N.; Bühlmann, P.; and Peters, J. 2019. Invariant causal prediction for sequential data. *Journal of the American Statistical Association*, 114(527): 1264–1276.
- Pham, K. H.; and Luengo-Oroz, M. 2022. Predictive modeling of movements of refugees and internally displaced people: Towards a computational framework. arXiv:2201.08006.
- Pople, A.; Hill, R.; Dercon, S.; and Brunckhorst, B. 2021. Anticipatory cash transfers in climate disaster response. Technical report, Centre for the Study of African Economies.
- Proma, A. M.; Islam, M. S.; Ciko, S.; Baten, R. A.; and Hoque, E. 2022. NADBenchmarks – a compilation of Benchmark Datasets for Machine Learning Tasks related to Natural Disasters. arXiv:2212.10735.
- Prosperi, M.; Guo, Y.; Sperrin, M.; Koopman, J. S.; Min, J. S.; He, X.; Rich, S.; Wang, M.; Buchan, I. E.; and Bian, J. 2020. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence*, 2(7): 369–375.
- Robinson, C.; Hohman, F.; and Dilkina, B. 2017. A deep learning approach for population estimation from satellite imagery. In *Proceedings of the 1st ACM SIGSPATIAL Workshop on Geospatial Humanities*, 47–54.
- Rolnick, D.; Donti, P. L.; Kaack, L. H.; Kochanski, K.; Lacoste, A.; Sankaran, K.; Ross, A. S.; Milojevic-Dupont, N.; Jaques, N.; Waldman-Brown, A.; et al. 2022. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)*, 55(2): 1–96.
- Rosenbaum, P. R.; and Rubin, D. B. 1983. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society: Series B (Methodological)*, 45(2): 212–218.
- Rudolph, K. E.; Schmidt, N. M.; Glymour, M. M.; Crowder, R.; Galin, J.; Ahern, J.; and Osypuk, T. L. 2018. Composition or context: using transportability to understand drivers of site differences in a large-scale housing experiment. *Epidemiology (Cambridge, Mass.)*, 29(2): 199.
- Runge, J.; Bathiany, S.; Bollt, E.; Camps-Valls, G.; Coumou, D.; Deyle, E.; Glymour, C.; Kretschmer, M.; Mahecha, M.; Muñoz-Marí, J.; et al. 2019. Inferring causation from time series in Earth system sciences, *Nat. Commun.*, 10, 2553.
- Sheth, P.; Shah, R.; Sabo, J.; Candan, K. S.; and Liu, H. 2022. STCD: A Spatio-Temporal Causal Discovery Framework for Hydrological Systems. In *2022 IEEE International Conference on Big Data (Big Data)*, 5578–5583. IEEE.
- Subbaswamy, A.; Chen, B.; and Saria, S. 2022. A unifying causal framework for analyzing dataset shift-stable learning algorithms. *Journal of Causal Inference*, 10(1): 64–89.
- Subbaswamy, A.; Schulam, P.; and Saria, S. 2019. Preventing failures due to dataset shift: Learning predictive models that transport. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 3118–3127. PMLR.
- Thalheimer, L.; Simperingham, E.; and Jjemba, E. W. 2022. The role of anticipatory humanitarian action to reduce disaster displacement. *Environmental Research Letters*, 17(1): 014043.
- Van den Homberg, M. J.; Gevaert, C. M.; and Georgiadou, Y. 2020. The changing face of accountability in humanitarianism: Using artificial intelligence for anticipatory action. *Politics and Governance*, 8(4): 456–467.
- Van der Heyden, J.; Demarest, S.; Van Herck, K.; De Bacquer, D.; Tafforeau, J.; and Van Oyen, H. 2014. Association between variables used in the field substitution and post-stratification adjustment in the Belgian health interview survey and non-response. *International journal of public health*, 59: 197–206.
- Vosper, E.; Watson, P.; Harris, L.; McRae, A.; Santos-Rodriguez, R.; Aitchison, L.; and Mitchell, D. 2023. Deep Learning for Downscaling Tropical Cyclone Rainfall to Hazard-relevant Spatial Scales. *Journal of Geophysical Research: Atmospheres*, e2022JD038163.

- Wang, L.; Chen, J.; and Marathe, M. 2019. DEFSI: Deep learning based epidemic forecasting with synthetic information. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 9607–9612.
- Weingärtner, L.; and Wilkinson, E. 2019. Anticipatory crisis financing and action: concepts, initiatives and evidence. Technical report, Centre for Disaster Protection.
- Xie, M.; Jean, N.; Burke, M.; Lobell, D.; and Ermon, S. 2016. Transfer learning from deep features for remote sensing and poverty mapping. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- Yadlowsky, S.; Namkoong, H.; Basu, S.; Duchi, J.; and Tian, L. 2022. Bounds on the conditional and average treatment effect with unobserved confounding factors. *The Annals of Statistics*, 50(5): 2587–2615.
- Yeh, C.; Perez, A.; Driscoll, A.; Azzari, G.; Tang, Z.; Lobell, D.; Ermon, S.; and Burke, M. 2020. Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. *Nature communications*, 11(1): 2583.
- Yu, J.; Liu, J.; Baek, J.-W.; Fong, C.; and Fu, M. 2022. Impact-based forecasting for improving the capacity of typhoon-related disaster risk reduction in typhoon committee region. *Tropical Cyclone Research and Review*, 11(3): 163–173.
- Zhang, Y.; Long, M.; Chen, K.; Xing, L.; Jin, R.; Jordan, M. I.; and Wang, J. 2023. Skilful nowcasting of extreme precipitation with NowcastNet. *Nature*, 619(7970): 526–532. Number: 7970 Publisher: Nature Publishing Group.
- Zhao, F.; and Zhang, C. 2020. Building damage evaluation from satellite imagery using deep learning. In *2020 IEEE 21st International Conference on Information Reuse and Integration for Data Science (IRI)*, 82–89. IEEE.
- Zigler, C. M.; and Dominici, F. 2014. Point: clarifying policy evidence with potential-outcomes thinking—beyond exposure-response estimation in air pollution epidemiology. *American journal of epidemiology*, 180(12): 1133–1140.