
ML for Flood Forecasting at Scale

Sella Nevo¹ Vova Anisimov¹ Gal Elidan² Ran El-Yaniv³ Pete Giencke⁵
Yotam Gigi² Avinatan Hassidim⁴ Zach Moshe¹ Mor Schlesinger⁵ Guy Shalev¹
Ajai Tirumali¹ Ami Wiesel² Oleg Zlydenko⁵ Yossi Matias^{1,5}

Abstract

Effective riverine flood forecasting at scale is hindered by a multitude of factors, most notably the need to rely on human calibration in current methodology, the limited amount of data for a specific location, and the computational difficulty of building continent/global level models that are sufficiently accurate. Machine learning (ML) is primed to be useful in this scenario: learned models often surpass human experts in complex high-dimensional scenarios, and the framework of transfer or multitask learning is an appealing solution for leveraging local signals to achieve improved global performance. We propose to build on these strengths and develop ML systems for timely and accurate riverine flood prediction.

Floods are the most common and deadly natural disaster in the world. Every year, floods cause from thousands to tens of thousands of fatalities [5, 21, 4, 22, 15], affect hundreds of millions of people [15, 22, 4], and cause tens of billions of dollars worth of damages [5, 4]. These numbers have only been increasing in recent decades [23]. Indeed, the UN charter notes floods to be one of the key motivators for formulating the sustainable development goals (SDGs), and directly challenges us: "They knew that earthquakes and floods were inevitable, but that the high death tolls were not." [3]

Early warning systems for floods, even with limited lead time and accuracy, can reduce both fatalities and economic damages by over a third, and in some cases almost by half [2, 26, 1]. Unfortunately, the majority of human costs that are due to flooding are concentrated in developing countries [15], which often lack effective and actionable early warning systems due to limited data collection, funding, or professional expertise [28]. The result is that across multiple countries, thousands die on average every year, and relief and mitigation efforts have very limited information to rely on.

We focus on riverine floods which are responsible for much of the effect on human life, and for which we believe the opportunity for ML-based predictive models is substantial. Perhaps the biggest challenge in scaling current hydrologic modeling efforts to global scale is the parameter calibration process - an optimization process aimed at matching model outputs to ground-truth measurements. The standard existing methodology includes significant manual work, which is arduous and time consuming [11, 16], difficult to teach [11], and is often over-parameterized [10]. The result is that models that appear to work well at construction time, are often not able to generalize because of the complex and dynamic nature of the problem [8]. Indeed, there have been many attempts to automate this process, yet these efforts have not yet been able to reach a level of accuracy and reliability sufficient for use in operational systems [11, 24]. Practical use has thus been limited to either semi-automatic processes which include manual intervention or interaction, [11, 17], or to aggressive constraining of the model that leads to inferior results [9, 16].

The field of Machine Learning (ML) is incredibly well-placed to make a difference since many of the core challenges that stand in the way of automatic calibration of hydrologic models are central to the ML arsenal. In fact, most large-scale ML systems these days replaced and improved on earlier inferior system that were built and calibrated by hand. There have been previous efforts at using ML

¹ Google Research ² Google Research and The Hebrew University ³ Google Research and Technion - Israel Institute of Technology ⁴ Google Research and Bar-Ilan University ⁵ Google Crisis Response

for flood prediction but these have been typically limited to a handful of basins or even a single one [27, 29, 14, 18, 12, 20, 19, 13]. More often than not, the learned model were not able to compete with manual calibration or achieve sufficient reliability for operational use. We believe that substantially better models can be learned by utilizing data at a larger scale. We also believe that recent advances in inductive transfer, transfer learning and multitask learning (see [25] for a recent survey) can overcome the difficulty of learning effective models from multiple related sources.

One of the main reasons the potential of ML for flood prediction has not yet been realized is that there are few organizations (or even communities) that have all the necessary components for success under one roof: knowledge and experience in operational flood forecasting; the ability to centralize data at a global scale; the required expertise in ML; the ability to effectively deploy predictive information.

In an attempt to change this, we developed and recently deployed an operational hydro-dynamic model, in a pilot area around Patna, India [6], with the explicit goal of preparing the ground for integrating ML models into the process. In this model, the physical processes are relatively well understood for several decades now [31], and relatively little calibration is required. It receives real-time measurements and short-term forecasts of river water levels in real-time, and produces an inundation map estimating the flood extent. Based on past events, our model's flood extent estimates are accurate down to a spatial resolution of 300 meters, achieving over 90% recall and over 75% precision. These results have been borne out by our first alerts produced in real time for the 2018 monsoon season. While our initial results are promising for this "simple" component of flood prediction, we did run into challenges we believe machine learning is well placed to solve, including high computational costs of the physics-based simulations, inaccuracies due to error-prone and biased inputs, and manual calibration that still dominates the amount of work required for each new location.

The above launched hydro-dynamic model is just one component of a full flood forecasting system, and the best understood one at that. We believe ML can improve the quality of multiple components, including ones that are not well simulated by physics-based models, from absorption and snow-melt models to river discharge estimation. Toward making this possible, we are collecting, organizing and combining open data sets from different sources to make this problem more accessible for the ML community. Together with the launched hydro-dynamic model, we believe this sets the ground for the development, assessment and deployment of ML models at the global, continent or country scale.

Concretely, we believe ML can provide large contributions to operational flood forecasting systems along each of the following axes:

ML-based hydro-dynamic modeling Computational costs still limit the scope and the resolution of hydro-dynamic models. These costs could potentially be reduced by up to orders of magnitude by integrating ML-based derivation methods, as has been experimentally done for other fluid dynamics systems (see e.g. [7]).

Remote discharge estimation Limited data is perhaps the most substantial obstacle for incorporating ML into the hydrology field. Gauges measuring discharge and water levels are relatively scarce, on the order of 100,000 globally. Yet the quantity, quality and variety of satellites constantly imaging all rivers is rising at a rapid pace. Improving our ability to estimate river discharge *without* in-situ measurements is a task that is both critical for the hydrology field, and incredibly well suited for ML. See [30] for our preliminary work in this area.

ML hydrologic models Data-driven hydrology models often under-perform, mainly due to lack of sufficient data: for each basin there are typically only hundreds to thousands of measurements. Yet, hydrologic phenomena follows the same physical principles everywhere, and it is reasonable to assume that by jointly considering the multiple tasks of flood prediction at multiple locations we can do better. We conjecture that by utilizing recent advances in transfer learning, we can do just that and effectively learn at a global scale.

Hybrid physics-ML hydrologic models An alternative to the pure ML approach above, is to adopt a hybrid approach where an ML model is integrated with classical physics-based components, which can contain important prior knowledge about the structure of the domain. Intuitively, one could view this approach as allowing the conceptual components to capture the majority of the hydrologic modeling, while the ML model is responsible for calibration, error corrections, and perhaps additional processes that were not well modelled.

In summary, we believe ML is particularly suited to improve effective flood prediction at scale, and we are excited to catalyze more academic and non-academic efforts in this direction.

References

- [1] World Bank - Global Assessment Report on Costs and Benefits of Early Warning Systems. https://www.preventionweb.net/english/hyogo/gar/2011/en/bgdocs/Rogers_&_Tsirkunov_2011.pdf, 2011. [Online; accessed 30-09-2018].
- [2] World Health Organization (WHO) - Global Report on Drowning. http://www.who.int/violence_injury_prevention/global_report_drowning/Final_report_full_web.pdf, 2014. [Online; accessed 30-09-2018].
- [3] United Nations Development Programme (UNDP) - Sustainable Development Goals. <https://shop.undp.org/pages/the-sustainable-development-goals>, 2015. [Online; accessed 30-09-2018].
- [4] United Nations Office for Disaster Risk Reduction (UNISDR) - The Human Cost of Weather Related Disasters. https://www.unisdr.org/2015/docs/climatechange/COP21_WeatherDisastersReport_2015_FINAL.pdf, 2015. [Online; accessed 30-09-2018].
- [5] The Centre for Research on the Epidemiology of Disasters (CRED) - Natural Disasters 2017. https://cred.be/sites/default/files/adsr_2017.pdf, 2017. [Online; accessed 30-09-2018].
- [6] Keeping people safe with AI-enabled flood forecasting. <https://www.blog.google/products/search/helping-keep-people-safe-ai-enabled-flood-forecasting/>, 2018. [Online; accessed 30-09-2018].
- [7] Y. Bar-Sinai, S. Hoyer, J. Hickey, and M. P. Brenner. Data-driven discretization: a method for systematic coarse graining of partial differential equations. *arXiv preprint arXiv:1808.04930*, 2018.
- [8] K. Beven. Changing ideas in hydrology—the case of physically-based models. *Journal of hydrology*, 105(1-2):157–172, 1989.
- [9] K. J. Beven. *Rainfall-runoff modelling: the primer*. John Wiley & Sons, 2011.
- [10] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [11] D. P. Boyle, H. V. Gupta, and S. Sorooshian. Toward improved calibration of hydrologic models: Combining the strengths of manual and automatic methods. *Water Resources Research*, 36(12):3663–3674, 2000.
- [12] M. Campolo, P. Andreussi, and A. Soldati. River flood forecasting with a neural network model. *Water resources research*, 35(4):1191–1197, 1999.
- [13] X. Chen, K.-W. Chau, and A. Busari. A comparative study of population-based optimization algorithms for downstream river flow forecasting by a hybrid neural network model. *Engineering Applications of Artificial Intelligence*, 46:258–268, 2015.
- [14] C. Dawson and R. Wilby. Hydrological modelling using artificial neural networks. *Progress in physical Geography*, 25(1):80–108, 2001.
- [15] S. Doocy, A. Daniels, C. Packer, A. Dick, and T. D. Kirsch. The human impact of earthquakes: a historical review of events 1980-2009 and systematic literature review. *PLoS currents*, 5, 2013.
- [16] K. Eckhardt and J. Arnold. Automatic calibration of a distributed catchment model. *Journal of hydrology*, 251(1-2):103–109, 2001.
- [17] H. V. Gupta, S. Sorooshian, and P. O. Yapo. Status of automatic calibration for hydrologic models: Comparison with multilevel expert calibration. *Journal of Hydrologic Engineering*, 4(2):135–143, 1999.

- [18] K.-I. Hsu, H. V. Gupta, X. Gao, S. Sorooshian, and B. Imam. Self-organizing linear output map (solo): An artificial neural network suitable for hydrologic modeling and analysis. *Water Resources Research*, 38(12):38–1, 2002.
- [19] K.-I. Hsu, H. V. Gupta, and S. Sorooshian. Artificial neural network modeling of the rainfall-runoff process. *Water resources research*, 31(10):2517–2530, 1995.
- [20] A. Jain and A. M. Kumar. Hybrid neural network models for hydrologic time series forecasting. *Applied Soft Computing*, 7(2):585–592, 2007.
- [21] S. Jonkman. Loss of life caused by floods: an overview of mortality statistics for worldwide floods. *DCI-233-6*, 2003.
- [22] S. N. Jonkman. Global perspectives on loss of human life caused by floods. *Natural hazards*, 34(2):151–175, 2005.
- [23] T. Loster. Flood trends and global change. In *Proceedings IIASA Conf on Global Change and Catastrophe Management: Flood Risks in Europe*, 1999.
- [24] H. Madsen, G. Wilson, and H. C. Ammentorp. Comparison of different automated strategies for calibration of rainfall-runoff models. *Journal of Hydrology*, 261(1-4):48–59, 2002.
- [25] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 2010.
- [26] P. J. Pilon et al. Guidelines for reducing flood losses. In *Guidelines for reducing flood losses*. Naciones Unidas, 1998.
- [27] D. L. Shrestha and D. P. Solomatine. Machine learning approaches for estimation of prediction interval for the model output. *Neural Networks*, 19(2):225–235, 2006.
- [28] D. Strömberg. Natural disasters, economic development, and humanitarian aid. *Journal of Economic perspectives*, 21(3):199–222, 2007.
- [29] A. S. Tokar and P. A. Johnson. Rainfall-runoff modeling using artificial neural networks. *Journal of Hydrologic Engineering*, 4(3):232–239, 1999.
- [30] A. Wiesel, A. Hassidim, G. Elidan, G. Shalev, M. Schlesinger, O. Zlydenko, R. El-Yaniv, S. Nevo, Y. Matias, Y. Gigi, and Z. Moshe. MI for flood forecasting at scale. In *Proceedings of the NIPS AI for Social Good Workshop*.
- [31] W. Wu. *Computational river dynamics*. CRC Press, 2007.