SCC-IRG Track 2 - Project Overview:
Transfer Learning for Fair and Scalable Flood Management in Smart Connected Cities

1 Overview

The U.S. National Climate Assessment [14] predicts that the intensity and frequency of extreme precipitation events will increase by 15% over the midwestern United States. Flood damage costs in the US were nearly $17 billion/year between 2010 and 2018 [1] and this flooding will certainly worsen under future climate scenarios [4, 9]. There is, therefore, a critical need for tools that midwestern cities can use to mitigate the flooding risks created by these extreme storms. This project's objective is to develop those tools for smart cities [5] that use Information, Communication, and Control (ICC) technologies to meet its citizens' needs.

This project uses an embedded sensor network called CSOnet [13] to meet its objectives. CSOnet embeds Internet-connected sensors and actuators into a city's underground wastewater (sewer) network to reduce combined sewer overflow (CSO) events during rain storms. This is done through real-time monitoring and control of the wastewater flows in the sewer pipes. Successful commercialization [18] of CSOnet technology has led to CSOnet deployments in a number of U.S. cities. This project will work with these CSOnet-enabled cities to train deep learning models that predict the likelihood of surface and basement flooding in each city neighborhood. These models would then be used for the real-time control and long-term management of urban flooding.

This project addresses technological and social science challenges that face smart cities using deep learning [10] to enhance their resilience to extreme climate. The main *technological challenges* involves training models that can be used to control wastewater flows and whose knowledge can be easily transferred between cities of all sizes. That challenge will be overcome using transfer learning [17] in which a pre-trained flood model is retuned using real-time data from the city's CSOnet system. The main *social science challenge* involves training models that can recommend which neighborhoods should receive wastewater infrastructure (I/F) improvements, and to do so in a manner that is perceived as fair with respect to a sensitive attribute such as neighborhood racial or socio-economic makeup. That challenge will be overcome using models that employ a *fair learning representation layer* [19] to statistically decorrelate neural network decisions from a sensitive neighborhood attribute.

2 Technical Approach

The project's technological challenge is concerned with how deep learning can be used to predict and control exigent flooding due to on-going storms using knowledge that is easily and inexpensively transferred between different cities. Deep neural networks outperform physics-based simulation models in flood forecast accuracy [3, 6, 21]. A recent model architecture (U-FLOOD) [11] has been shown to generalize to catchments not present in its training data [7]. This model has an encoder-decoder architecture in which the encoder downscales images of the urban landscape onto a lower dimensional set of internal features, before upscaling those features back onto a flood map. Our hypothesis is that the outputs of the U-FLOOD encoder form *fundamental features* that are critical for generating the flood maps of *any city*. If true, this would mean that a given city's flood model can be based on a single U-FLOOD encoder. So this project's flood models will have the architecture shown in Fig. 1 where the U-FLOOD encoder acts as a pre-trained *model base* used by any city, followed by a *model head* that is retrained using data from the given city. In particular, we propose training the model head using real-time sensor data from the city's CSOnet system. The model head is a sequence-to-sequence transformer that outputs a sequence of future sequence of CSOnet sensor data. That prediction would then be used for the real-time control of wastewater sewer flows to mitigate the threat of exigent flooding due to an on-going storm.

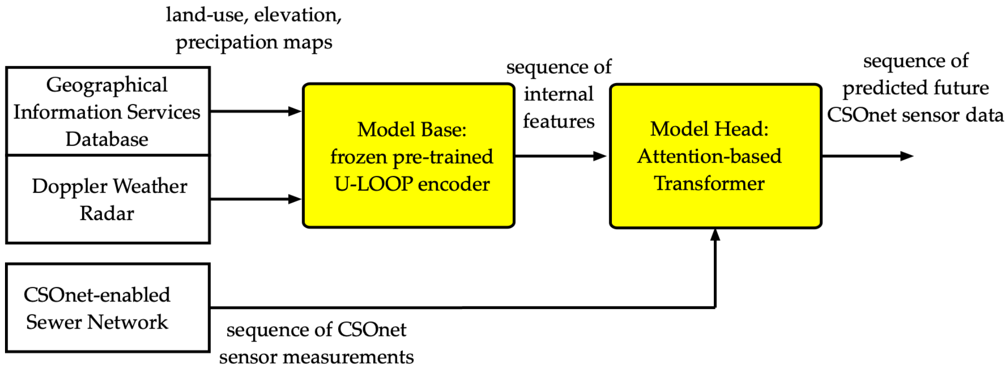

Figure 1: Flood Model Architecture

The project's social science challenge is concerned with how deep learning can be used to *improve wastewater infrastructure (I/F) in a manner that is fair with respect to sensitive neighborhood attributes such as race and socio-economic makeup.*. Deep neural networks can use historical data to train a model that predicts how human decision makers might act in response to future inputs. Our hypothesis is that these deep neural network models can be retrained to be *fair* in a well-defined statistical sense. In particular, we will use the model architecture shown in Fig. 2. The model takes a 2D tensor whose x-direction itemizes distinct city neighborhoods and whose y axis itemizes the neighborhood attributes commonly used in decided whether a given neighborhood's wastewater I/F will be upgraded. We would work with our city partners to identify that attribute list, but we expect it to include things such as neighborhood real estate values, population, types of commercial businesses, flood map, state of existing I/F, crime rate, etc. The model would then map that tensor of neighborhood attributes onto two distinct target outputs; the likelihood that a given neighborhood has its wastewater I/F upgraded and the likelihood that the neighborhood is classified as being disadvantaged with respect to a sensitive attribute such as race or socio-economic standing. This project will work with the small business entity (HydroDigital LLC), owning the IP, for CSOnet to assess this approach's potential for knowledge transfer between cities and to assess how well it performs in mitigating flooding threats due to extreme climate in the two target cities (SBN and KC).

The model architecture in Fig. 2 consists of a *fair representation layer*, a *policy critic*, and a *equity critic*. Following [19], the fair representation layer outputs a set of lower dimensional features measured with respect to past policy decisions and the sensitive neighborhood attribute. These feature vectors are then used by two different classifier models; a *policy critic* that outputs the most likely decision made by past human decision makers and an *equity critic* that outputs the most city's most likely classification of each neighborhood with respect to the sensitive attribute. Fig. 2 shows that we train this model in two steps. The first step trains the model on historical data to set up the fair representation layer and critics. In the second step, we freeze the weights of the representation layer and equity critic and then retrain the policy critic with a loss function that has been regularized with respect to a fairness metric such as statistical parity or equal opportunity [2]. The project will work with municipal decision makers to evaluate how well this deep learning approach to fair decision making addresses the equity issues faced by our two target cities (SBN

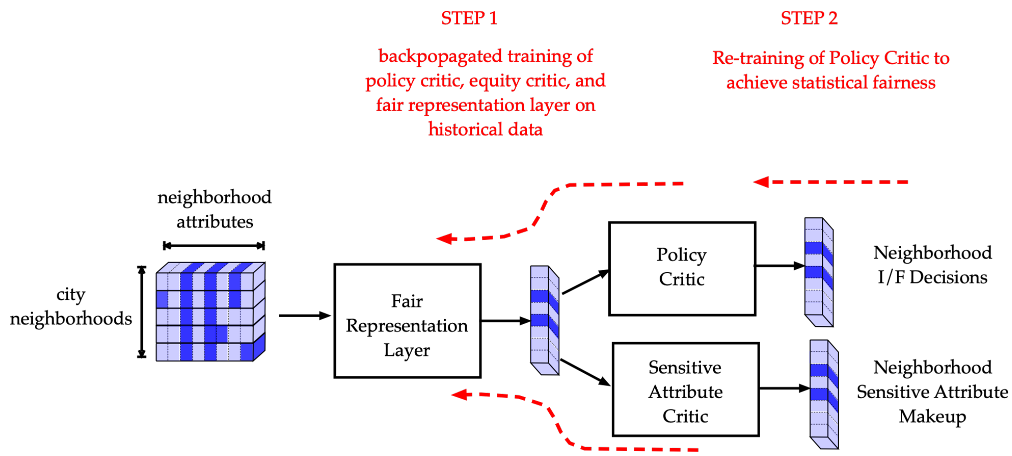

Figure 2: fig:Retraining I/F upgrade decisions for statistical fairness

and KC) when they develop flood abatement policies.

3 Integrative Research Plan

This project will follow a research plan that integrates the technological and social sciences work being performed to answer the research questions raised by this project. The project's work will be performed by four teams of scientists and engineers; deep learning researchers led by Prof. Michael Lemmon at the University of Notre Dame (UND), embedded system engineers led by Dr. Luis Montestruque at HydroDigital LLC, and municipal engineers/decision makers in two midwestern US cities. One of the city teams will be from South Bend, Indiana (SBN) where the first CSOnet site was built in 2009. We have tentatively chosen the second city team be from Kansas City, Missouri (KC).

3.1 Research Questions

The research questions to be addressed by the proposed work are enumerated below.

1. **Neural Network Knowledge Transfer and Scalability:** How do we take the knowledge learned by neural network models trained on physics-based simulations and *transfer that knowledge between cities of varying size* to ensure the control and management outcomes achieved using these models are comparable to outcomes obtained using models trained from scratch on these cities?

2. **Datasets and Model Architectures for Flood Control/Management:** What type of *datasets and deep learning models* are needed and how should these models be used in developing *real-time control and long-term management policies* that mitigate the threat of urban flooding due to extreme precipitation?

3. **Statistical Fairness in Developing Smart City Flood Abatement Policies:** Which *neighborhood attributes* and what notions of *statistical fairness* are of greatest value to municipal decision makers when they formulate *flood abatement policies* that equitably benefit all city neighborhoods regardless of racial or socio-economic makeup?

4. **Smart City Resilience to Climate-driven Pluvial Flooding:** What *blueprints can smart city engineers follow* in using these deep learning methods to enhance their city's *resilience to pluvial flooding* that will occur because of the more intense and frequent precipitation events expected under climate change?

These four research questions will be answered in an integrative manner through six inter-related project tasks. The relationship between these tasks is shown in Fig. 3. These tasks are discussed in more detail below.

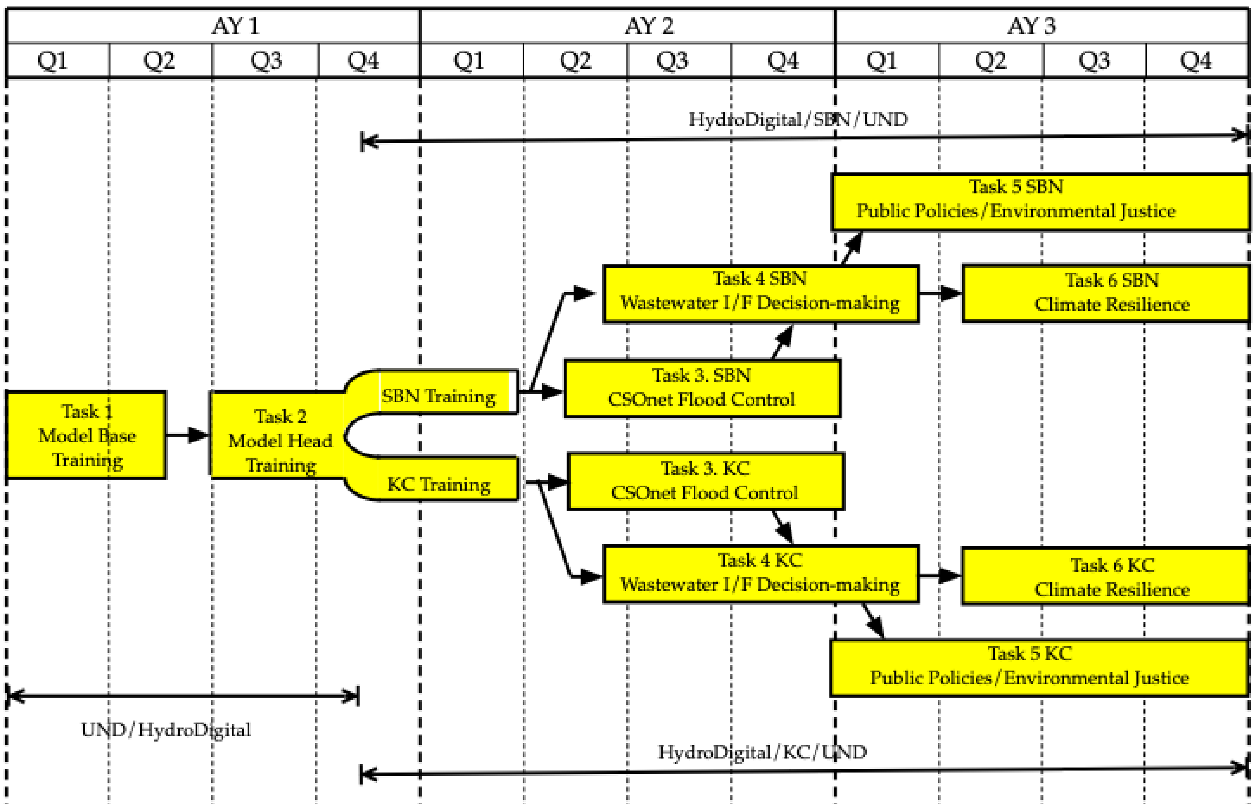

Figure 3: Integrative Research Plan

3.2 Task 1 - Pre-training Model Base from Hydrological Simulation

This task will be performed primarily by the UND and HydroDigital team with some data assistance from SBN and KC. UND will perform the training of the model from training data that has been generated by a 2D hydrological simulation (CA-ffe) for a region similar in size to SBN. UND will generate a number of urban landscapes and HydroDigital will provide a number of storm scenarios (precipitation maps). SBN and KC will provide historical data regarding land-use, terrain elevation, and prior flood history. The main part of this task involves training a U-FLOOD model using training data generated by the CA-ffe simulation model. UND will evaluate how well the trained model generalizes to the two cities, SBN and KC. The main deliverable will be the U-FLOOD encoder, which will be used in Task 2.

3.3 Task 2 - Retraining Model Head for Specific Cities

The first half of this task be performed by UND and HydroDigital and will involve developing the transformer model head and its training procedures. After that, the task bifurcates into two parts to be conducted primarily by the city teams (SBN and KC). In this second part, the city teams will use the flood model developed by UND/HydroDigital in the first two tasks to train models specific to their cities. UND will assess the accuracy of these two city models. The main deliverable will the the SBN and KC models to be used in task 3.

3.4 Task 3 - Real-time CSOnet Flood Control for On-going Storms:

Task 3 will be conducted primarily by HydroDigital for both the SBN and KC models. UND will act as consultant to HydroDigital for this task. The city teams will work with HydroDigital to integrate the models from task 2 with each city's CSOnet system. Both cities will monitor how the resulting system performs during their rainy season to evaluate how well the CSOnet system reduces the risk of localized urban flooding in each city. The main deliverable will be a study assessing whether to model developed in Task 1 and 2 can indeed be transferred between different cities.

3.5 Task 4 - Neural Network Modeling of Long-term Wastewater I/F Decision-Making

This task will be performed by UND with information regarding neighborhood attributes and historical infrastructure decisions in the two cities (SBN and KC). UND will use the data provided by the cities to train a model that predicts the past decision making history of both cities. The main deliverable from this task will be the decision making model in Fig. 2 for each project city along with an evaluation of the model's prediction accuracy.

3.6 Task 5 - Public Policies for Environmental Justice

This task will be performed by UND with information from SBN/KC regarding how each city classifies their neighborhoods with regard to racial and socio-economic makeup. UND will develop methods for retraining their model's policy critic to achieve a desired level of statistical fairness. UND will work with both cities to evaluate the extent to which each city sees the various notions of statistical fairness helping them achieve environmental equity for their mixture of neighborhoods. The main deliverable from this task will be a study assessing the degree to which statistical notions of fairness can be used to assess environmental justice.

3.7 Task 6 - Evaluating Climate Resilience for Project Smart Cities

This task will be performed by SBN and KC to define how they measure climate resilience with respect to flooding and to evaluate the extent to which the public policies arising out of tasks 4 and 5 enhance each city's resilience to those extreme precipitation events. UND and HydroDigital will work together to develop a blueprint or plan by which the results of this project can be easily disseminated to other midwestern cities of varying size. The main deliverables will be the blueprint as well as the assessment of how well this project enhanced each cities climate resilience.

References

[1] ASFPM (association of State Floodplain Managers). A cost analysis for completing and maintaining the nation's nfip flood map inventory. ASFPM (assocaition of State Floodplain Managers), 2020.

[2] Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness in machine learning. *Nips tutorial*, 1:2017, 2017.

[3] Dieu Tien Bui, Nhat-Duc Hoang, Francisco Martínez-Álvarez, Phuong-Thao Thi Ngo, Pham Viet Hoa, Tien Dat Pham, Pijush Samui, and Romulus Costache. A novel deep learning neural network approach for predicting flash flood susceptibility: A case study at a high frequency tropical storm area. *Science of The Total Environment*, 701:134413, 2020.

[4] Heejun Chang, Arun Pallathadka, Jason Sauer, Nancy B Grimm, Rae Zimmerman, Chingwen Cheng, David M Iwaniec, Yeowon Kim, Robert Lloyd, Timon McPhearson, et al. Assessment of urban flood vulnerability using the social-ecological-technological systems framework in six us cities. *Sustainable Cities and Society*, 68:102786, 2021.

[5] Mark Deakin and Husam Al Waer. From intelligent to smart cities. *Intelligent Buildings International*, 3(3):140–152, 2011.

[6] Asmamaw Gebrehiwot, Leila Hashemi-Beni, Gary Thompson, Parisa Kordjamshidi, and Thomas E Langan. Deep convolutional neural network for flood extent mapping using unmanned aerial vehicles data. *Sensors*, 19(7):1486, 2019.

[7] Zifeng Guo, Vahid Moosavi, and João P Leitão. Data-driven rapid flood prediction mapping with catchment generalizability. *Journal of Hydrology*, 609:127726, 2022.

[8] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

[9] Kenneth E Kunkel, Thomas R Karl, Michael F Squires, Xungang Yin, Steve T Stegall, and David R Easterling. Precipitation extremes: Trends and relationships with average precipitation and precipitable water in the contiguous united states. *Journal of Applied Meteorology and Climatology*, 59(1):125–142, 2020.

[10] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[11] Roland Löwe, Julian Böhm, David Getreuer Jensen, Jorge Leandro, and Søren Højmark Rasmussen. U-flood–topographic deep learning for predicting urban pluvial flood water depth. *Journal of Hydrology*, 603:126898, 2021.

[12] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.

[13] Luis Montestruque and M Lemmon. CSONET: a metropolitan scale wireless sensor-actuator network. In *International Workshop on Mobile Device and Urban Sensing (MODUS)*, 2008.

[14] DR Reidmiller, Christopher W Avery, David R Easterling, Kenneth E Kunkel, KLM Lewis, Thomas K Maycock, and Brooke C Stewart. Fourth national climate assessment. *Volume II: Impacts, Risks, and Adaptation in the United States*, 440, 2018.

[15] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.

[16] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information*

processing systems, 30, 2017.

[17] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016.

[18] A. Zaleski. One city's fight to solve its sewage problem with sensors. *MIT Technology Review*, May, 2021.

[19] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International conference on machine learning*, pages 325–333. PMLR, 2013.

[20] Yu Zhang, Peter Tiňo, Aleš Leonardis, and Ke Tang. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(5):726–742, 2021.

[21] Gang Zhao, Bo Pang, Zongxue Xu, Dingzhi Peng, and Depeng Zuo. Urban flood susceptibility assessment based on convolutional neural networks. *Journal of Hydrology*, 590:125235, 2020.