Early warning of complex climate risk with integrated artificial intelligence

Markus Reichstein

mreichstein@bgc-jena.mpg.de

Max Planck Institute for Biogeochemistry, Jena https://orcid.org/0000-0001-5736-1112

Dorothea Frank

Max-Planck-Institute for Biogeochemistry Jena

Vitus Benson

Max-Planck-Institute for Biogeochemistry

Gustau Camps-Valls

Universitat de Valéncia https://orcid.org/0000-0003-1683-2138

Joachim Denzler

University of Jena https://orcid.org/0000-0002-3193-3300

Kai Kornhuber

Columbia University https://orcid.org/0000-0001-5466-2059

Bernhard Schoelkopf

Max Planck Institute for Intelligent Systems https://orcid.org/0000-0002-8177-0925

Ricardo Vinuesa

KTH Royal Institute of Technology https://orcid.org/0000-0001-6570-5499

Boran Han

Amazon Web Services https://orcid.org/0000-0001-9567-7062

Carina Fearnley

University College London https://orcid.org/0000-0003-0102-1391

Nasim Rahaman

Mila - Quebec Artificial Intelligence Institute / Max Planck Institute for Intelligent Systems

José María Tarraga

Universitat de Valéncia

Jan Blunk

Friedrich Schiller University Jena https://orcid.org/0009-0001-7037-3545

Karen Dall

German Red Cross

Giulia Martini

World Food Program https://orcid.org/0000-0002-9611-2191

Naomi Nganga
Kenia Red Cross
Danielle Robinson
Amazon Web Services

Article

Keywords:

Posted Date: April 17th, 2024

DOI: https://doi.org/10.21203/rs.3.rs-4248340/v1

Additional Declarations: There is **NO** Competing Interest.

Early warning of complex climate risk with integrated artificial intelligence

Markus Reichstein[1,2,3] Vitus Benson[2,3,4] Gustau Camps-Valls[5]

Han Boran[1], Carina Fearnley[6], Kai Kornhuber[7], Nasim Rahaman[8], Bernhard Schöllkopf[8], José María Tárraga[5], Ricardo Vinuesa[9]

Jan Blunk[10], Karen Dall[11], Joachim Denzler[10,2], Dorothea Frank[3], Giulia Martini[12], Naomi Nganga[13], Danielle Maddix Robinson[1]

1 Amazon Web Services

2 ELLIS Unit Jena

3 Max-Planck-Institute for Biogeochemistry

4 ETH Zurich

5 University of Valencia

6 University College London

7 Columbia University

8 Max-Planck-Institute for Intelligent Systems

9 KTH Royal Institute of Technology

10 University of Jena

11 German Red Cross

12 World Food Program

13 Kenya Red Cross

Preface paragraph

As our climate undergoes significant shifts, both the Earth and human societies are increasingly exposed to disasters and stress. This situation underscores the critical need for robust Early Warning Systems (EWS), which are intricately designed to monitor, assess, and relay information about impending risks and hazards. EWS are vital in promoting resilient and sustainable development, yet they encounter substantial challenges in forecasting hazards and impacts, communicating risks, and in the efficiency of decision-making processes. In this perspective, we examine these challenges and explore the transformative role of integrated Artificial Intelligence

Foundation Models (AI FMs), especially focusing on the capabilities of Large Multi-Modal Models (LMMs). We discuss the power of these models in developing a Multi-Hazard Early Warning System (MHEWS), combining Meteorological and Geospatial FMs for impact prediction in a comprehensive approach. Emphasizing a user-centric strategy, this paper highlights the importance of intuitive interfaces and incorporating community feedback to enhance crisis management. Given the complex nature of climate risks, we emphasize the need for causal representations in AI models, to avoid conclusions and predictions based on spurious correlations. Additionally, we introduce the concept of decadal EWSs, which aims to provide longer term, yet spatially resolved forecasts. By leveraging climate ensembles and generative approaches, these advancements aim to provide proactive solutions for evolving climate dynamics and societal responses.

1 Early-Warning Systems for Complex Climate Risks

Early-warning systems (EWS) are an essential component of risk-reduction strategies for climate and environmental hazards and thus should be an central element of resilient sustainable-development strategies[1]. The United Nations and the World Meteorological Organization recognize the importance of these EWS and installed efforts to develop them via the Early-Warnings-for-All Initiative launched in 2022, also related to Target G of the UN Sendai Framework 2015-2030[2]. There are numerous past cases proving the value of EWSs for saving lives and livelihoods[3-5]. One key example is the investment in research and implementation in tsunami warnings following the Indian Ocean tsunami in 2004[6]. Focused collaboration has resulted in more robust, international, and technologically advanced warnings that have saved many lives since 2004, including during the 2011 Tōhoku tsunami[7] . However, it is essential to recognize that complex risks from surging climate and weather extremes involving multiple hazards as either concurrent or cascading events pose significant additional challenges to developing effective EWSs[8]. Projected increases in severity and frequency under unmitigated greenhouse gas emissions on top of changing exposure and vulnerability will make these efforts even more important for future climate risk adaptation[9], because relying on past norms and guidelines will prove inappropriate under non-stationary risks.

The United Nations[10] define EWS as *'An integrated system of hazard monitoring, forecasting and prediction, disaster risk assessment, communication and preparedness activities systems and processes that enables individuals, communities, governments, businesses and others to take timely action to reduce disaster risks in advance of hazardous events'*. Yet, current EWS tend to emphasize the hazard prediction (e.g. weather) compared to impact prediction and communication (a recent example is the German 2021 flood, where weather forecast and warning was timely and correct, but the impact not anticipated[11] and adequate preventive measures not taken. Lately, there has been more focus towards impact-based forecasts and warnings (IBFW). However, early studies suggest that these make little difference to outcomes as they still only provide the impact information rather than what actions to take in response[12]. Moreover, it has been proposed to design IBFWs for individual members of the public, which allows for a more fine-grained treatment of vulnerability and coping capacity, thus increasing the strength of future warnings[13].

The accuracy and effectiveness of EWSs depend not just on the quality of data gathered from sensors, process understanding, and the ability to predict hazards accurately and assess their potential impact, but also the speed and effectiveness of communication, and the ability to make timely and effective decisions, e.g. implemented as Anticipatory Action in the humanitarian domain[14].. All of this requires preparedness to enable the EWS to be sustainable, effective, and enable the end users to take early actions to enhance their safety and reduce economic and social losses. An important aspect for EWS is the relevant time-scale. Early warnings typically vary from seconds, to tens of thousands of years, but for climatic hazards time scales are generally on the hourly to weekly time-scale for more rapid onset hazards (e.g. storms), and have longer time-scales for slow-onset hazards (e.g. drought, desertification). Longer-time scales beyond a year are not considered on classical EWS but are very relevant for conscious spatial and infrastructure planning and societal preparedness, especially in the context of climate change[15]. The diversity of relevant aspects for early warning results in tremendous complexity and challenges in implementing effective MHEWS, while some require a large collective effort to make progress, for others, AI can offer the necessary leap forward.

2 Main challenges with Early-Warning Systems

Weather-related EWSs operate along a warning chain involving observations, forecasts (weather, hazard, impact), communication and decision making[16], and should be continuously evaluated (Figure 1). As in the case of a chain, the efficacy largely depends on the weakest link.

Hazard forecasts to a large extent rely on numerical weather prediction (NWP), which have improved tremendously over the past decades[17]. However, challenges remain: For example, with fast-onset disasters such as storms and floods, accurate forecasting of precipitation is necessary, a fact that involves resolving convection, which is computationally slow in NWP models. Hence, for such disasters, lead times can be fairly short, sometimes too short for effective action[18]. For instance, in late 2021 the tropical cyclone Rai struck the Philippines. It had undergone an unforeseen very rapid intensification in the hours before landfall, and hence early action mechanisms had not been triggered[19]. On the other hand, early warning for slow-onset disasters, e.g. droughts, builds upon sub-seasonal to seasonal forecasting[20]. At these time scales, predictability is driven by boundary conditions such as the sea-surface temperatures and the land-surface soil moisture and with the chaotic nature of the atmosphere, seasonal forecasts suffer from a lot of uncertainty (cf. Table 1), although for large-scale extremes like the recent Horn of Africa drought years some predictive skill has been achieved[21]. Yet for more localized extremes, there is seldom enough certainty on spatio-temporal extents to enable effective early action more than two weeks in advance.

Crucially, weather or hazard forecasts are not sufficient, because the same weather can have very different impacts. For instance, this was evidenced in Germany 2021 (Table 1), where a few weeks before the devastating Central European floods there was a similar meteorological event in North Eastern Germany with almost no impact[22]. The reason is a completely different landscape, which is less hilly and has more sandy soils, allowing for faster infiltration of rain. Yet predictions of impacts are challenging because they result from the interaction of the weather system with ecosystems/landscapes and societal systems[23]. This needs to consider sub km-scale, often m-scale, local context, and variables that are outside the physical climate system. In addition, even hydrodynamic models partially failed in 2021 because, complexities such as debris flow and geohydromorphological dynamics have not been considered. Moreover, societal impact forecasts need to build upon maps of exposure and vulnerability[4]. Such maps are often coarse in resolution due to a general sparsity of gridded socioeconomic data, although down-scaling attempts have yielded promising results[24]. Combining them with the physical variables obtained from the hazard forecast is not trivial, especially since dynamic aspects of vulnerabilities are not often considered[25].

Another major challenge are concurrent, compounding and cascading events, hampered by the lack of connections across the various thematic, institutional and regional silos. This is particularly critical, as amplifying cross-border effects such as impacts on supply chains, water management and disaster response capacities are important. Even more broadly, teleconnections, e.g. due to trade, river systems and atmospheric transport, are barely integrated into EWS, as their inclusion requires advanced data sharing, real-time communication, and predictive models that can account for these long-distance relations, information sources and impacts[26].

Furthermore, an ideal EWS strives to harness the full spectrum of available observations, yet present systems face notable limitations in achieving this. For instance, feeding satellite information to a physical model often requires an observational operator, as what is measured is only a proxy of what is modeled. This becomes increasingly challenging in regimes with low signal-to-noise ratio, uneven and sparse data sampling, scarcity of measurements, and wide diversity of quality, quantity and granularity of data. Hence, existing EWS rarely leverage all available data. For instance, EWS for floods and storms do not assimilate all locally available radar, gauge and satellite information, but instead focus on a few data modalities and resolutions [27]. Furthermore, potentially informative sources of data for food security EWS, e.g. from social media or economic factors, are not typically exploited in their entirety[28]

Communication of warnings, especially to the affected population, i.e. the last mile[29], is another critical aspect[30]. Numerous cases over history demonstrate that even if the EWS produces actionable forecasts, communication failed[31]. Most recently, the Mediterranean storm Daniel lead to severe rainfall and flooding in Libya with over 4,300 people dead and many more displaced[32]. While a lack of communication was certainly not the only reason for this devastating outcome, it surely contributed, given even TV weather reports predicted the landfall at least four days ahead[33] Global initiatives such as the Common Alerting Protocol have been useful to standardize warning data enabling media outlets and cell phone broadcasters to issues warnings[34]. Still, affected communities can have very different needs on EWS information, which can be best achieved by involving them already in the creation of the EWS[35], yet doing that on a global scale is hard. Here, an additional opportunity arises, often local traditional or indigenous knowledge is overseen in EWS, but can be a useful source[36], for instance when it comes to the challenge of inclusiveness. Warnings should be inclusive, not just for direct ethical reasons, but also because inclusive EWSs help to save more lives, preserve livelihoods and prevent greater economic losses, and impact longer-term equitable growth and prosperity. Designing inclusive EWSs requires to consider the peculiarities of diverse communities and their needs, ideally through their involvement from the beginning[37]. This is a challenge for current EWS, as adapting them to local conditions is costly, continuously iterating them through adaptive learning often restricted by rigid constraints and in contrast a one-size-fits-all model is cheap to implement and maintain (but not optimal for the local communities).

Last but not least, an ideal EWS should also take the expected impact of decisions based on the warnings into account, i.e. the response itself as a risk factor[38]. This can lead to a highly non-trivial decision-making feedback loop. In other words, based on the decisions that are taken upon a warning, vulnerabilities and impacts may change, which therefore would change the warning.

For example, an air pollution health impact forecast would be most useful if it considered not only the number of affected people if no action were taken, but also the same number under an evacuation plan in place. This requires the modeling of sociology and psychology, which, especially in combination with physical modeling, is challenging. Furthermore, after a disaster or an avoided disaster, the effect of interventions needs to be understood. Essentially, this requires constructing models that can operate with counterfactuals to compare "what would have happened" with "what actually happened". For example, the food security impact of droughts is often dampened by the markets through food imports. Given this mediating effect, it is challenging to estimate what the impact of an EWS and derived anticipatory action is. Here, Earth observation has been identified as useful to monitor, evaluate, account to and learn (MEAL) from anticipatory action[39], but turning such ideas into real and operational practice is still needed.

In summary, current EWSs face challenges including limitations in forecasting accuracy for fast and slow-onset disasters, difficulties in predicting impacts due to local environmental and societal variables, underutilization of diverse data sources, and challenges in effectively communicating warnings to varied communities. In addition, the complex task of incorporating societal and psychological aspects of potential warnings into the decision-making process is critical. In all of these challenges, developments in AI promise to advance the field.

3 Addressing the current limitations with existing ML: opportunities and challenges

Machine learning methods, especially deep neural networks (DNNs), have demonstrated to successfully tackle important limitations of existing EWS. So far, these efforts have nevertheless mostly focused on hazard forecasts, and have yet to trickle down the whole early warning chain.

Weather forecasting with deep neural network emulators, known as **Meteorological Foundation Models (FMs)**, has gained traction in recent years. Trained on extensive historical observations using self-supervised learning, these models predict the next time step, ensuring fast and accurate numerical weather prediction (NWP). Existing meteorological FMs may be classified broadly into two categories: global medium-range forecasting models [40-44] trained on the ERA5 3D reanalysis dataset[45] and regional precipitation nowcasting models[46-48] [49-51] trained on satellite radar observations. Meteorological FMs outperform traditional approaches on small-scale phenomena related to storms and rainfall: GraphCast[52] and PanguWeather[42] greatly reduce the tracking errors of tropical cyclones and NowCastNet[51] is skillful also for extreme precipitation events, something which was previously considered intractable. Most meteorological FMs are built on top of approaches from Computer Vision[53]: Convolutional Neural Networks or Vision Transformers [54],and recently Graph Neural Networks[41,55] are also gaining traction.

Moving forward, meteorological FMs need now take on a probabilistic perspective, which can enable the generation of large ensembles and thereby of sharp worst-case scenarios useful for early warning. They might transition to utilizing raw multimodal observations, moving away from depending on existing data assimilation for a coherent reanalysis dataset. Additionally, meteorological FMs should place more emphasis on the challenging yet important subseasonal-to-seasonal time scale for predictability. So far, most work on seasonal forecasts with ML focuses on coarse indicators. Especially the ENSO phenomenon, which greatly influences agricultural weather in Africa and South America, can be predicted well with deep neural networks [56-59]. Such phenomena are crude descriptions of the actual dynamics unfolding in the Earth system [60], which is why moving to dense meteorological FMs is both promising and necessary [21].

On land, distinguishing between hazard and impact forecasts blurs. Consider the European Floods use case: traditional hydrological models for stream flow forecasts needed basin-specific tuning. However, treating stream flow forecasts as a time series, Long Short-Term Neural Networks (LSTMs) trained on cross-basin data outperform previous methods[61,62]. Google FloodHub operationalizes this, providing flood forecasts in over 80 countries, offering early warnings to large populations[63].

Traditional EWSs issue warnings based on administrative boundaries, which may not align with the disaster's actual spatial pattern. AI enables dense spatio-temporal predictions, a fact that has exhibited important implications for example in the Horn of Africa drought use case, where humanitarians can now transition from district-level vegetation health forecasts[64] to using the EarthNet AI models [65-67], which provide predictions for individual fields and communities by using high-resolution satellite imagery. Further developments of such approaches may be called **Geospatial FMs**: They leverage the vast availability of satellite data through self-supervised learning. Here we posit that geospatial FMs, particularly those focusing on forecasting, can have a large impact on EWSs. For instance, like the EarthNet models [65-67], the EarthPT foundation model[68] can out-of-the-box predict NDVI. Hence, this allows for highly targeted triggers for anticipatory action. Similarly, wildfire risk follows localized patterns, yet current risk maps are typically provided at regional level, a fact that has been heavily criticized for instance in Greece[69]. Now, deep learning-based approaches can deliver spatially explicit maps of wildfire danger[70]. The upcoming geospatial FMs, like the Prithvi FM[71] for burned area segmentation and the Presto FM[72] for fuel moisture estimation, show promise for prediction. However, the application of generative AI for dense spatio-temporal prediction remains under explored. As lead time increases, uncertainty in drought forecasting and wildfire risk estimation grows, and diffusion models[73] may offer sharp and plausible predictions, avoiding unphysical mean predictions.

Socio-economic variables, which are crucial for understanding vulnerabilities and impacts, are often limited to coarse administrative levels and infrequent sampling intervals. Nevertheless, ML can successfully leverage them to predict drought impacts in the Horn of Africa: the WFP HungerMap utilizes XGBoost regression-tree models[74] to nowcast[75] and forecast[76] food insecurity. In socioeconomic models, interpretability is crucial to generate trust from decision makers. This has been commonly achieved through computing SHAP values[77] after model training, for instance to discover drivers of displacement[78]. Now, causal machine learning is gaining traction with recent work leveraging causal inference methods[79] to predict displacement in Somalia[80], providing insights for humanitarian response and planning. Furthermore, symbolic regression is used to discover symbolic expressions directly from data, allowing for interpretatability[81-84]. Looking ahead, transformers, like the Perceiver-IO[85] and 4M[86] models, show promise in handling multiple modalities and long-range teleconnections. These models can combine diverse data, such as county-level agricultural data and high-resolution satellite imagery, for applications like crop yield prediction[87]. Transformers also demonstrate success in leveraging circulation-driven teleconnections to forecast global wildfire risk[70], paving the way for improved EWS leveraging similar cross-boundary effects such as supply chains or streamflow.

The impact of AI on communication and the last mile challenge in EWSs remains uncertain. While it becomes challenging to envision how AI could assist vulnerable communities struggling with basic needs, there are specific communication challenges where AI may excel. Adjusting warnings to the local context, providing interactivity, and doing so in native languages are areas where AI, especially ChatBots, could be beneficial. These large language models (LLMs) are now good at in-context learning[88] meaning they can explain warnings in reasonable narratives and offer context-specific answers, as seen in flood disaster reporting[89]. Applying such models globally, especially in the global south with diverse languages, presents challenges [90]. Yet, recent work on multilingual language models[91,92] holds the potential to democratize EWS for historically under-served communities, facilitating easier communication of warnings in multiple languages at minimal cost. The next frontier lies in spoken language, where radio, identified by UN Global Pulse as a powerful tool for social listening[93], can be leveraged by speech models supporting over 1000+ languages[94] as additional data, paving the way for user-centric EWS.

AI is not a silver bullet to solve all problems – rather the usage of AI comes with its own challenges. In particular supervised ML, which seems most promising for EWS, suffers from problems with biases. Most prominently, supervised AI methods are particularly strong if the testing scenario is as close as possible to the training data — that is if there is no distribution shift between training and testing data[95]. However, in Earth-related datasets there is a sampling bias towards certain time periods and locations[96] and the climate extremes relevant for EWS are rare (though increasingly less so), hence distribution shifts are a common problem. Here, simulation[97], climate analogues[98], and space-for-time substitution[99] may be useful to obtain additional data samples and dampen distribution shifts. Furthermore, the data sets seldom contain all relevant variables (omitted variable bias) and often have data impurities, be it low sensor precision for physical variables or privacy and representativity issues with socioeconomic data. The former can be approached with multifidelity models involving data uncertainty[100], while the latter can be tackled through citizen science, that is involving volunteers in data collection, an approach that is successful in conservation ecology[101] and humanitarian mapping[102].

Data-related challenges can sometimes be even further amplified by inductive biases in DNNs. Inductive biases are modeling assumptions (e.g. locality or recurrence) that are needed to make deep learning work[103], but often cause unwanted side-effects. For instance most DNNs have a spectral bias — they prefer low-resolution features and omit high-frequency information[104]. This can lead to shortcut learning[105]: e.g. an image classification DNN may pick up non-causal features, such as using a green background to predict a cow, hence failing if a cow is in front of a blue background. Causal representation learning aims at building DNNs that are only using causal links in the data and thus capture the complex cause-and-effect relationships in the system[106]. While this area is still under active research, developing ML models and especially FMs for EWS should aim to observe causality. For instance, a causal model can evaluate the impact of interventions and thus understand the differences between the observed world and what might have occurred without humanitarian action. Additionally, causal modelsare interpretable by construction;, for example, a causal model for wildfires could disentangle complex driving relations beyond spurious correlations, and could reveal an increased wildfire danger if, in addition to high heat, there is also prolonged water stress and specific socioeconomic conditions.

Moving towards FMs for EWS, an inherent property of generative models becomes a challenge: the generation of seemingly credible yet incorrect information termed hallucination[107]. In particular, EWSs that process natural language are vulnerable to hallucinate. Recent large language models[108,109] have tackled hallucination through scale and through Reinforcement Learning from Human Feedback (RLHF)[110]. The idea being larger data sets and expert knowledge can help improve correctness. Especially for a multimodal EWS FM, domain expertise ranging from meteorology, environmental sciences to sociology and politics is necessary to help refine the accuracy and relevance of generated information. Here, increasing the explainability of FMs by outputting intermediate results sequentially in a step-by-step manner is both necessary for validation and has been shown to decrease hallucination[111,112].

As it is, DNNs are frequently described as “black boxes”, due to their property of approximating an high-dimensional function that can not be easily interpreted by humans. This may hinder their applicability in policy where trust is necessary. Especially at the beginning of the deployment of an AI model, trust can be earned if the reasoning behind model predictions can be explained[113]. For instance, if one were to study the effects of different humanitarian interventions modeled with what-if scenarios, which requires considering both direct effects and feedbacks, it is hard to imagine high trust in the predictions, if underlying reasons are not understood. So far, most approaches to explainable AI focus on post-hoc explanations[114] that are generated with a fully trained model. However, these models are not trained for generating faithful attributions, so it is not ensured that the post-hoc explanations are interpretable[115]. Hence, other ways of directly encoding interpretable parts into the model, for instance through hybrid modeling, may be more suitable[116]. An interesting approach here is building up agent-based modeling, where stakeholders are represented as independent agents interacting with each other. Each agent could be an AI, with learned behavior, allowing to model socio-economic systems directly from raw data[117,118].

On a more general note, and from a practical perspective, taking action is often hindered by political and administrative constraints, unrelated to scientific evidence and technological options. Furthermore, a lack of infrastructure and the prevalence of conflict in many of the most vulnerable areas limits the realized value of technological solutions. Even knowing these general limits, the range of opportunities for AI in EWSs is large. EWS need to be people-centric[119], considering local context and ideally work for multiple different hazards. AI can help integrate multi-modal data streams, increase prediction capabilities, enable more targeted impact forecasts and help with the communication, which is particularly challenging, as it needs to consider the different backgrounds of all users of an EWS. In order to succeed, involving stakeholders from many different disciplines, for example economists, social scientists, policy experts and psychologists, is probably a good starting point giving AI in EWS an edge over traditional approaches.

4 Vision: Foundational EWS

The early warning chain is a complex interconnected system with parts from different silos that build upon each other. Modern AI FMs may now offer the opportunity to overcome those silos and develop an integrated EWS. Such a system would be multi-hazard and multi-impact and span the whole warning chain including communication and decision making. In a first step, FMs that already exist (**Meteorological FMs** and **Geospatial FMs**), can be fine-tuned to improve individual pillars such as hazard or impact forecasting (Fig. 2). Second, they may be combined into **Impact FMs**, that can work with weather, geo-spatial and socioeconomic data to produce impact-forecasts in an integrated way. Finally, **Early Warning FMs** interface the impact model with natural language, photos and videos, hence bridging all modalities involved in the warning chain and thus achieving a fully integrated MHEWS.

Both Impact and Early Warning FMs need to work with multiple modalities. Initially, FMs were pre-trained on large and diverse datasets of single modalities, and subsequently fine-tuned for various tasks. They demonstrate remarkable "zero-shot" generalization, often accurately predicting outcomes for tasks they were not explicitly trained for. Further, the representations learned by these models enable them to adapt to new tasks using a minimal amount of additional data, in a process called fine-tuning. A wide variety of models used in natural language processing[108,109,120,121], computer vision[122-124] and speech processing[125] have provided empirical evidence for zero-shot generalization and fine-tuning across various data modalities. Recently, the pre-training approach has also seen success when combining multiple modalities[126-129], resulting in large multi-modal models (LMMs) that exhibit rich joint performance over text, image and speech. This development signifies a step towards more holistic, integrated AI systems that can process and interpret diverse data types. Instead of bridging images, text and speech, LMMs for EWS (Early Warning FMs) primarily process time-series data and maps of diverse physical and socioeconomic variables, various sensors, many resolutions, alongside natural language and photos for interactivity. Techniques like feature-wise linear modulation[130] and cross-attention[73,131] can be employed to fuse these modalities. Some of the geospatial FMs already use those methods to have first multi-modal capabilities[72].

However, such a model system for Early warning will have to obey important characteristics listed in Box 1. The design which best fits those criteria will need to be researched in the coming years. Crucially, models should be capable of simulating what-if scenarios, providing foresight into potential outcomes and enabling proactive decision-making (e.g. what happens if a floodgate gets opened, or if a village is evacuated). To this end, in addition to merely representing statistical dependencies, the models should learn causal representations that support prediction under interventions and that have been argued to be more robust with respect to ubiquitous distribution shifts[106]. Causal representations structure problems in terms of mechanisms underlying the data-generating process, and lend themselves well to building hybrid models or causal digital twins that combine components learned from data with other sources of knowledge, such as surrogate models or simulations[132]. This can be in conflict with the popular end-to-end learning, but it provides sequential “checkpoints” (e.g. prediction, then communication) which may additionally enhance model trust. Further considerations that may inform the design of those models is the level of explainability, transparency and recourse that they should provide. If decisions are taken by humans based on information provided by AI systems, it is crucial for a human to be able to understand the causes underlying a recommendation[113].

For the challenge of robustness, several training methodologies are promising. The "contrastive" learning technique is particularly prominent, aiming to harmonize data representations across varying modalities. It builds a latent space, where similar information is lumped closely together, irrespective of the underlying data modality, successfully integrating audio, images, text, and sensor data[126-129]. Alternatively, the latent space can also be build such that representations of complementary data inputs are predictive of each other, which enhances accuracy for uni-modal[133] and multi-modal[134] models alike. In parallel with advancements in large language models [108,121], reinforcement learning can be employed for further refinement. This process fine-tunes the pre-trained model through a rewards-based mechanism, focusing on the precision of its decisions and recommendations. In addition, this technique can be used to adapt the warning to the local cultural context and make it inclusive by thoughtful design under non-crises conditions, avoiding human biases readily occurring under stress[135]. Such a strategy is invaluable in crisis scenarios, leading to more effective and informed interventions.

In addition, to fully realize the potential of AI systems in EWS, they should be context-specific and designed with user-friendly interfaces that can be readily deployed by front-line humanitarian workers with limited technical expertise. Ideally, such systems should be able to function as a contributing participant in crisis meetings, providing real-time data analysis, predictive modeling, and actionable insights to inform decision-making. Hence, the resulting system in its entirety would comprise one or many FMs doing the heavy lifting, complemented with domain specific models and expertise for robustness, interpretability and intervention analysis and amended with text and image processing for user-centric communication, interactivity and active learning feedbacks (Fig. 3).

However, one key limiting factor is the collection and management of training data for these models. Unlike for text or images, there are no large, harmonized datasets for multi-modal EWS. Instead, many different sources need to be combined. While some, such as ESA's Sentinel 2 satellite imagery[136] or ECMWFs ERA5 meteorological reanalysis[45] are large publicly available datasets in analysis-ready formats, others are not. An Impact FM will require gridded and tabular socio-economic data such as crop yield or hospitalizations, which are seldom standardized across administrative boundaries and also of much smaller quantity (Gigabytes instead of Petabytes). For Early Warning FMs, alongside standard text datasets, more specialized resources such as humanitarian reports or press articles, which are scattered across the internet and not always public, need to be collected. Hence, the multi-modal systems for EWS will largely rely on pre-training with those large-scale datasets that do exist, and then careful addition of sparser data sources, e.g. using positional meta data or building upon natural language as a mediating modality.

Validation and verification of EWSs for climate hazards pose important challenges. On the one hand, FMs present unique challenges in validation due to their unprecedented versatility. Traditional AI models in geoscience are often designed for specific tasks, such as predicting rainfall patterns or detecting signs of drought from satellite imagery. These models are validated against these specific use cases. However, FMs are capable of performing a wide range of tasks, including those that might be specified by an end-user for the first time, like predicting the impact of an unforeseen climatic event. This broader scope makes it inherently more challenging to anticipate all potential failure modes. Developers and regulators will need to clearly communicate the tested use cases for which FMs warning systems are validated and caution users against 'off-label usage' that ventures into new, untested territories. The broad capabilities of such Early Warning FMs require regulatory foresight, necessitating adaptations in institutional and governmental policies, and may also influence insurance and liability frameworks. This complexity may require assessments by multidisciplinary teams including climatologists, meteorologists, environmental scientists, and other specialists, making the fact-checking process more challenging both during validation and post-deployment. On th eother hand, in order to facilitate the verification of Early Warning FMs outputs, developers should ideally incorporate explainability techniques. For instance, outputs could include references or links to underlying data sources or scientific literature that support the model's predictions. This would allow experts to more efficiently verify the accuracy and reliability of Early Warning FMs predictions. Furthermore, it is crucial for Early Warning FMs to accurately express uncertainties in their predictions to prevent overconfident and potentially misleading statements.

Last but not least, current early warning systems, which are based on impacts caused by concrete weather conditions in the next hours to weeks, should be complemented by decadal time-scale early warning systems. Developing a decadal time-scale EWS for climate and weather risks is essential due to the increasing variability and extremity of weather patterns caused by climate change. Decadal EWS should guide effective adaptation measure, more targeted than what can be inferred from general climate change metric and a general precautionary principle. This involves identifying vulnerable regions and sectors, planning infrastructure developments, and formulating policies that are resilient to long-term climatic changes. Effective communication strategies are needed to convey long-term risks and adaptations to governments, businesses, and communities, ensuring preparedness. Of course, reliable forecasts are a prerequisite, too. Hence, generally similar challenges as the ones mentioned above for short-term EWS need to be addressed, yet with an important addition: For this important challenge probabilistic forecasts are highly relevant. These forecasts present a range of possible outcomes with associated probabilities, offering a more nuanced understanding of long-term risks. In addition, there has been a trade-off between the spatial resolution of predictions and the time scale over which they are made – longer time scales typically meant broader, less detailed spatial predictions (Fig. 4). However, probabilistic AI may disrupt this norm, and generate high-resolution forecasts even for extended time scales. This can be possible, because AI can effectively model and account for the even aleatoric uncertainties inherent in long-term forecasts. This approach provides a more detailed and nuanced understanding of potential future scenarios, even at a granular spatial level. For instance, it can provide localized climate-risk assessments for specific regions or cities far into the future, something that was traditionally challenging due to the broad-brush approach required for long-term forecasts. In essence, probabilistic AI breaks the conventional link between spatial and temporal scales in forecasting, enabling more precise and detailed long-term predictions. This advancement is crucial for effective risk assessment and planning in the context of climate change and weather variability.

In summary, a multi-modal **Early Warning FM** holds promise for enhancing crisis response mechanisms. By processing and interpreting a diverse array of data types, these models can provide valuable insights and predictions in real time. To maximize their utility, it is essential to design Early Warning FMs with user-friendly interfaces that can be easily utilized by non-technical personnel. This would enable these systems to actively contribute to crisis meetings, offering real-time data analysis, predictive modeling, and actionable insights. The ability to simulate potential scenarios further enriches their contribution, offering foresight into possible outcomes and aiding in proactive decision-making. Future research should focus on the practical implementation of these models in real-world EWSs, including the development of robust evaluation metrics, and potentially benchmark tasks and simulations to assess their performance and impact. Overcoming challenges related to acquisition and management of high-quality training data will be critical to realize the potential of EWS FMs.

5 Further outlook

The future of Early Warning Systems (EWSs) lies in their ability to incorporate the response, the societal feedback, as well asconsidering long-term systemic impacts as risk factors beyond the classical hazard-exposure-vulnerability paradigm[38]. Traditional EWSs often focus on immediate responses to crises, but there is a growing recognition of the need for systems that can predict and mitigate long-term consequences, especially those that might result in misguided or counterproductive responses.

We foresee at least five important aspects that can help addressing this future challenge:

1. Agent-based approaches: This involves simulating the actions and interactions of autonomous agents (individuals, groups, or entities) to assess their effects on the system as a whole. By modeling how different agents in society might respond to warnings and crises, EWSs can predict and plan for a range of human behaviors and their potential impacts over time. This can help in understanding complex social dynamics and in designing more effective warning messages and interventions.
2. ML-Based inverse inference: Inverse ML modeling focuses on deducing the underlying causes, states or parameters from observed effects. Applied to EWSs, inverse modeling can help understand why certain crisis situations evolve as they do by analyzing the outcomes and working backward to identify the initial conditions or decisions that led to

them (avoiding assumptions like rational behavior). This can be crucial in identifying long-term trends and systemic risks that might not be apparent from direct observation.

3. Gamification and the data deluge:: Gamification involves using game-design elements in non-game contexts to engage people and encourage participation. In the context of EWSs, gamification can be used to gather data from the public, such as through apps that turn data collection into a game-like experience, e.g. with examples from previous disasters. In addition, the interactive EWS-FM chatbot can be used to gather a wealth of important, representative and high-quality data. This can lead to more comprehensive and real-time data collection, providing a richer dataset for predicting long-term effects and societal responses to warnings.
4. Integration into policy for long duration warnings: Integrating EWSs with policy-making processes for long-term planning is crucial. EWS should not only be about immediate alerts but also about providing data and insights that can guide long-term policy decisions and planning. This means developing systems that can offer insights into how different policy decisions might play out over extended periods, helping policymakers to understand the potential long-term impacts of their decisions and plan accordingly.
5. Towards collaborative efforts: Enhancing EWS requires global and interdisciplinary efforts among technical, domain knowledge, and community experts for understanding and validating the complex dynamics triggering displacement. Cooperation between national and international organizations, funding entities, and sectors like development and humanitarianism is vital for data collection, preservation, and innovative data-driven approaches. Encouraging collaboration among donors, humanitarian organizations, and academia can further transparency in cataloging predictive models, contributing to improved preparedness and response to hazards.

6 Conclusion

Integrated AI will lead to paradigm shifts in EWS. First, AI, especially Meteorological FMs, are already revolutionizing weather and hazard forecasts, leading to enormous improvements in lead times and resolution of warnings before disaster strikes. Second, multi-modal AI, materialized through Impact FMs, can leverage geospatial and socio-economic data to assess vulnerabilities and tear down previously existing silos impeding effective impact-based warnings. Third, multi-hazard EWS, soon being implemented across the globe for the UN EarlyWarningsForAll initiative, should not miss the opportunities of multi-modal AI, circumventing the rigidity of existing systems, and incorporating the whole early warning chain, including communication, into an Early Warning

FM that allows users a ChatBot-like interactive experience with warnings. Finally, causal machine learning, once mature, will enable interpretable analysis of the Early Warning FM and lead to more effective decision making, as the effects of interventions can be predicted through What-If scenarios. The success of AI in EWS crucially depends on the availability and quality of training data and the careful development of responsible, accountable and trustworthy AI methods. These are too important aspects to be left solely to the private sector, but instead should be co-developed by humanitarians and AI experts in public institutions. Only this way, using AI will democratize the access to EWS and improve the livelihoods of all people, irrespective of their backgrounds.

References:

1 Pal, I., Shaw, R., Shrestha, S., Djalante, R. & Cavuilati, R. A. W. C. *Toward sustainable development: Risk-informed and disaster-resilient development in Asia*. 1-20 (2021).

2 UNDRR & WMO. “Global status of multi-hazard early warning systems: Target G”, United Nations Office for Disaster Risk Reduction. 1-27 (UNDRR, WMO, Geneva, 2022).

3 Kelman, I. & Glantz, M. H. in *Reducing Disaster: Early Warning Systems For Climate Change* (ed Ashbindu Zommers Singh, Zinta) Ch. 5, (2014).

4 Golding, B. *Towards the “Perfect” Weather Warning*. 1 edn, (Springer Cham, 2022).

5 WMO. *Cyclone Amphan highlights the value of multi-hazard early warnings*, <https://public.wmo.int/en/media/news/cyclone-amphan-highlights-value-of-multi-hazard-early-warnings> (2020).

6 Fakhruddin, S. H. M. *Recovery from the Indian Ocean Tsunami*. 1 edn, 59-72 (2015).

7 Fujinawa, Y. & Noda, Y. Japan's earthquake early warning system on 11 March 2011: performance, shortcomings, and changes. *Earthquake Spectra* **29**, 341-368 (2013).

8 Garcia, C. & Fearnley, C. J. Evaluating critical links in early warning systems for natural hazards. *Environmental Hazards* **11**, 123-137, doi:10.1080/17477891.2011.609877 (2012).

9 Rentschler, J. *et al.* Global evidence of rapid urban growth in flood zones since 1985. *Nature* **622**, 87-92 (2023).

10 UNDRR. Global Assessment Report on Disaster Risk Reduction. 425 p. (United Nations Office for Disaster Risk Reduction (UNDRR), Geneva, Switzerland, 2019).

11 Thieken, A. H. *et al.* Performance of the flood warning system in Germany in July 2021 – insights from affected residents. *EGUsphere* **2022**, 1-26, doi:10.5194/egusphere-2022-244 (2022).

12 Potter, S. H. *et al.* The influence of impact-based severe weather warnings on risk perceptions and intended protective actions. *International Journal of Disaster Risk Reduction* **30**, 34-43, doi:10.1016/j.ijdrr.2018.03.031 (2018).

13 Potter, S., Harrison, S. & Kreft, P. The benefits and challenges of implementing impact-based severe weather warning systems: Perspectives of weather, flood, and emergency management personnel. *Weather, climate, and society* **13**, 303-314 (2021).

14 Chaves-Gonzalez, J. *et al.* Anticipatory action: Lessons for the future. *Frontiers in Climate* **4**, 932336 (2022).

15 Zommers, Z. & Alverson, K. *Resilience: The Science of Adaptation to Climate Change*. 22 (Elsevier, 2018).

16 Šakić Trogrlić, R. *et al.* in *Towards the “Perfect” Weather Warning* Ch. Chapter 2, 11-46 (2022).

17 Bauer, P., Thorpe, A. & Brunet, G. The quiet revolution of numerical weather prediction. *Nature* **525**, 47-55, doi:10.1038/nature14956 (2015).

18 Dunstone, N. *et al.* Windows of opportunity for predicting seasonal climate extremes highlighted by the Pakistan floods of 2022. *Nat Commun* **14**, 6544, doi:10.1038/s41467-023-42377-1 (2023).

19 OCHA, U. After Action Review on the CERF Anticipatory Action for Typhoons. https://reliefweb.int/attachments/af219dbc-97e3-4df2-9126-e0b30f0e4083/220422%20Report%20on%20the%20After%20Action%20Review%20for%20CERF%20Anticipatory%20Action.pdf. (2022).

20 Funk, C. & Shukla, S. *Drought early warning and forecasting: theory and practice*. (Elsevier, 2020).

21 Funk, C. *et al.* Frequent but Predictable Droughts in East Africa Driven by a Walker Circulation Intensification. *Earth's Future* **11**, e2022EF003454, doi:https://doi.org/10.1029/2022EF003454 (2023).

22 DWD. The weather in Germany 2021 - https://www.dwd.de/EN/press/press_release/EN/2021/20210830_the_weather_in_germany_in_summer_2021_news.html. (2021).

23 Reichstein, M., Frank, D., Sillmann, J. & Sippel, S. in *Climate Extremes and Their Implications for Impact and Risk Assessment* (eds Jana Sillmann, Sebastian Sippel, & Simone Russo) 341-353 (Elsevier, 2020).

24 Lloyd, C. T., Sorichetta, A. & Tatem, A. J. High resolution global gridded data for use in population studies. *Sci. Data* **4**, 1-17 (2017).

25 Jurgilevich, A., Räsänen, A., Groundstroem, F. & Juhola, S. A systematic review of dynamics in climate risk and vulnerability assessments. *Environmental Research Letters* **12**, 013002 (2017).

26 Krishna Prabhakar, S. V. R. Implications of regional droughts and transboundary drought risks on drought monitoring and early warning: a review. *Climate* **10**, 124 (2022).

27 Jain, S. K. *et al.* A Brief review of flood forecasting techniques and their applications. *International Journal of River Basin Management* **16**, 329-344 (2018).

28 Wu, D. & Cui, Y. Disaster early warning and damage assessment analysis using social media data and geo-location information. doi:10.1016/j.dss.2018.04.005 (2018).

29 De Perez, E. C. *et al.* Learning from the past in moving to the future: Invest in communication and response to weather early warnings to reduce death and damage. *Climate Risk Management* **38**, 100461 (2022).

30 Sufri, S., Dwirahmadi, F., Phung, D. & Rutherford, S. A systematic review of community engagement (CE) in disaster early warning systems (EWSs). *Progress in Disaster Science* **5**, 100058 (2020).

31 Fakhruddin, B., Clark, H., Robinson, L. & Hieber-Girardet, L. Should I stay or should I go now? Why risk communication is the critical component in disaster risk reduction. *Progress in disaster science* **8**, 100139 (2020).

32 Corps, I. M. Libya Flooding: Situation Report #9. (2023).

33 BBC. *Storm Daniel lashes Greece*, <https://www.bbc.com/weather/av/66730977> (2023).

34 Dalela, P. K., Saldhi, A., Bhave, P. & Tyagi, V. Common alerting protocol compliant emergency warning and alert system for legacy broadcasting networks. in *2020 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)* 1-5 (IEEE.

35 Kelman, I., Glantz, M. H., Singh, A. & Zommers, Z. *Reducing Disaster: Early Warning Systems For Climate Change*. 89–108 (Springer, Dordrecht, 2014).

36 Macherera, M. & Chimbari, M. J. A review of studies on community based early warning systems. *Jàmbá: Journal of Disaster Risk Studies* **8** (2016).

37 Yore, R., Fearnley, C., Fordham, M. & Kelman, I. Designing Inclusive, Accessible Early Warning Systems: Good Practices and Entry Points. (GFDRR; UCL Warning Research Centre; World Bank, 2023).
38 Simpson, N. P. *et al.* A framework for complex climate change risk assessment. *One Earth* **4**, 489-501, doi:10.1016/j.oneear.2021.03.005 (2021).
39 Enenkel, M., Dall, K., Huyck, C. K., McClain, S. N. & Bell, V. Monitoring, evaluation, accountability, and learning (MEAL) in anticipatory action—earth observation as a game changer. *Frontiers in Climate* **4**, doi:10.3389/fclim.2022.923852 (2022).
40 Kochkov, D. *et al.* Neural General Circulation Models. doi:10.48550/arXiv.2311.07222 (2023).
41 Lam, R. *et al.* Learning skillful medium-range global weather forecasting. *Science* **0**, eadi2336, doi:10.1126/science.adi2336 (2023).
42 Bi, K. *et al.* Accurate medium-range global weather forecasting with 3D neural networks. *Nature* **619**, 533-538, doi:10.1038/s41586-023-06185-3 (2023).
43 Pathak, J. *et al.* FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators. doi:10.48550/arXiv.2202.11214 (2022).
44 Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K. & Grover, A. ClimaX: A Foundation Model for Weather and Climate. *1st workshop on Synergy of Scientific and Machine Learning Modeling, SynS & ML ICML* (2023).
45 Hersbach, H. *et al.* The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* **146**, 1999-2049, doi:https://doi.org/10.1002/qj.3803 (2020).
46 Ravuri, S. *et al.* Skilful precipitation nowcasting using deep generative models of radar. *Nature* **597**, 672-677, doi:10.1038/s41586-021-03854-z (2021).
47 Gao, Z. *et al.* Earthformer: Exploring Space-Time Transformers for Earth System Forecasting. *Advances in Neural Information Processing Systems* **35**, 25390-25403 (2022).
48 Espeholt, L. *et al.* Deep learning for twelve hour precipitation forecasts. *Nature Communications* **13**, 5145, doi:10.1038/s41467-022-32483-x (2022).
49 Asperti, A. *et al.* Precipitation nowcasting with generative diffusion models. doi:10.48550/arXiv.2308.06733 (2023).
50 Shi, X. *et al.* Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems* **28** (2015).
51 Zhang, Y. *et al.* Skilful nowcasting of extreme precipitation with NowcastNet. *Nature* **619**, 526-532, doi:10.1038/s41586-023-06184-4 (2023).
52 Lam, R. *et al.* Learning skillful medium-range global weather forecasting. *Science*, eadi2336 (2023).
53 Thuemmel, J. *et al.* Inductive biases in deep learning models for weather prediction. doi:10.48550/arXiv.2304.04664 (2023).
54 Dosovitskiy, A. *et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale*, <https://openreview.net/forum?id=YicbFdNTTy> (2023).
55 Keisler, R. Forecasting Global Weather with Graph Neural Networks. doi:10.48550/arXiv.2202.07575 (2022).
56 Ham, Y.-G., Kim, J.-H. & Luo, J.-J. Deep learning for multi-year ENSO forecasts. *Nature* **573**, 568-572, doi:10.1038/s41586-019-1559-7 (2019).
57 Petersik, P. J. & Dijkstra, H. A. in *Geophys Res Lett* Vol. 47 8 (2020).
58 Dijkstra, H., Petersik, P., Hernández-García, E. & López, C. The Application of Machine Learning Techniques to Improve El Niño Prediction Skill. *Frontiers in Physics*, 13, doi:10.3389/fphy.2019.00153 (2019).
59 Cachay, S. R. *et al.* The World as a Graph: Improving El Ni\~no Forecasts with Graph Neural Networks. doi:10.48550/arXiv.2104.05089 (2021).

60 Schlör, J., Strnad, F., Capotondi, A. & Goswami, B. A multi-modal representation of El Ni\~no Southern Oscillation Diversity. doi:10.48550/arXiv.2307.11552 (2023).

61 Kratzert, F. *et al.* Towards learning universal, regional, and local hydrological behaviors via machine learning applied to large-sample datasets. *Hydrol. Earth Syst. Sci.* **23**, 5089-5110, doi:10.5194/hess-23-5089-2019 (2019).

62 Nearing, G. *et al.* AI Increases Global Access to Reliable Flood Forecasts. *arXiv preprint arXiv:2307.16104* (2023).

63 Nevo, S. *et al.* Flood forecasting with machine learning models in an operational framework. *Hydrol. Earth Syst. Sci.* **26**, 4013-4032, doi:10.5194/hess-26-4013-2022 (2022).

64 Barrett, A. B. *et al.* Forecasting vegetation condition for drought early warning systems in pastoral communities in Kenya. *Remote Sensing of Environment* **248**, 111886, doi:https://doi.org/10.1016/j.rse.2020.111886 (2020).

65 Benson, V. *et al.* Forecasting localized weather impacts on vegetation as seen from space with meteo-guided video prediction. (2023).

66 Robin, C. *et al.* in *Artificial Intelligence for Humanitarian Assistance and Disaster Response workshop at NeurIPS* (2022).

67 Requena-Mesa, C., Benson, V., Reichstein, M., Runge, J. & Denzler, J. EarthNet2021: A large-scale dataset and challenge for Earth surface forecasting as a guided video prediction task. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 1132-1142 (

68 Smith, M. J., Fleming, L. & Geach, J. E. EarthPT: a foundation model for Earth Observation. *arXiv preprint arXiv:2309.07207* (2023).

69 Zafeiropoulos, K. *Greece's map for predicting wildfires is anachronistic and inadequate*, <https://miir.gr/en/greece-s-map-for-predicting-wildfires-is-anachronistic-and-inadequate/> (2023).

70 Prapas, I. *et al.* TeleViT: Teleconnection-driven Transformers Improve Subseasonal to Seasonal Wildfire Forecasting. 6 (2023).

71 Jakubik, J. *et al.* Foundation Models for Generalist Geospatial Artificial Intelligence. *arXiv preprint arXiv:2310.18660* (2023).

72 Tseng, G., Zvonkov, I., Purohit, M., Rolnick, D. & Kerner, H. Lightweight, Pre-trained Transformers for Remote Sensing Timeseries. *arXiv preprint arXiv:2304.14065* (2023).

73 Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-Resolution Image Synthesis With Latent Diffusion Models. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 10684-10695 (2022).

74 Chen, T. & Guestrin, C. in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 785–794 (Association for Computing Machinery, San Francisco, California, USA, 2016).

75 Martini, G. *et al.* Machine learning can guide food security efforts when primary data are not available. *Nature Food* **3**, 716-728, doi:10.1038/s43016-022-00587-8 (2022).

76 Foini, P., Tizzoni, M., Martini, G., Paolotti, D. & Omodei, E. On the forecastability of food insecurity. *Scientific Reports* **13**, 2793 (2023).

77 Lundberg, S. M. & Lee, S.-I. A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems30.* (Curran Associates, Inc., 2017).

78 Ronco, M. *et al.* Exploring interactions between socioeconomic context and natural hazards on human population displacement. *Nat Commun* **14**, 8004, doi:10.1038/s41467-023-43809-8 (2023).

79 Runge, J. *et al.* Inferring causation from time series in Earth system sciences. *Nature Communications* **10**, 2553, doi:10.1038/s41467-019-10105-3 (2019).

80 Tárraga Habas, J. M., Sevillano Marco, E. & Miranda, M. T. The state-of-the-art on Drought Displacement Modelling. (iDMC internal displacement monitoring centre, 2022).

81 Brunton, S. L., Proctor, J. L. & Kutz, J. N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences* **113**, 3932-3937, doi:10.1073/pnas.1517384113 (2016).

82 Cranmer, M. Interpretable Machine Learning for Science with PySR and SymbolicRegression.jl. doi:10.48550/arXiv.2305.01582 (2023).

83 Cranmer, M. *et al.* Discovering Symbolic Models from Deep Learning with Inductive Biases. *Advances in Neural Information Processing Systems33.* 17429–17442 (Curran Associates, Inc., 2020).

84 Camps-Valls, G. *et al.* Discovering causal relations and equations from data. *Physics Reports* **1044**, 1-68, doi:10.1016/j.physrep.2023.10.005 (2023).

85 Jaegle, A. *et al.* in *ICLR 2022* (2022).

86 Mizrahi, D. *et al.* 4M: Massively Multimodal Masked Modeling. in *Neural Information Processing SystemsAdvances in Neural Information Processing Systems* (arXiv, 2023-12-11).

87 Lin, F. *et al.* in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* pp. 5774-5784 (2023).

88 Brown, T. *et al.* in *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)* (2020).

89 Colverd, G., Darm, P., Silverberg, L. & Kasmanoff, N. FloodBrain: Flood Disaster Reporting by Web-based Retrieval Augmented Generation with an LLM. in *6th Workshop on Artificial Intelligence for Humanitarian Assistance and Disaster Response (NeurIPS 2023)* (arXiv, 2023-11-05).

90 Naudé, W. & Vinuesa, R. Data deprivations, data gaps and digital divides: Lessons from the COVID-19 pandemic. *Big Data & Society* **8**, 20539517211025545, doi:10.1177/20539517211025545 (2021).

91 Ogueji, K., Zhu, Y. & Lin, J. Small Data? No Problem! Exploring the Viability of Pretrained Multilingual Language Models for Low-resourced Languages. in *MRL 2021Proceedings of the 1st Workshop on Multilingual Representation Learning* (eds Duygu Ataman *et al.*) 116–126 (Association for Computational Linguistics, 2021-11).

92 Wei, X. *et al.* PolyLM: An Open Source Polyglot Large Language Model. doi:10.48550/arXiv.2307.06018 (2023).

93 Luccioni, A. S., Pham, K. H., Lam, C. S. N., Aylett-Bullock, J. & Luengo-Oroz, M. Ensuring the Inclusive Use of NLP in the Global Response to COVID-19. *Machine Learning and Principles and Practice of Knowledge Discovery in Databases* (eds Michael Kamp *et al.*) 259-266 (Springer International Publishing, 2021).

94 Pratap, V. *et al.* Scaling Speech Technology to 1,000+ Languages. doi:10.48550/arXiv.2305.13516 (2023).

95 Schneider, S. *et al.* Improving robustness against common corruptions by covariate shift adaptation. *Advances in neural information processing systems* **33**, 11539-11551 (2020).

96 Meyer, H. & Pebesma, E. Machine learning-based global maps of ecological variables and the challenge of assessing them. *Nature Communications* **13**, 1-4 (2022).

97 Kekić, A. *et al.* Evaluating vaccine allocation strategies using simulation-assisted causal modeling. *Patterns* **4**, 100739, doi:https://doi.org/10.1016/j.patter.2023.100739 (2023).

98 Hallegatte, S., Hourcade, J.-C. & Ambrosi, P. Using climate analogues for assessing climate change economic impacts in urban areas. *Climatic Change* **82**, 47-60, doi:10.1007/s10584-006-9161-z (2007).

99 Pickett, S. T. A. in *Long-Term Studies in Ecology: Approaches and Alternatives* (ed Gene E. Likens) 110-135 (Springer New York, 1989).

100 Morimoto, M., Fukami, K., Maulik, R., Vinuesa, R. & Fukagata, K. Assessments of epistemic uncertainty using Gaussian stochastic weight averaging for fluid-flow regression. *Physica D: Nonlinear Phenomena* **440**, 133454, doi:10.1016/j.physd.2022.133454 (2022).

101 Silvertown, J. A new dawn for citizen science. *Trends in Ecology & Evolution* **24**, 467-471, doi:10.1016/j.tree.2009.03.017 (2009).

102 Herfort, B., Lautenbach, S., Porto de Albuquerque, J., Anderson, J. & Zipf, A. The evolution of humanitarian mapping within the OpenStreetMap community. *Scientific Reports* **11**, 3037, doi:10.1038/s41598-021-82404-z (2021).

103 Goyal, A. & Bengio, Y. Inductive biases for deep learning of higher-level cognition. *Proceedings of the Royal Society A* **478**, 20210068 (2022).

104 Ackaah-Gyasi, K. N., Valdez, S., Gao, Y. & Zhang, L. Exploring Spectral Bias in Time Series Long Sequence Forecasting. (2023).

105 Geirhos, R. *et al.* Shortcut learning in deep neural networks. *Nature Machine Intelligence* **2**, 665-673 (2020).

106 Schölkopf, B. *et al.* Toward Causal Representation Learning. *Proceedings of the IEEE* **109**, 612-634 (2021).

107 Ji, Z. *et al.* Survey of hallucination in natural language generation. *ACM Computing Surveys* **55**, 1-38 (2023).

108 OpenAI *et al.* GPT-4 Technical Report. doi:10.48550/arXiv.2303.08774 (2023).

109 Chowdhery, A. *et al.* PaLM: Scaling Language Modeling with Pathways. *Journal of Machine Learning Research* **24**, 1-113 (2023).

110 Ouyang, L. *et al.* Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* **35**, 27730-27744 (2022).

111 Dhuliawala, S. *et al.* Chain-of-verification reduces hallucination in large language models. *arXiv preprint arXiv:2309.11495* (2023).

112 Pan, L. *et al.* Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies. *arXiv preprint arXiv:2308.03188* (2023).

113 Vinuesa, R. & Sirmacek, B. Interpretable deep-learning models to help achieve the Sustainable Development Goals. *Nature Machine Intelligence* **3**, 926-926 (2021).

114 Samek, W., Montavon, G., Lapuschkin, S., Anders, C. J. & Müller, K.-R. Explaining deep neural networks and beyond: A review of methods and applications. *Proceedings of the IEEE* **109**, 247-278 (2021).

115 Ronco, M. & Camps-Valls, G. Role of locality, fidelity and symmetry regularization in learning explainable representations. *Neurocomputing* **562**, 126884, doi:https://doi.org/10.1016/j.neucom.2023.126884 (2023).

116 Reichstein, M. *et al.* Deep learning and process understanding for data-driven Earth system science. *Nature* **566**, 195-204, doi:10.1038/s41586-019-0912-1 (2019).

117 Zheng, S., Trott, A., Srinivasa, S., Parkes, D. C. & Socher, R. The AI Economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science advances* **8**, eabk2607 (2022).

118 Park, J. S. *et al.* Generative agents: Interactive simulacra of human behavior. in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* 1-22 (

119 Percy, F., Lourdes, M. & Macasil, K. GUIDANCE DOCUMENT on People-Centered Risk-Informed Early Warning Systems. 48 (2023).

120 Radford, A., Narasimhan, K., Salimans, T. & Sutskever, I. Improving Language Understanding by Generative Pre-Training. (2018).

121 Touvron, H. *et al.* Llama 2: Open Foundation and Fine-Tuned Chat Models. doi:10.48550/arXiv.2307.09288 (2023).

122 Dosovitskiy, A. *et al.* An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. in *International Conference on Learning Representations* (2020/10/02).

123 Kirillov, A. *et al.* Segment Anything. in *Proceedings of the IEEE/CVF International Conference on Computer Vision* 4015-4026 (2023).

124 Oquab, M. *et al.* DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research*, doi:10.48550/arXiv.2304.07193 (2023).

125 Radford, A. *et al.* Robust Speech Recognition via Large-Scale Weak Supervision. in *International Conference on Machine LearningProceedings of the 40th International Conference on Machine Learning* 28492-28518 (PMLR, 2023-07-03).

126 Radford, A. *et al.* Learning Transferable Visual Models From Natural Language Supervision. in *International Conference on Machine LearningProceedings of the 38th International Conference on Machine Learning* 8748-8763 (PMLR, 2021-07-01).

127 Girdhar, R. *et al.* ImageBind: One Embedding Space To Bind Them All. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 15180-15190 (2023).

128 Alayrac, J.-B. *et al.* Flamingo: a Visual Language Model for Few-Shot Learning. *Advances in Neural Information Processing Systems* **35**, 23716-23736 (2022).

129 Liu, H., Li, C., Wu, Q. & Lee, Y. J. Visual Instruction Tuning. in *Neural Information Processing SystemsAdvances in Neural Information Processing Systems* (arXiv, 2023-12-11).

130 Perez, E., Strub, F., Vries, H. d., Dumoulin, V. & Courville, A. FiLM: Visual Reasoning with a General Conditioning Layer. *Proceedings of the AAAI Conference on Artificial Intelligence* **32**, doi:10.1609/aaai.v32i1.11671 (2018).

131 Zhang, Y. & Yan, J. Crossformer: Transformer Utilizing Cross-Dimension Dependency for Multivariate Time Series Forecasting. in *The Eleventh International Conference on Learning Representations* (2022/09/29).

132 Kekić, A. *et al.* Evaluating vaccine allocation strategies using simulation-assisted causal modeling. *Patterns* **4**, doi:10.1016/j.patter.2023.100739 (2023).

133 Oord, A. v. d., Li, Y. & Vinyals, O. Representation Learning with Contrastive Predictive Coding. doi:10.48550/arXiv.1807.03748 (2019).

134 Assran, M. *et al.* Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture. in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* 15619-15629 (2023-06).

135 Yu, R. Stress potentiates decision biases: A stress induced deliberation-to-intuition (SIDI) model. *Neurobiology of stress* **3**, 83-95 (2016).

136 Louis, J. *et al.* SENTINEL-2 SEN2COR: L2A Processor for Users. in *ESA Living Planet Symposium 2016Proceedings Living Planet Symposium 2016SP-740.* (ed L. Ouwehand) 1-8 (Spacebooks Online, 2016-08).

137 Mohr, S. *et al.* A multi-disciplinary analysis of the exceptional flood event of July 2021 in central Europe. Part 1: Event description and analysis. *NHESSD 2, 5079–5111, 2014 Seasonal forecasting of fire over Kalimantan, Indonesia A. C. Spessa et al. Title Page Abstract Introduction Conclusions References Tables Figures J I J I Back Close Full Sc* **2022**, 1-44 (2022).

138 Ludwig, P. *et al.* A multi-disciplinary analysis of the exceptional flood event of July 2021 in central Europe–Part 2: Historical context and relation to climate change. *Natural Hazards and Earth System Sciences* **23**, 1287-1311 (2023).

139 Koks, E., Van Ginkel, K., Van Marle, M. & Lemnitzer, A. Brief communication: Critical infrastructure impacts of the 2021 mid-July western European flood event. *NHESSD 2, 5079–5111, 2014 Seasonal forecasting of fire over Kalimantan, Indonesia A. C. Spessa et al. Title Page Abstract Introduction Conclusions References Tables Figures J I J I Back Close Full Sc* **2021**, 1-11 (2021).

140 Mohr, S. *et al.* A multi-disciplinary analysis of the exceptional flood event of July 2021 in central Europe – Part 1: Event description and analysis. *Nat. Hazards Earth Syst. Sci.* **23**, 525-551, doi:10.5194/nhess-23-525-2023 (2023).

141 Cornwall, W. Europe's deadly floods leave scientists stunned. *Science* **373**, 372-373, doi:doi:10.1126/science.373.6553.372 (2021).

142 Economist, T. Canada's wildfires have burnt an area 16 times larger than normal. https://www.economist.com/the-americas/2023/06/15/canadas-wildfires-have-burnt-an-area-16-times-larger-than-normal. (2023).

143 Kelly, J. Record breaking PM2.5 pollution levels in NYC in early June 2023 regular occurrence in over 350 cities worldwide. https://energyandcleanair.org/record-breaking-pm2-5-pollution-levels-in-nyc-in-early-june-2023-regular-occurrence-in-over-350-cities-worldwide/#:~:text=Record%20breaking%20PM2.-,5%20pollution%20levels%20in%20NYC%20in%20early%20June%202023,in%20over%20350%.). (2023).

144 https://www.nyc.gov/assets/doh/downloads/pdf/eode/eode-air-quality-impact.pdf, N. Y. C. D. o. H. a. M. H. Air Pollution and the Health of New Yorkers. 40.

145 CEMS. *FAQ on EFAS and the recent flood events*, <https://www.efas.eu/en/news/faq-efas-and-recent-flood-events> (2021).

146 Kochkov, D. *et al.* Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences* **118**, e2101784118 (2021).

147 Robin, C. *et al.* preprint - Learning to forecast vegetation greenness at fine resolution over Africa with ConvLSTMs. (2022).

148 Balashankar, A., Subramanian, L. & Fraiberger, S. P. Predicting food crises using news streams. *Science Advances* **9**, eabm3449 (2023).

<u>Key references</u>

Brown 2020 **This paper introduced GPT-3, the first large language model capable of in-context learning.**

Funk 2020 **This book collects the state-of-the-art in drought early warning without ML.**

Geirhos 2020 **This paper introduced the concept of shortcut learning which need to be tackled through careful evaluation for deep neural networks to generalize well.**

Golding 2022 **This book collects the state-of-the-art in EWS for weather-related fast-onset disasters without ML.**

Lam 2023 **This paper introduced GraphCast, the first deep neural network beating the state-of-the-art medium-range weather forecasting model IFS.**

Martini 2022 **This paper introduced the models behind the WFP HungerMap, the first fully ML-based vulnerability assessment of food insecurity.**

Nevo 2022 **This paper introduced the Google FloodHub, which is the first global initiative leveraging machine learning for an operational EWS.**

Prapas 2023 **This paper introduced TeleViT, a transformer model that is able to leverage teleconnections to improve wildfire early warning.**

Radford 2021 **This paper introduced CLIP, the first image-text model doing multi-modal learning with a contrastive loss function.**

Schölkopf 2021 **This paper introduced Causal Representation Learning, an new paradigm for machine learning, where the latent space is not solely based on correlations, but on causation.**

Box: Essential Criteria for AI systems in effective early warning systems

1) Accuracy and Timeliness:

   Scientific Challenge: Developing AI algorithms that can accurately and quickly detect and predict multiple hazards, considering their complex and evolving nature. Aspects of data latency (real-time) are similar as for classical EWS (and addressed in observation and analysis systems)

2) Scalability and Integration:

   Scientific Challenge: Designing AI systems that can efficiently handle the large volume of data from diverse sources and integrate seamlessly with existing data-acquisition and communication networks. The multi-modal Early Warning FM approach (Fig. 2) is poised to address this. In addition, federated learning is emerging as a powerful technique for training models in decentralized environments with non-iid data, such as Earth system science datasets. It addresses challenges of bandwidth, cost, and data privacy by enabling collaborative model training while keeping data decentralized, offering advantages for handling heterogeneous data, optimizing non-iid data, and accommodating diverse environments. Consider the scalability of AI to provide warnings to regions or event individual needs.

3) Adaptability, Reliability, and Robustness:

   Scientific Challenge: Building AI models and systems that can adapt to changing conditions, while ensuring reliability and robustness in the face of technical failures, network disruptions, or adverse environmental conditions. Developing algorithms that can dynamically update models based on new data and continuously learn from evolving hazards and human response patterns. Few-shot learning capabilities of FMs via in-context learning without re-training offer promise. Implementing redundancy and backup mechanisms, along with fault-tolerant architectures, can enhance system reliability and robustness. Additionally, exploring anomaly detection and anomaly response techniques can help identify and address system failures or abnormal behavior. As a back up, it would be vital to also explore what would happen if technology fails and automated actions do not trigger as expected (accommodating failure). Unprecedented versatility of EWS-FM present unique challenges for validation.

4) Transparency, Explainability, and Ethical Considerations:

   Scientific Challenge: Developing AI models and algorithms that are transparent, explainable, and avoid biases in decision-making processes, ensuring ethical and fair early warning systems (human-value-based AI). Employing explainable AI techniques, such as rule-based systems or interpretable machine-learning models, to make the decision-making process more transparent and understandable. Incorporating fairness metrics and conducting regular audits can help identify and mitigate biases in the data, algorithms, and decision outputs. Developing systems that account for inclusiveness, and for those that have little data (asylums, migrants, persecuted etc.). Leveraging EWS-FMs ability to present complex climate impact data in accessible formats.

5) User-Centric Design and Accessibility:

Scientific Challenge: Designing AI-based early warning systems that are user-friendly and accessible to all individuals subject to Early Warning, including those with diverse language preferences, disabilities, and educational and socio-economic backgrounds. Applying user-centered design principles and conducting user studies to understand the needs, capabilities, and limitations of different user groups. Developing multi-modal interfaces that cater to diverse communication channels, such as visual, auditory, or tactile, can enhance accessibility. Additionally, leveraging natural language processing techniques to support multiple languages and incorporating assistive technologies can improve inclusivity.

6) <u>Stakeholder Collaboration and Post-Event Analysis:</u>

Scientific Challenge: Facilitating effective collaboration among stakeholders and conducting comprehensive post-event analysis to understand the long-term consequences of hazards. Employing data analytics and visualization techniques to analyze post-event data, including damage assessments, recovery needs, and socio-economic impacts. Implementing collaborative platforms and information-sharing mechanisms can enhance stakeholder engagement and knowledge exchange. Additionally, conducting interdisciplinary research and co-producing with social scientists, economists, and urban planners to provide holistic insights for post-event analysis and future mitigation strategies.

Table 1: Overview of early warning short-comings for past events and how AI could help to address the short-comings in the future.

	EUROPEAN FLOODS 2022[137-139]	HORN OF AFRICA DROUGHT 2016-2017 AND 2019FF	MIDWEST HEAT-FIRE-SMOKE DISASTER 2023
GENERAL SITUATION	o July 2021 heavy rainfall in Germany, Belgium, France and UK [140] o Flash floods afterwards [141] o Biggest impact in small river Ahr: 196 dead	o Five failed rainfall seasons 2019-2023 o Devastating multi-year drought in Somalia, Ethiopia and Kenya o Over 36 Mio people affected (__REF__ https://reliefweb.int/report/ethiopia/horn-africa-drought-regional-humanitarian-overview-call-action-revised-24-august-2022)	o Very warm and dry spring o Wildfires in Canada [142] o June 2023 smoke travels to midwest: record-shattering airpolution in NYC[143] o Large impacts on Health[144]
EXISTING EWS	o EFAS european flood alerts were send out[145] o Warning chain included neighboring counties and the media	o FEWSNet food security classification forecast based on ENSO forecasts o SPI forecasts from ICPAC & VCI forecasts from RCMRD o Anticipatory action based on those triggers + vulnerability data, e.g. by Kenya Red Cross on County level	o Airnow.gov airquality forecast for 2 days in advance o Communication to public very late o Measures taken were canceling of public events etc.
OBSERVATIONS & WEATHER FORECAST	Shortcomings: River Gauges were damaged by the high water levels Weather forecast was already quite good 2 days in advance, not the main limiting factor Role of AI and challenges: Global Meteorological FM offer orders of magnitude higher speed allowing larger ensembles and thus better probabilistic estimates. Meteorological FM radar map forecasts are better at heavy rainfall events than weather models[14] Video gauges (not in contact with the water) AI calibrated to predict water levels and flows may be more robust (minor point), but calibration needed	Shortcomings: Lack of hydrometeorological stations in parts of the area (e.g. Somalia) => insufficient precipitation data Biases in spatialized precipitation estimates (fusing station and satellite data) Lack of food insecurity data, i.e. vulnerability Because of longer lead times and the superposition of teleconnections (ENSO, IOD) forecasts are uncertain Local-scale forecasts for districts or villages a major challenge confidence levels are stated as one of the reasons the early decision making was delayed [23]. Because of inherent uncertainties, acting on seasonal forecasts is currently estimated to take up to a decade before a FbF system would generate value (due to false alarm cost etc)[12] Role of AI and challenges: Geospatial foundation models may lead to better interpolation and fusion of satellite and station data by maximising data input and via generalization and transfer from other areas. This may include new anti-causal learning approaches. For instance, precipitation can be inferred anti-causally from soil moisture observation. Similarly, Martini et al.[75] showed that vulnerability can be estimated from proxy data, using a global machine learning-	Shortcomings: Citizen science, especially cell phone data, is challenging to use in existing EWS. Weather forecast necessary for circulation, but ultimately it is the transport of smoke, that is relevant for air quality. Thus need integration of weather forecast with fire forecast for air quality forecasting. Role of AI and challenges: AI-based methods to integrate sparse and low quality sensor data, e.g. human mobility patterns to estimate health effects. AI-based airquality forecasts can leverage teleconnections —> longer lead times Meteorological FMs may be extended to include Chemical Transport Modeling, which offers order of magnitude speed ups in simulation of the dynamics of wildfire smoke.

		based model trained on primary data outside the area of consideration. Meteorological FMs may be extended towards both, seasonal forecast and local downscaling – research needed on both teleconnections and hyperlocal influence of soils and vegetation.	
HAZARD & IMPACT FORECAST	Shortcomings: Flood levels were not precisely and locally enough predicted, because of insufficient resolution and negligence of debris flow and morpho-dynamic processes. Impact forecast non-existent Detailed forecasts for smaller river basins were missing Role of AI and challenges: Machine learning–accelerated computational fluid dynamics[146] can overcome computational limitations of hydro-morpho-dynamic modelling. AI-based stream flow forecasts to capture local circumstances and debris flow & scale to ungauged basins AI guided forensic analysis of exposure and vulnerability using multi-modal approaches, local fine-tuning of geospatial foundation models. Limitations include knowledge of high-resolution morphology including bottlenecks such as bridges or channels and stochasticity of debris flow, and societal data for vulnerability assessment. Space-for-time generalization acceptable?	Shortcomings: Current EWS are centered on weather only focus on single variables, e.g. rainfall, not compound or cascading events relatively coarse-grained information, struggling in low-information and high-risk areas like pastoralist regions during 2016 drought, timing of onset biased Impacts such as vegetation or crop conditions not explicitly addressed, often based on simple hazard thresholds, not accounting for ecological and socio-economic conditions Need for holistic multi-hazard approach, including complex cascading effects (drought, fire, flood, locust outbreaks) Role of AI and challenges: Localized impact forecasts help increase the efficiency of existing funding. Directly leveraging high resolution Earth observation to map impacts on vegetation and harvests conceivable, e.g. with the EarthNet models[147]. Text information successfully used as additional feature to predict food crises[148] including socio-economic conditions. Language Foundation Models work without a notion of “scale”. For Hazard & Impact forecasting need to integrate data across many scales. Multi-modal transformer models like Perceiver IO[15] hint at one possible avenue, but range of scales treated there is limited compared to scales in Disaster Early Warning chain. Impact foundation models can improve prediction skills by integrating, weather, land surface and socio-economic variables	Shortcomings: Wildfire forecasts need to take into account memory effects, such as a previously very dry season, for fuel estimation. Airquality forecasts suffer from a lack of integration of teleconnections. There was room for predictability from the wildfire season in Canada, which was not leveraged in the midwest, as airquality warnings had only been issued 2 days in advance. PM2.5 values are still quite abstract for decision makers and the general public, hence could go one step further and make impact forecasts on health and economics. Role of AI and challenges: Wildfire modeling with recurrent neural networks and transformers to account for long-range memory effects Transformers for leveraging teleconnections, such as the circulation-driven pattern Impact foundation model for integrating weather -> socio-economic -> health But: This requires branching across Silos
WARNING COMMUNICATION, DECISION AND LONGER-TERM ADAPTATION	Shortcomings: Warnings were issued 1-2 days in advance, but with numbers (mm of rain, flood levels), which are not always effective with respect to making timely decisions[13] Role of AI and challenges: Conditioned on flood levels and digital elevation models, expected inundation areas and expected damages can be visualized with (generative) AI based maps and photo-realistic	Shortcomings: Issues with respect to interpretation of warning levels regarding the food security situation because of their abstract and coarse nature, resulting in disagreements over the appropriate response and course of action, and even leading to delays in funding and disengagement of some donors Status-quo bias in decision making with a tendency to resist action Lack of accessibility and transparency in early warning reports, mainly available in English rather than local languages, using complex scientific jargon, and neglecting information sharing with drought-affected communities,	Shortcomings: Silo thinking lead to communication of warning very late, and to top down decisions such as last minute closing of events. Health impacts are often related to insufficient protection of vulnerable people, which could be improved by better communication. Role of AI and challenges: Early warning foundation model for communication to local decision makers, event organizers etc.

	representations. In addition, generation of language based (written or audio for visually impaired) warnings are indicated. ChatBot for interactivity in these communications. Situation room: Early warning foundation model to enable local decisionmakers to have detailed access to relevant knowledge	limited their usefulness in helping vulnerable populations take necessary precautions Quantifying the impact of the interventions missing Identification of the root causes of vulnerability Insufficient finance, especially 2021-2023 Whole region is considered "high alert", but lack of granularity <u>Role of AI and challenges:</u> AI based forecasts of impacts (see above) allow more tangible, interpretable and fine-grained forecasts of food insecurity (and related issues like WASH or migration). Causal analysis allows for estimating effectively measure cause-effect estimations and thresholds of impacts the effect of interventions (more research needed). Causal representation learning to discover causal pathways and enable what-if modeling for planning & evaluation of anticipatory action AI can assist in optimizing the dissemination of warnings based on factors like geographical location, population density, and vulnerability indices, ensuring that the right information reaches the right people at the right time. Large-language models may efficiently translate warnings to natural language including local languages and inclusivity. Early warning foundation models for informing front-line humanitarian workers in their local language —> enabling interactivity + democratization	Longer-term AI-based risk assessments (decadal early warning) to understand the distribution of heat-fire-smoke events in the future and enable adaptation, especially through better infrastructure and education

Figures and tables

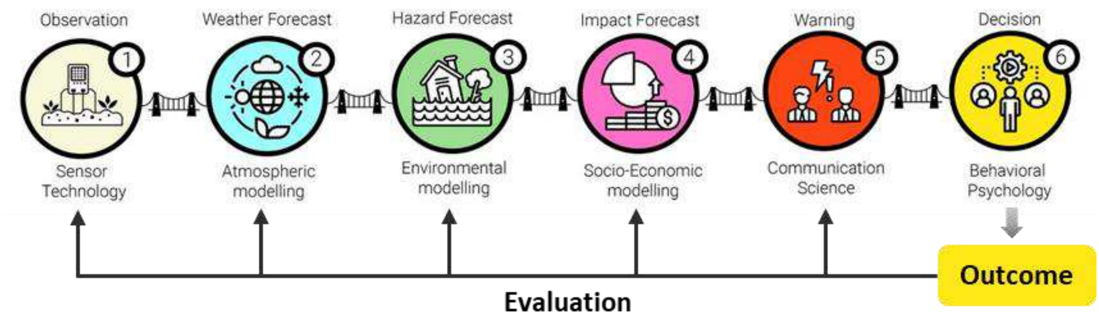

Fig. 1: **Early-warning chain from observation to decision.** "All six bridges of death" (pers. comm. Brian Golding, HiWeather) have to be crossed for an effective early warning. Figure modified after ref [16]

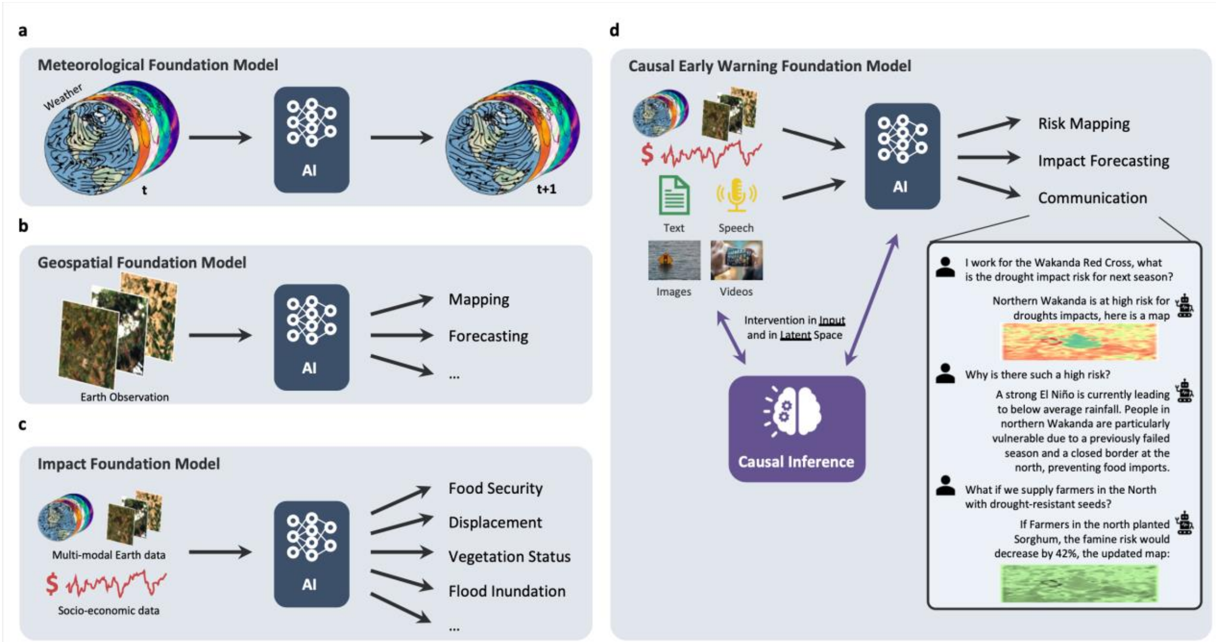

Fig. 2: **From domain specific to cross-domain foundation models for early warning.** While an Impact FM (c) integrates Meteorological FM (a) and Geospatial FM (b) elements, an effective Early Warning FM (d) needs to integrate unstructured human information in addition and needs to be integrated with a causal inference approach.

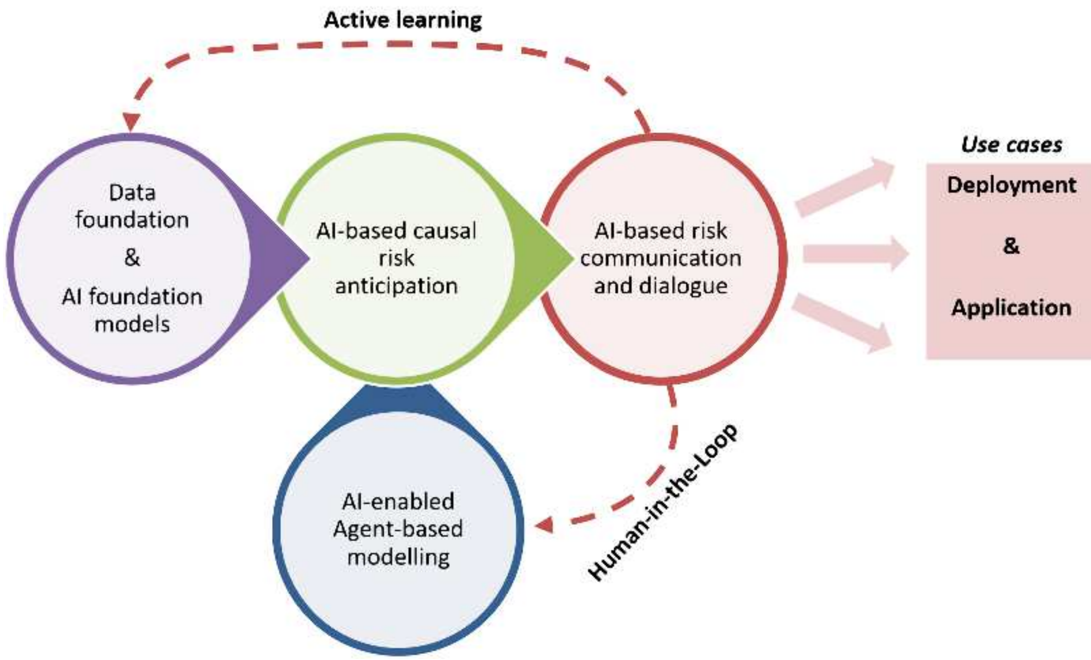

Fig. 3: **Integrative, AI-enabled strategy for Early Warning of complex climate risks including an interactive component**. The Early Warning FM leads to improved causal and data informed risk anticipation, followed by AI-based communication. Anticipating disaster Response as a risk factor[38] should be addressed via Agent-based modelling embracing AI for parameter estimations. Information from the user-interaction should feed back to the model improvement.

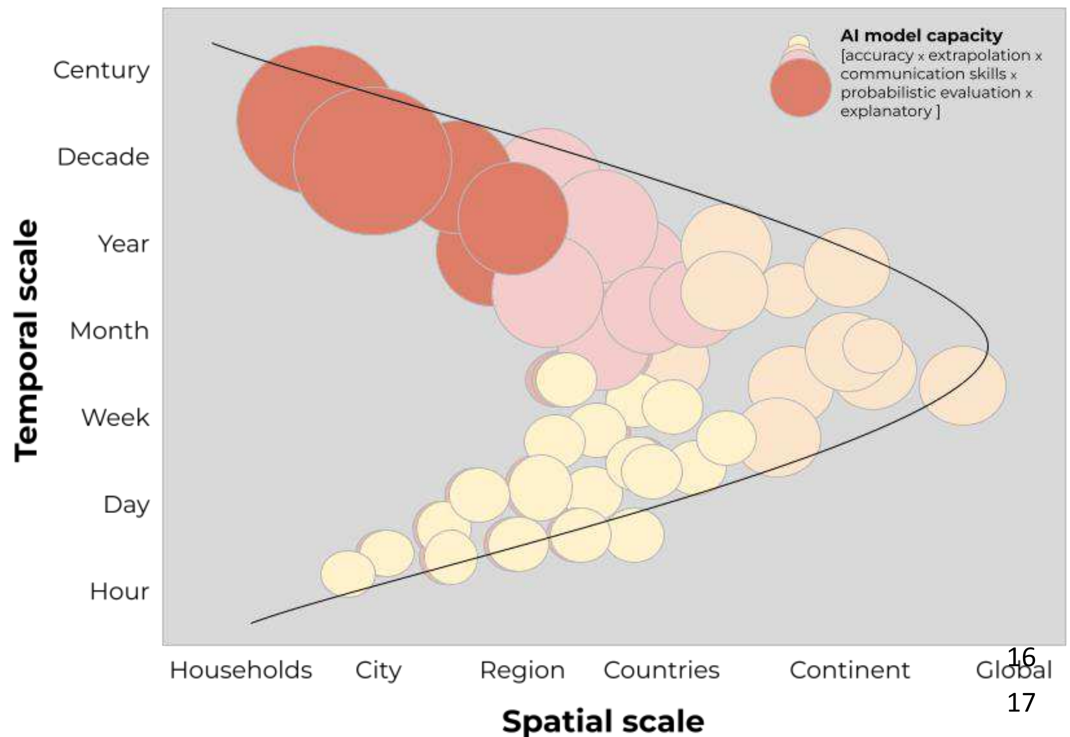

Fig. 4: **Spatial and temporal scales to be addressed with Early Warning systems.** Probabilistic models can "escape" the typical correlation of temporal and spatial scale and can make and communicate more local probabilistic risk assessments also at long time-scales. However, this also poses new challenges for robustness, explainability and communication.