Infrastructure Resilience for Climate Adaptation

Amrita Gupta*
School of Computational Science and Engineering
Georgia Institute of Technology
Atlanta, Georgia
agupta375@gatech.edu

Caleb Robinson*
School of Computational Science and Engineering
Georgia Institute of Technology
Atlanta, Georgia
dcrobins@gatech.edu

Bistra Dilkina
Department of Computer Science
University of Southern California
Los Angeles, California
dilkina@usc.edu

ABSTRACT

Developing and maintaining resilient transportation infrastructure is a key strategy for meeting several UN sustainable development goals in the face of climate change-driven extreme flooding events. We present a framework for performing data-driven vulnerability analysis for flooding on existing transportation networks, and use this analysis to inform decision-making about investments for climate adaptation. We apply this approach to study the potential impacts of severe flooding on regional mobility in Senegal, using a combination of flood hazard maps and a travel demand model based on call detail record data. We use the estimated number of infeasible trips as a direct measure of flooding-induced mobility impacts, as well as an objective for minimizing these impacts. We then compare three alternative road network upgrade strategies to assess the extent to which each strategy would preserve network functionality under a given flooding scenario. We illustrate that strategies driven solely by travel demand can lead to underinvestment in roads that are at risk of flooding, while solely focusing on repairing flooded road segments neglects the criticality of those repairs to mobility. For example, in a 100 year flooding scenario with a fixed budget, our strategy that considers both flooding and mobility data can achieve a 53% reduction in the number of infeasible trips, while a strategy that just considers flooding data achieves only a 38% reduction for the same cost. Our framework can be applied more broadly to integrate information from a variety of sources about climate hazards and potential human impacts to make better informed decisions about investments in critical infrastructure systems.

CCS CONCEPTS

• **Social and professional topics → Sustainability**; • **Theory of computation → Network optimization**; • **Applied computing →** *Transportation*;

KEYWORDS

climate resilience, computational sustainability, mobility

*Both authors contributed equally to this work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
COMPASS '18, June 20–22, 2018, Menlo Park and San Jose, CA, USA
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5816-3/18/06…$15.00
https://doi.org/10.1145/3209811.3209859

ACM Reference Format:
Amrita Gupta, Caleb Robinson, and Bistra Dilkina. 2018. Infrastructure Resilience for Climate Adaptation. In *COMPASS '18: ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS), June 20–22, 2018, Menlo Park and San Jose, CA, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3209811.3209859

1 INTRODUCTION

Climate change is a global concern that is already causing widespread disruptions to environmental and socioeconomic systems and human health [6, 14, 16]. Recent studies have assessed the risks posed by extreme weather events, rising sea levels, and altered temperature and precipitation regimes to essential infrastructure systems [8, 18]. Damages to infrastructure are of particular concern for developing countries, where investment in energy, water, communication and transport infrastructure is a key strategy for meeting several UN sustainable development goals. For instance, the estimated costs of road repair and maintenance across Africa under current climate change projections exceed $150 bn, which can significantly divert funding from initiatives for expansion and development [8]. Thus, climate change will exacerbate existing socioeconomic vulnerabilities and threaten the success of crucial development schemes unless steps are taken to proactively mitigate these costs.

Road networks are especially important to supporting socioeconomic development in the least developed countries, since they provide access to services like education and healthcare and enable overland trade flows that are integral to the growth of developing economies. Africa's infrastructure deficit has led to the establishment of initiatives like the Programme for Infrastructure Development in Africa [19], and indeed, Goal 9 of the United Nation's Sustainable Development Goals is to "build resilient infrastructure, promote sustainable industrialization and foster innovation". However, currently about 85% of the road network in Western Africa and 88% of that in south central Africa are comprised of unpaved roads [8], making them highly susceptible to damage from precipitation which accounts for 80% of their degradation [7].

In general, data-driven machine learning and optimization tools have been used to help policy makers at all scales of government make better informed decisions with regards to *sustainability* objectives. For example, deep learning models have been used to combine satellite imagery, night light data, and scarce household survey data in Africa to create spatially explicit poverty maps in African countries using only satellite imagery [11]. This type of information helps to close the data gap between the developed and developing world and lets policy makers allocate resources to aid in areas where the need is the greatest. Mathematical optimization tools have also been widely adopted by conservation planners designing wildlife conservation reserves that balance multiple objectives, respect budget restrictions, and meet spatial or connectivity requirements and conservation targets to achieve improved outcomes for endangered species [5, 9]. These tools help decision-makers develop systematic and cost-effective plans whose benefits and costs are explicitly quantified, which is critical for evaluating and reporting the success of these initiatives.

The first step in a framework for making investment recommendations for climate-resilient road infrastructure is analyzing the exposure and sensitivity of roads to climatic pressures. Secondly, the contribution of each road to regional accessibility must be assessed. Roads that are both critical thoroughfares and frequently rendered inoperative under historical or projected climatic conditions should be prioritized for weather-proofing and other upgrades. Lastly, after risk and mobility impacts are determined, optimization techniques can be used to determine explicit plans for allocating road maintenance funds under multiple sustainable development objectives, such as maximizing rural connectivity or minimizing the expected number of people isolated due to weather-related infrastructure failures. This approach has the potential to minimize the long-term cost of establishing a reliable road network while helping to buffer vulnerable populations from extreme weather events. This framework is similar in spirit to disaster response applications, in which the primary goal is typically to facilitate emergency evacuation [13] or protect critical assets [23]. These approaches focus on meeting immediate demands placed on the road network following a given disaster scenario, whereas our work focuses on recommending strategic upgrades to maintain the large-scale, normal functionality of the road transportation network under different scenarios.

In this work we implement the framework described in the previous paragraph using techniques from network science and optimization with compiled data sources, then demonstrate its application to improving the flood-resilience of the national road network in Senegal. Specifically, we find which roads in Senegal should be fortified against flooding to prevent losses to mobility, given a fixed budget for infrastructure investments. We integrate flooding data from Fathom.Global, mobility data from Orange S.A., and GIS data describing the major road network in Senegal to describe the potential impacts of different flooding scenarios on the road network and its functionality. We examine several strategies for recommending road upgrades that take different data dimensions into account. Our experiments show how solutions generated by the strategy that considers all of the available data is able to best increase the accessibility under different flooding scenarios, highlighting the need for both flooding and mobility data to prepare for flooding outcomes. In particular, high-resolution flood mapping has recently become available for the entire world, and call detail records are also globally available thanks to high cellular subscription rates in most countries. This highlights the promise of our approach for developing decision support tools for policy makers in Senegal and worldwide.

2 METHODS

2.1 Data

The process that we follow in this paper can be seen as a general framework for studying the effects of flooding on mobility and providing recommendations for how to maximize accessibility in *any* country. This framework requires three pieces of data that cover the study area: 1.) a GIS representation of the road network, 2.) a spatially-explicit measure of flood risk, and 3.) mobility data describing flows between different zones/regions. While this work is an application of our framework to Senegal, the same process can be applied to other places.

We first construct an undirected graph representation of the road network based on an ArcGIS shapefile [1], with edges E representing road segments and vertices V representing the latitude-longitude coordinates of the endpoints of these road segments. Each edge has a distance property representing the length of the corresponding road segment. For the Senegal national and regional road network, this results in a connected graph with 6,917 vertices and 7,175 edges. Open source geospatial data made available by projects like OpenStreetMap [4] can be used to generate these graph representations for road networks in other regions or at different resolutions.

The Fathom Global dataset [20, 22] provides flooding data for the entire globe at the resolution of approximately 90m^2. Specifically, this dataset provides flood depth rasters for floods of different severities characterized by return period λ, which is the estimated time interval between flooding events of a similar intensity. Intuitively, a 500-year flood is more severe and less likely to occur than a 100-year flood. Given a return period λ, the value of each cell of the flood depth raster is the maximum flood depth estimated by a hydrodynamic model for a flood of the specified severity. While our study area is limited to Senegal, this dataset is generated at a global scale and thus could be utilized for similar studies in other countries. There are also alternative freely-available data sources for flood mapping, such as from the NASA MODIS near real-time global flood mapping project [3].

Call detail record (CDR) data, consisting of time-indexed sequences of cell towers used by anonymized users, is an excellent source of ground-truth human mobility data. In Senegal, we use CDR data provided by Orange [17], the biggest mobile provider in Senegal with 1,666 cell towers across the country, with the data made available through the UN Data for Climate Action Challenge [2]. Since CDR data is recorded only in terms of cell tower used, we use the approximate latitude-longitude locations of the N cell towers to construct "cell tower zones", Voronoi regions containing all locations that are closer to a given cell tower than to any other tower (Figure 1(c)). Given the sequence of cell tower zones visited by each user in the dataset, e.g. $\{l_1, l_2, \cdots, l_K\}$ for a customer that moves through K zones, we consider each consecutive pair of cell tower zones (l_k, l_{k+1}) to be a trip from zone l_k to zone l_{k+1}. We obtain the inter-zone travel demand by constructing an $N \times N$ origin-destination (OD) zone trip matrix $\mathbf{T}$ whose (i, j)-th entry represents the total number of trips taken from cell tower zone i to zone j by all users in the dataset. For further applications, and in the absence of CDR data, analytical human mobility models, such

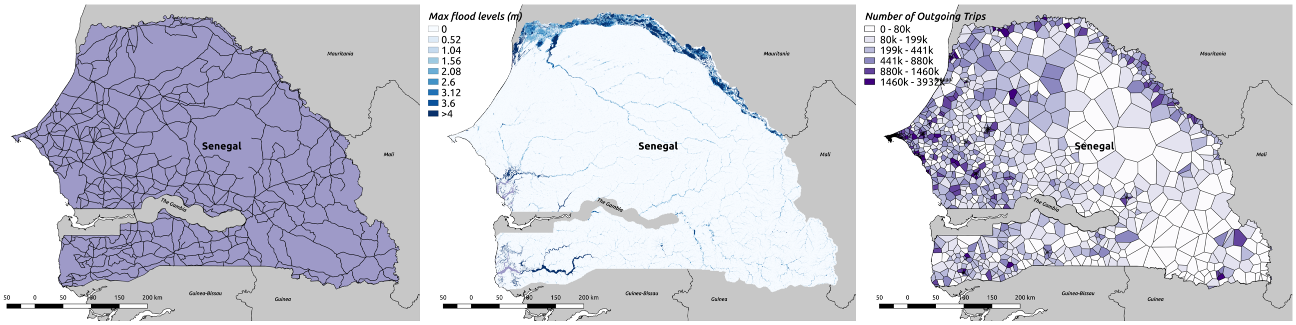

(a) Senegal national and regional highway system. **(b) Estimated flood depth for an 100-year flood.** **(c) Number of outgoing trips per cell tower zone.**

Figure 1: Data sources for estimating infrastructure exposure to flooding and potential impacts on mobility.

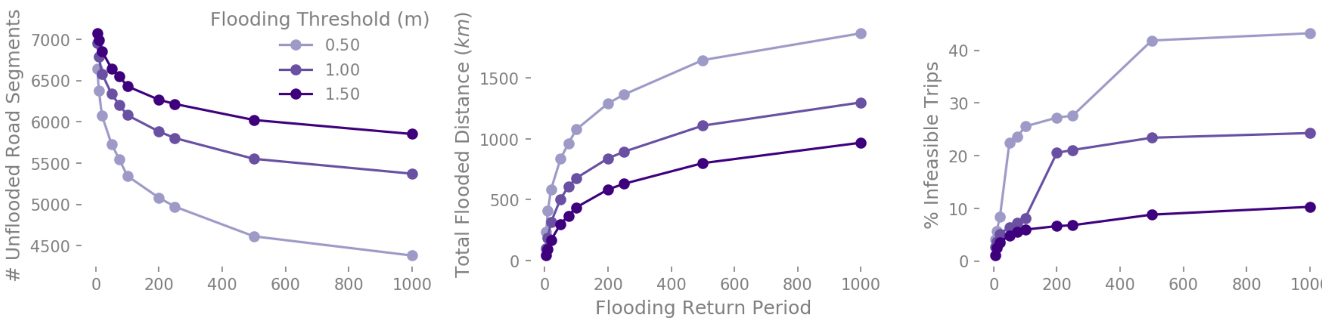

Figure 2: Graphs showing the effect of floods of increasing severity on the number of flooded road segments, total distance of flooding, and estimated percent of infeasible trips over the road network. Three flooding thresholds, 0.5, 1m, and 1.5m are shown.

as gravity models [12], can be used to approximate zone-to-zone trips given the population of each zone.

2.2 Vulnerability Assessment

The first step towards improving the flood resilience of the Senegal road network is to estimate the exposure of each road segment to flooding in a given return period scenario. To do this, we first apply a threshold of 1m to the flood depth raster and then obtain the geometric intersection of roads with raster cells with a maximum flood depth value $\geq$ 1m. Edges corresponding to road segments that pass through any such cells form a subset $E_F^\lambda \subseteq E$ of flooded roads. For each edge in this set we compute the flooded distance to estimate the length of the road segment that would need to be upgraded to make the road segment traversable again. We also obtain an *unflooded* subgraph $G_U^\lambda = (V, E_U^\lambda)$ consisting of the original set of vertices V and any edges corresponding to road segments not affected by flooding, representing the parts of the road network that are still traversable. Note, the two edge sets have no edges in common (i.e. $E_F^\lambda \cap E_U^\lambda = \emptyset$) and $E_F^\lambda \cup E_U^\lambda = E$.

Next, we estimate the accessibility impact of a given flood scenario. Various measures have been used to describe the performance of transportation networks in disasters, many of which are based on the idea of generalized travel cost [24], which quantifies the total amount of distance, time, money, etc. that must be traveled or spent in order to meet travel demand using a given transportation network. A degraded network will have a higher generalized travel cost due to the loss of paths leading to increased difficulty of traveling between pairs of locations. In this work, we consider binary costs, in which the cost associated with an attempted trip from one location to another is 0 if the trip is successful and 1 if there is no path between those locations in the degraded network. The decrease in the accessibility of the road network can then be quantified in terms of the number of trips that become infeasible due to flooding.

In order to connect the flooding effects on the road network to the zone-to-zone demand data, we assume that trips between pairs of zones happen between random origins and destinations within each zone. Specifically, for each zone i, we calculate Z_i, the set of vertices from the road network that are within its boundaries. Then, we assume that a trip leaving zone i and arriving at zone j is equally likely to start from any vertex within Z_i, and similarly equally likely to end at any vertex in zone Z_j. For a given road network $G^\lambda = (V, E^\lambda)$, we compute two $N \times N$ matrices $\mathbf{C}^0(G)$ and $\mathbf{C}^1(G)$. In $\mathbf{C}^1(G)$ entry (i, j) stores the number of pairs of vertices (u, v) with $u \in Z_i$ and $v \in Z_j$ between which a path exists in graph G. Similarly, in $\mathbf{C}^0(G)$ the (i, j)-th entry stores the number of pairs of vertices (u, v) between which there is no path in graph G. Then, based on our assumption the fraction of unsuccessful trips from zone i to j is given by $\mathbf{C}^0_{ij}(G)/\left(\mathbf{C}^0_{ij}(G) + \mathbf{C}^1_{ij}(G)\right)$. Then, we let

$$I(G) = \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{T_{ij}\mathbf{C}^0_{ij}(G)}{\mathbf{C}^0_{ij}(G) + \mathbf{C}^1_{ij}(G)}$$

be the number of infeasible trips. Here the number of infeasible trips between an origin i and destination j is counted as the total number of trips from i to j, multiplied the fraction of infeasible paths between the two zones.

2.3 Optimizing Accessibility in a Given Flooding Scenario

One plausible goal for improving the flood-resilience of the national road network is to strategically fortify road segments against flooding (e.g. with surfacing, elevation, etc.) such that the normal functioning of the network is preserved as much as possible. We propose minimizing the estimated number of trips made infeasible due to flooding as a way to achieve this goal. We suppose that we have a fixed road maintenance budget available for financing upgrades, and further that the monetary cost of fortifying a road segment against flooding of a given severity is proportional to the length of the road segment that is estimated to become flooded in that scenario. Formally:

Given: A graph G^λ_U representing the flooded parts of the road network, E^λ_F a set of flooded roads, costs $c(e)$ for upgrading each road $e \in E^\lambda_F$ and budget B.
Find: A feasible intervention plan consisting of edges $U \subseteq E^\lambda_F$ such that $\sum_{e \in U} c(e) \leq B$, that minimizes $\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{T_{ij}\mathbf{C}^0_{ij}(G')}{\mathbf{C}^0_{ij}(G')+\mathbf{C}^1_{ij}(G')}$, where $G' = (V, U \cup E^\lambda_U)$.

Note that the quantities $\mathbf{C}^0_{ij}(G')$ and $\mathbf{C}^1_{ij}(G')$ are dependent on our road upgrade decisions U. Given a candidate set U, computing $\mathbf{C}^0(G')$ and $\mathbf{C}^1(G')$ involves checking for the existence of paths between pairs of vertices in the upgraded road network, which can be done in a polynomial number of calls to BFS or DFS. However, the number of possible combinations of roads we can upgrade is exponential in $|E^\lambda_F|$, making an exhaustive search for an optimal feasible solution intractable. One solution strategy for solving combinatorial optimizations problems like the above problem is to formulate a mixed integer linear program in which decisions about whether or not to upgrade each road segment are encoded in binary decision variables and the existence of paths between pairs of nodes is encoded using network flow constraints (e.g. see [10]). Then, sophisticated commercial LP solvers such as CPLEX or Gurobi may be applied to solve the problem using a combination of heuristics and the branch-and-bound algorithm. Alternatively, it is possible to iteratively construct the set U of roads to upgrade following, e.g. a greedy heuristic [15], and metaheuristic optimization techniques such as GRASP (greedy randomized adaptive search procedure) or ant colony optimization could be employed to improve the performance of the heuristic algorithm in practice. When the objective function exhibits the property of *submodularity* or diminishing returns, the greedy heuristic has proven approximation guarantees [15]. Although our objective function does not exhibit submodularity, we nevertheless adopt a greedy strategy, selecting the edge that affords the largest improvement in the number of feasible trips per unit cost at each iteration. Specifically, at the k-th step we add the edge e with the largest value $\frac{I(G'_{k-1})-I(G'_k)}{c(e)}$, where $G'_0 = G^\lambda_U$ and $G'_k = G'_{k-1} \cup \{e\}$. Note that even this simple greedy algorithm is computationally expensive: we will start with $|E^\lambda_F|$ potential choices, where evaluating each choice will require recomputing all-pairs shortest paths. Once the "best" edge is selected and added to the graph, we must repeat this process with $|E^\lambda_F| - 1$ candidate choices, and so on until the given budget is used.

3 EXPERIMENTS

3.1 Baselines

We consider a set of baseline infrastructure development schemes to assess demand- and flooding-aware road upgrades under a given budget. These baselines, the **mobility** and **flooding** methods, are chosen to contrast against the potential benefits afforded by incorporating *both* spatially-explicit flooding and travel demand data to minimize the number of infeasible trips due to flooding, which we refer to in this section as **mobility + flooding**.

Mobility: Transportation infrastructure development plans will likely prioritize investments in assets that contribute to greater mobility and connectivity, e.g. by facilitating travel demand between major urban centers. As a simplistic example of such a strategy, we consider upgrading roads in order of their usage, estimated by assigning trips to routes and determining which road segments get the most trips. That is, we construct the set U of roads to upgrade by adding edges in decreasing order of $\frac{\text{(estimated number of trips over)e}}{\text{(distance of)e}}$ until we exhaust the budget. All roads in the road network are candidates for upgrades, and the amount of investment in each road segment is set proportional to the length of the road. Specifically, we set $\beta = \frac{1}{|E^\lambda_F|} \cdot \sum_{e \in E^\lambda_F} \frac{\text{(flooded distance in } e)}{\text{(distance of } e)}$. Then, for all edges $e \in E$ we set $c(e) = \beta \cdot \text{(distance of } e)$. Intuitively, we set the cost of upgrading a road segment to be on average the same amount it would cost to fortify in the flooded scenario, if it was flooded. When spatially-explicit flooding information is unavailable or not used to assess costs, it is possible to 1) select roads that are not flooded to upgrade in order to improve some other accessibility objective, and 2) upgrade a road that gets flooded, but to allocate a lower investment than the amount required to adequately fortify the road against the flooding projected to occur, in which case the road remains impassable under the flooding scenario.

Flooding: An alternative goal may be to improve the flood-resilience of the road network by ensuring that the maximum amount of the network remains operational in a flood of a given severity. This equates to maximizing the number of edges that can be fortified with a fixed budget. Here, only roads that will be flooded in the given flood setting are candidates for upgrading, and we set $c(e) = \text{(flooded distance in } e)$ for all such edges $e \in E^\lambda_F$. Since a road segment is considered

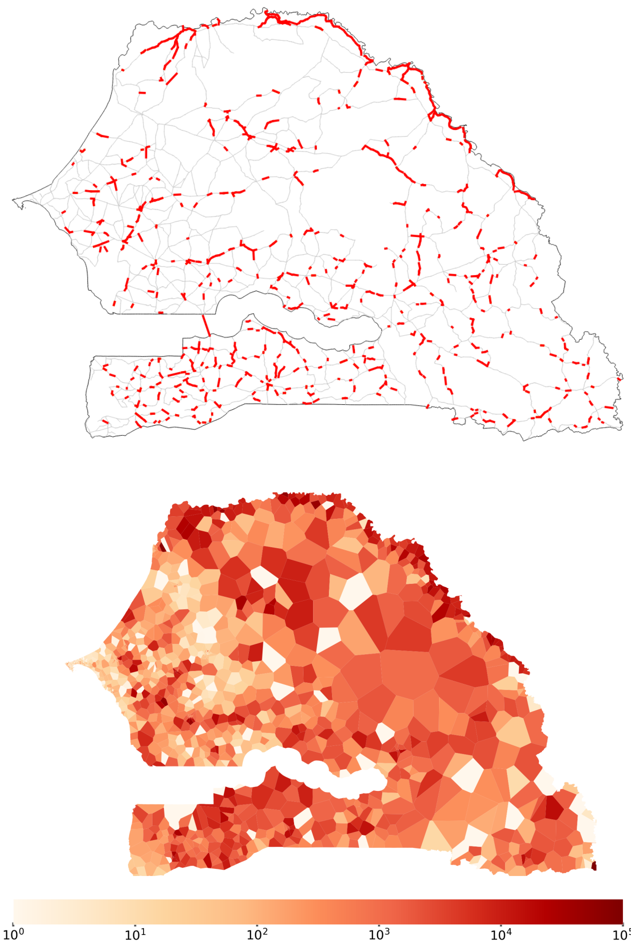

Figure 3: Impacts of the $\lambda = 100$ flooding scenario and a threshold of 1m. The (top) figure shows the road network of Senegal, with red links indicating roads that are impassable due to flooding. The (bottom) figure shows the estimated number of outgoing trips per zone that cannot be completed due to the damages in the road network.

impassable even if a small fraction of it becomes flooded, we construct the set of roads to upgrade, U, by adding edges in increasing order of $c(e)$ until the budget is exhausted.

3.2 Impacts of Flooding on Interregional Transport in Senegal

3.2.1 Impacts of Different Flood Severities. First, we examine how floods of different severities affect the existing Senegal road network and the estimated number of trips that can be completed over the road network in each scenario. The functionality of the road network depends on both its *exposure* to flooding (or the magnitude of the flood event) as well as its *sensitivity* (characterized here by the flooding threshold), two key components of vulnerability analysis. Figure 2 shows how as the severity of the flood increases, the number of traversable road segments falls, where the size of this effect varies with the flooding threshold considered. Most of the road segment failures occur even during floods of relatively low severities; e.g. approximately 25% of road segments fail during a 1000-year flood due to flooding over 1m, but a 100-year flood is enough to cause 15% of road segments to fail to the same level and is 10 times more likely to occur. The total flooded distance also increases rapidly for low to moderate flood severities. Assuming that costs of road network repair, maintenance or fortification will be proportional to the length of roads that need to be fixed, this indicates that attempting to fortify the entire network against even relatively minor, frequent flooding events will quickly become very costly. This highlights the need for a cost-effective infrastructure development plan that will prioritize roads to fortify in order to recover the maximum effective functionality of the network, e.g. in terms of the number of trips that remain feasible.

As the flood severity increases, the percentage of trips that cannot be completed also increases, but in a non-smooth fashion relative to the road network flooding. This suggests there are different types of trips that become infeasible at different levels of flooding, with some becoming infeasible even in milder flooding scenarios and others becoming infeasible only when flooding becomes much more widespread.

3.2.2 Efficacy of Alternative Road Upgrade Schemes. Next, we compare the performance of the baseline methods, **mobility** and **flooding**, with the greedy optimization method, **mobility + flooding**, for choosing road upgrades that maximize the number of trips that can be successfully completed using the road network given flooding of severity λ. We focus on the 100-year flood scenario and consider flood depths of 1m or greater (Figure 3). In this setting, we see that road segments throughout the Senegal highway system become impassable, especially along the northern border along the Senegal River and south of the Gambia along the Casamance River. Mobility in these regions is severely impacted as can be seen by the large numbers of unsuccessful outgoing trips from zones in these areas.

Fortifying the roads with the highest estimated number of trips, i.e. using the **mobility** baseline method to pick roads to fortify, does not significantly reduce the number of infeasible trips, primarily because most of the roads selected for investment are not flooded. The **flooding** baseline method upgrades the largest number of flooded roads possible with the available budget and allows more trips to remain feasible under the specified flooding setting simply by maximizing the number of links that remain operational in the network. This method does not take into consideration either the level of use or demand on each link or the possible existence of alternative routes, both of which are factors in how critical a road segment is for regional connectivity. Roads that are flooded but enable relatively few trips may be considered less critical to supporting mobility than those that are heavily used. Similarly, flooded roads may not necessarily prevent trips from being completed if there are other paths that can be taken from the origin to the destination. These effects can only be considered by using both flooding exposure and mobility impacts to inform investment decisions, as in the **mobility + flooding** method. Figure 4 shows that the **mobility + flooding** approach can yield a budget allocation

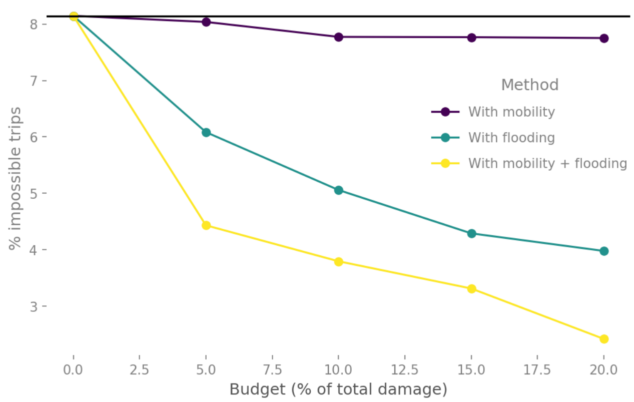

Figure 4: Effect of alternative road investment methods with varying budgets on the percent of impossible trips under the $\lambda = 100$ flooding scenario and 1m threshold. Budgets are shown as percentage of the cost of repairing *all* roads in the given flooding scenario. The black horizontal line shows the baseline percentage of impossible trips when no road investments are made.

that recovers more effective functionality of the road network (in terms of the number of successful trips made) even though the number of links that are restored is lower than the other methods (e.g. in the 5% budget case, the **mobility + flooding** method restores 175 edges, while the **mobility** and **flooding** methods restore 179 and 359 edges respectively).

3.2.3 Differential Regional Benefits. We compare the spatial distribution of roads recommended for repair by the **mobility + flooding** method to those chosen by the baseline methods. Figure 5 shows the roads that are selected for upgrades by each method, as well as the improvements to the number of possible outgoing trips per zone in the $\lambda = 100$ flood scenario. While the **flooding** method chooses almost twice as many individual road segments as the **mobility + flooding** method, the additional edges do not confer as much benefit to mobility. For instance, at the 5% budget level, there are 49 edges chosen by the **mobility + flooding** method that are *not* selected by the *flooding* method, which then must be responsible for the 2% difference in the improvements between both methods. These 49 roads are focused around higher population areas in the north and regions around the capital Dakar on the western tip of Senegal. Specifically, in Figure 5(b) we observe that there are larger improvements to the number of outgoing trips in the high population areas in Saint-Louis on the north western coastline, and the Diourbel and Kaolack urban regions to the east of Dakar while there are few improvements in the south-eastern sparsely populated portion of the country. In contrast, Figure 5(d) shows that the **flooding** method results in more improvements in the south east portion, with fewer improvements in the same higher populated areas. It is also evident from the spatial distribution of the repaired roads in each case that incorporating travel demand results in chosen roads that are closer to major settlements. We can see in Figure 5(e) that the **mobility** method suggests road fortifications along the main N1 and N7 highways as those roads carry the largest number of trips. This solution, which may be beneficial under different objectives, does not help restore connectivity, as these roads are not compromised due to flooding (by either not being flooded or having alternate unflooded paths available for use).

4 DISCUSSION

We have proposed a framework to quantify the impacts of flooding on human mobility and subsequently provide road fortification suggestions that maximize accessibility over the road network. We demonstrate the application of this framework in Senegal. Our results suggest that fortification decisions that are made only considering a single data dimension (i.e. *flooding* or *mobility*) fail to capture the information needed to improve multi-dimensional objectives (i.e. number of possible trips, which is a function of *flooding and mobility*). In other words, the solutions that are found through methods that only use a single data source do not have large intersections with those that use multiple data sources. We cannot expect to improve the climate resilience of the existing road network in a cost-effective manner without specifically including information on the spatial distribution of floods as well as the movements of the population.

Our framework can be used to develop a decision support tool to help policy makers and urban planners systematically identify climate change-related risks to infrastructure. Spatially explicit data or models characterizing the distribution, magnitude or probability of relevant climate stresses provide information that can be used both to estimate damages and set budgets, as well as to create proactive management actions. Then, a decision support system could evaluate a given plan in terms of multiple quantitative objectives. Alternative recommendations generated using mathematical optimization can then be compared and used to subsequently prioritize investments. In particular, planners can easily examine the effects of preparing for a given flood magnitude or improving infrastructure up to a certain sensitivity level by varying these inputs to the decision support system.

One key component of our work is the emphasis on the interaction between decision making and data. Indeed, without the access to *and incorporation of* relevant and disaggregated data, decision making tools will be limited in their capacity to provide relevant recommendations. Our results highlight this point in the context of recommending infrastructure upgrades to ensure mobility under different flooding scenarios in Senegal, but this applies to many sustainability related decision making applications.

Considering further data sources, we can extend and adapt our work in a number of ways to further build climate change resilience into development planning. One obvious refinement of our current approach is to replace the "flood risk data" with a predictive module that can take weather data as input and predict spatially explicit flood extents and depths, which can then be used in combination with climate change models to evaluate flood-related mobility impacts under different climate change scenarios. Vulnerability to other climate change-related hazards, such as droughts or wildfires, can also be incorporated to develop a more complete picture of climate change exposure and risks faced in different regions. Other components of our framework that can be improved upon include

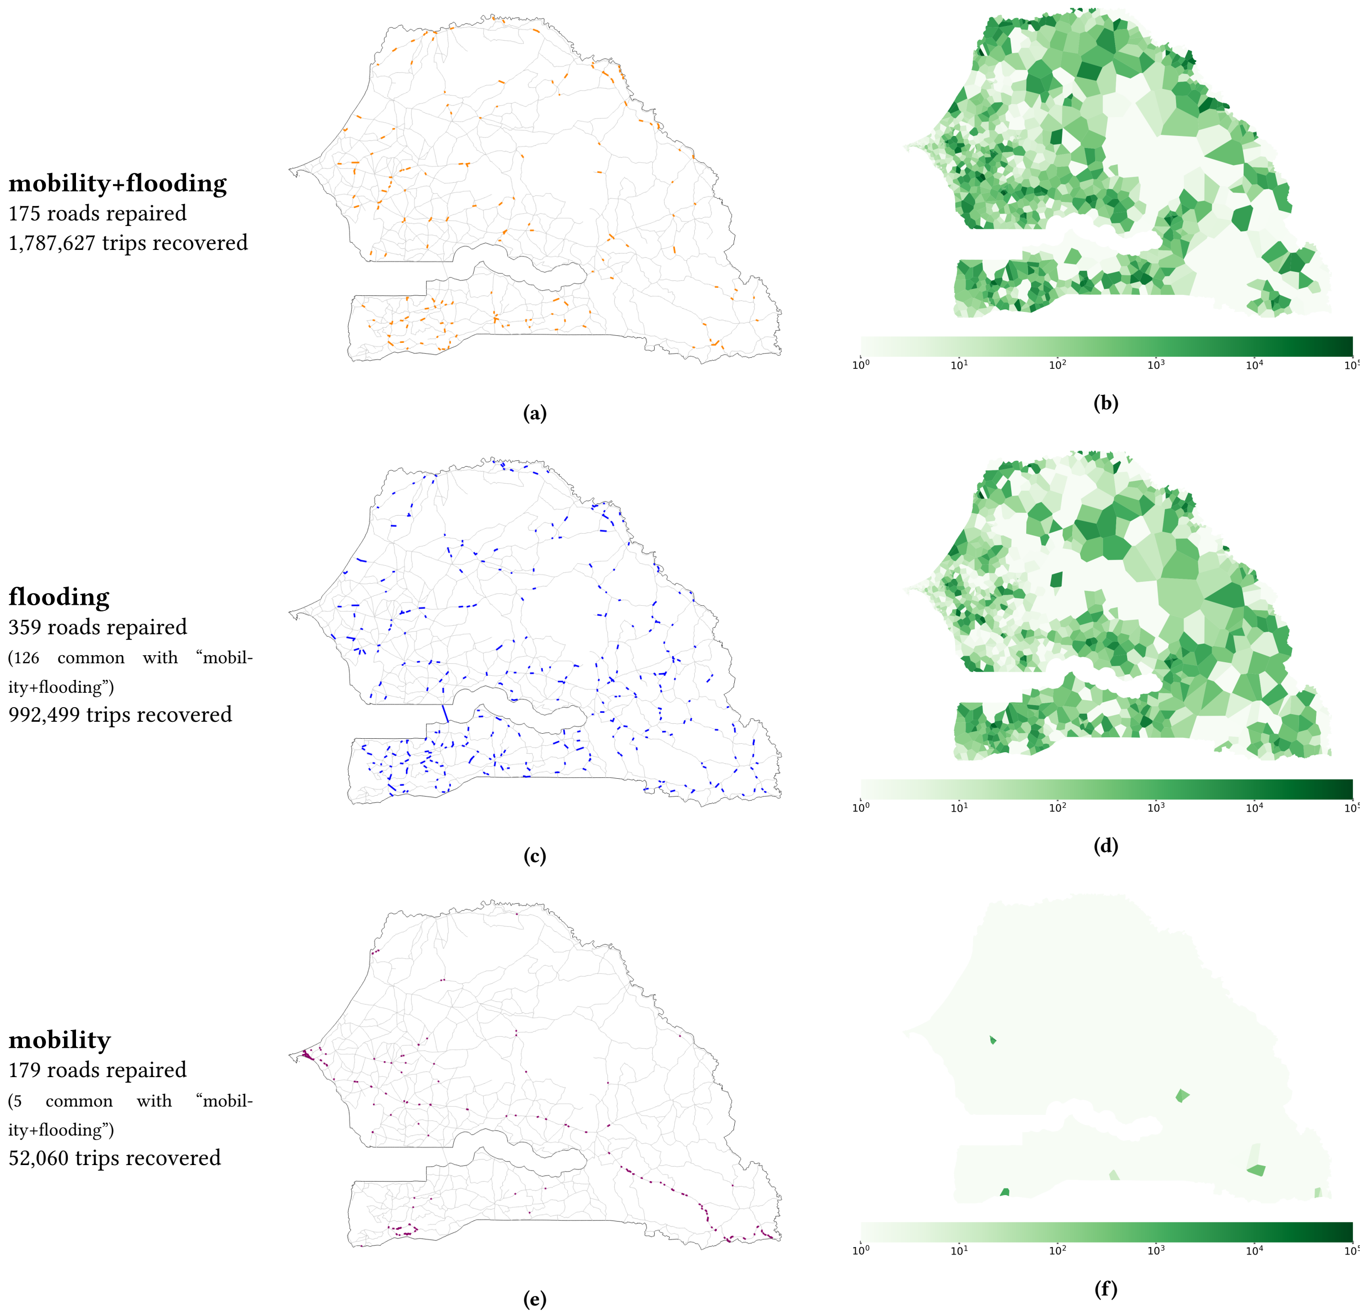

Figure 5: The roads chosen for improvement (left column) and corresponding increases in numbers of outgoing feasible trips per zone (right column) for each of the three methods discussed in the text under a 5% budget. Each map in the right column shows the difference in the number of infeasible outgoing trips per zone between the baseline scenario and the fortified scenario.

our definition of accessibility, which can be modified to weigh rural regions more heavily or prioritize routes needed for the distribution of disaster relief resources (i.e. utilizing spatially explicit population data). If ground truth travel demand data (which we derive from CDR data) is unavailable, gravity or radiation models of human mobility [12, 21] can be used as a substitute. In general, modeled inputs can be useful to test "what if" scenarios. However, access to ground truth data is essential to generate truly context-specific recommendations in sustainability applications. Finally, the optimization method we proposed can be improved upon to find better solutions, whether through a mathematical formulation, or through algorithmic improvements. Similarly the optimization objective could be changed to balance multiple policy goals, or to find solutions that are robust to uncertainties in the input data sources.

ACKNOWLEDGMENTS

This work is supported by the NSF grant CCF-1522054 (COMPUST-NET: Expanding Horizons of Computational Sustainability). We also thank Orange and Sonatel for access to the mobility data, and Fathom.Global for access to the flooding data.

REFERENCES

[1] 2013. Senegal Roads - ArcGIS. https://www.arcgis.com/home/item.html?id=7a4486a4903b42c3ba3ce3a750c3bcca Accessed: 2017-06-22.

[2] 2017. Data For Climate Action Challenge Data. http://dataforclimateaction.org/home/data/ Accessed: 2018-03-01.

[3] 2018. NASA MODIS. https://floodmap.modaps.eosdis.nasa.gov/ Accessed: 2018-03-01.

[4] 2018. OpenStreetMap. https://www.openstreetmap.org/about Accessed: 2018-03-01.

[5] Ian R Ball, Hugh P Possingham, and M Watts. 2009. Marxan and relatives: software for spatial conservation prioritisation. *Spatial conservation prioritisation: Quantitative methods and computational tools* (2009), 185–195.

[6] David S Battisti and Rosamond L Naylor. 2009. Historical warnings of future food insecurity with unprecedented seasonal heat. *Science* 323, 5911 (2009), 240–244.

[7] Paul Chinowsky and Channing Arndt. 2012. Climate change and roads: A dynamic stressor–response model. *Review of Development Economics* 16, 3 (2012), 448–462.

[8] Paul Chinowsky, Amy Schweikert, Niko Strzepek, Kyle Manahan, Kenneth Strzepek, and C Adam Schlosser. 2013. Climate change adaptation advantage for African road infrastructure. *Climatic change* 117, 1-2 (2013), 345–361.

[9] Bistra Dilkina and Carla P Gomes. 2010. Solving connected subgraph problems in wildlife conservation. In *International Conference on Integration of Artificial Intelligence (AI) and Operations Research (OR) Techniques in Constraint Programming*. Springer, 102–116.

[10] Bistra Dilkina, Rachel Houtman, Carla P Gomes, Claire A Montgomery, Kevin S McKelvey, Katherine Kendall, Tabitha A Graves, Richard Bernstein, and Michael K Schwartz. 2017. Trade-offs and efficiencies in optimal budget-constrained multispecies corridor networks. *Conservation Biology* 31, 1 (2017), 192–202.

[11] Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. 2016. Combining satellite imagery and machine learning to predict poverty. *Science* 353, 6301 (2016), 790–794.

[12] Chaogui Kang, Yu Liu, Diansheng Guo, and Kun Qin. 2015. A generalized radiation model for human mobility: spatial scale, searching direction and trip constraint. *PloS one* 10, 11 (2015), e0143500.

[13] Kanal Kumar, Julia Romanski, and Pascal Van Hentenryck. 2016. Optimizing Infrastructure Enhancements for Evacuation Planning.. In *AAAI*. 3864–3870.

[14] Corey Lesk, Pedram Rowhani, and Navin Ramankutty. 2016. Influence of extreme weather disasters on global crop production. *Nature* 529, 7584 (2016), 84.

[15] Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne VanBriesen, and Natalie Glance. 2007. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 420–429.

[16] Anthony J McMichael. 2015. Extreme weather events and infectious disease outbreaks. *Virulence* 6, 6 (2015), 543–547.

[17] Orange Data for Development Challenge. 2013. Book of abstract : scientific paper. http://www.d4d.orange.com/fr/content/download/43453/406503/version/1/file/D4DChallengeSenegal_Book_of_Abstracts_Scientific_Papers.pdf. Accessed: 2017-06-24.

[18] Mathaios Panteli and Pierluigi Mancarella. 2015. Modeling and evaluating the resilience of critical electrical power infrastructure to extreme weather events. *IEEE Systems Journal* (2015).

[19] Jean Ping, Donald Kaberuka, and Abdoulie Janneh. 2011. Programme for infrastructure development in Africa. *Interconnecting, Intergrating & transforming a continent* (2011).

[20] Christopher C Sampson, Andrew M Smith, Paul D Bates, Jeffrey C Neal, Lorenzo Alfieri, and Jim E Freer. 2015. A high-resolution global flood hazard model. *Water resources research* 51, 9 (2015), 7358–7381.

[21] Filippo Simini, Marta C González, Amos Maritan, and Albert-László Barabási. 2012. A universal model for mobility and migration patterns. *Nature* 484, 7392 (2012), 96.

[22] SSBN. 2017. SSBN-Global Senegal. http://www.ssbn.co.uk/ssbnglobal. Accessed: 2017-06-25.

[23] Byron Tasseff, Russell Bent, and Pascal Van Hentenryck. 2016. Optimal Flood Mitigation over Flood Propagation Approximations. In *International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. Springer, 358–373.

[24] Michael AP Taylor, Somenahalli VC Sekhar, and Glen M D'Este. 2006. Application of accessibility based methods for vulnerability analysis of strategic road networks. *Networks and Spatial Economics* 6, 3-4 (2006), 267–291.