

Data Extrapolation in Social Sensing for Disaster Response

Siyu Gu*, Chenji Pan*, Hengchang Liu[†], Shen Li*, Shaohan Hu*, Lu Su*, Shiguang Wang*, Dong Wang*,
Tanvir Amin*, Ramesh Govindan[¶], Charu Aggarwal^{||}, Raghu Ganti^{||}, Mudhakar Srivatsa^{||}, Amotz Barnoy**,
Peter Terlecky**, Tarek Abdelzaher*

*University of Illinois at Urbana-Champaign, [†]University of Science and Technology of China,

[¶]University of Southern California, ^{||}IBM T. J. Watson Research Center, **City University of New York

Abstract—This paper complements the large body of social sensing literature by developing means for augmenting sensing data with inference results that “fill-in” missing pieces. Unlike trend-extrapolation methods, we focus on prediction in disaster scenarios where disruptive trend changes occur. A set of prediction heuristics (and a standard trend extrapolation algorithm) are compared that use either predominantly-spatial or predominantly-temporal correlations for data extrapolation purposes. The evaluation shows that none of them do well consistently. This is because monitored system state, in the aftermath of disasters, alternates between periods of relative calm and periods of disruptive change (e.g., aftershocks). A good prediction algorithm, therefore, needs to intelligently combine time-based data extrapolation during periods of calm, and spatial data extrapolation during periods of change. The paper develops such an algorithm. The algorithm is tested using data collected during the New York City crisis in the aftermath of Hurricane Sandy in November 2012. Results show that consistently good predictions are achieved. The work is unique in addressing the bi-modal nature of damage propagation in complex systems subjected to stress, and offers a simple solution to the problem.

I. INTRODUCTION

This paper addresses the problem of data extrapolation in participatory sensing applications, in the face of disruptive pattern changes, such as those that occur during natural disasters. We consider cases where resource limitations or accessibility constraints prevent attainment of full real-time coverage of the measured data space, hence calling for data extrapolation.

Many participatory sensing applications were investigated in recent years [1]–[6]. In participatory sensing, sources measure application-related state at locations of interest then usually report it at a later time (e.g., when they encounter a WiFi access point a few hours later). Hence, at any given time, the latest state of some points of interest may be unknown. Incomplete real-time coverage may also arise due to scarcity of sensing resources. For example, volunteers in a disaster-response application may survey and report locations of damage. If there are fewer volunteers than damage locations, the state of some of these locations will not be immediately reported. In such scenarios, one question is: can we infer the missing data?

Many time-series data extrapolation approaches are based on the assumption that past trends are predictive of future values. These approaches do not do well when disruptive changes occur. For example, a history of no traffic congestion on main highways of some city does not offer a good traffic

predictor if a natural disaster causes a mass evacuation. An alternative recourse is to consider only spatial correlations. For example, certain city streets tend to get flooded together after heavy rain (e.g., because they are at the same low elevation), and certain blocks tend to run out of power together after a thunderstorm (e.g., because they share the same power lines). Understanding such correlations can thus help infer state at some locations from state at others when disruptive changes (such as a flood or a power outage) occur.

In this paper, we show that system state in post-disaster scenarios alternates between periods of calm (when the past is a good predictor of the future) and periods of sudden change, as new parts of the infrastructure are damaged (e.g., due to aftershocks) or repaired. Hence, data extrapolation algorithms that rely predominantly on spatial correlations or predominantly on temporal correlations tend not to work consistently well, as the relative importance weights of temporal versus spatial correlations change significantly between periods of calm and periods of change. Instead, we show that such algorithms must switch intelligently between two extrapolation modes with different emphasis on temporal versus spatial correlations.

Of special interest is the case where correlations needed for extrapolation are themselves not known in advance, but are rather learned on the fly. The need for joint learning and extrapolation distinguishes this paper from some existing work [7]–[9] that predicts missing sensor values assuming a *previously known* correlation structure between sensors, or a known temporal pattern.

We apply the results to an example case study of a New York City crisis in the aftermath of Hurricane Sandy. Many gas stations, pharmacies, and grocery stores around New York City were closed after the hurricane, resulting in severe supply shortage that lasted several days. The outages were correlated, since different stores shared suppliers or power. Our study shows the degree to which extrapolation could infer gas, food, and medical supply availability during the crisis in the absence of complete and fresh information.

The remainder of this paper is organized as follows. We present the general system design and illustrate prediction challenges in Section II. A new algorithm that addresses these challenges via appropriate switching between spatial and temporal extrapolation is presented in Section III. An evaluation is presented in Section IV. Section V reviews related work. We conclude the paper in Section VI.

II. SYSTEM MODEL

We consider a model of participatory sensing applications in which the reported state is binary. It is desired to obtain the state of several points of interest (PoIs). A central collection node (e.g., the command center) collects the state from participants who make observations and report them later.

The time when participants report their observations may vary. Measurements that are older than some threshold, are deemed stale. Hence, at any given time, there may be “blind points” in the PoI map generated by participants, where fresh information is not available. The challenge is to infer the missing state automatically and accurately.

The main contribution of this work lies in addressing the extrapolation problem in scenarios consistent with disaster response. Two main challenges characterize those scenarios:

- *Disruptive change*: By definition, disasters are unique disruptive events that invalidate normal data trends, making prediction based on historical (time-series) trends largely incorrect.
- *Scarcity of training data*: Since disasters are rare and generally unique, there is very little training data that one can rely on. To understand the worst case, we restrict the prediction algorithm to use only training data available from the current disaster itself. This scarcity of data severely limits the complexity of prediction models that can be used.

We consider applications where today’s information matters the most and people prefer undertaking some actions based on best-effort guessing to obtaining exact data at a certain delay. For example, in the case of finding gas stations around New York City that are operational after hurricane Sandy, if one needed to fill up their car now, yesterday’s gas availability would be of less use. The challenge is therefore to infer the *current* missing PoI state.

We assume that old (and hence potentially stale) information on PoI state is available. For example, in disaster response scenarios, volunteers might physically report back to the command center daily, which makes yesterday’s information available at the center. We call the maximum reporting latency, a *cycle*. Hence, by definition, the backend server knows the state of all PoI sites in previous cycles, but has only partial information in the current cycle. This assumption simplifies our algorithmic treatment. It can easily be relaxed allowing for information gaps in previous cycles as well, since such gaps can always be filled in using the same extrapolation algorithm, applied to past state.

A. Problem Statement and Solution Challenges

More formally, our participatory sensing system can be characterized by a weighted graph $G = (V, E)$, $|V| = n$, $|E| = m$, where the node set V represents the n PoIs. We assume that set V is known and remains unchanged. The link set E represents the correlations among PoIs.

One way to compute links E , is to apply the Kendall’s Tau statistical method [10] to estimate correlations. More concretely, assume two PoIs, x and y , have data (x_1, x_2, \dots, x_n) and (y_1, y_2, \dots, y_n) . The Kendall’s Tau correlation coefficient, denoted by $KT(x, y)$, can be represented as:

$$KT(x, y) = 1 - \frac{1}{n} \sum_{i=1}^n XOR(x_i, y_i) \quad (1)$$

Each edge (x, y) between PoI nodes x and y has a weight, $w_{xy} = KT(x, y)$, representing the correlation value. The link set E may be reduced by setting a predefined threshold such that only links with correlations higher than the threshold are retained.

The extrapolation algorithm takes partial state of PoI sites in the current cycle, historical data of PoI sites in previous cycles, and the relationships (i.e., edges) learned so far as inputs. It then infers the current state of missing PoI sites.

As argued above, scarcity of training data renders complex prediction models, such as ARIMA and various data mining models [11], ineffective. For example, on the 4th day of a disaster, we have only 3 past training points, which might be fewer than the number of parameters in some models. This means that our prediction model would have to be very simple. Indeed a contribution of this work lies in arriving at a very simple model that works well with little data, as opposed to beating the current mature state of the art in time-series prediction from large data sets.

We first consider several obvious simple heuristics that can be used for extrapolation. To illustrate the impact of insufficient training data, we also consider ARIMA [11], a standard (and powerful) time series analysis method for non-stationary processes, commonly used in complex forecasting tasks, such as forecasting financial systems [12]. The performance of these solutions will determine whether or not a new extrapolation approach is needed.

- *Random*: It is the most trivial baseline in which the status of missing sites is guessed at random. It shows what happens when no intelligence is used in guessing.
- *BestProxy*: It uses the Kendall’s Tau method to find actual pairwise (spatial) correlations between PoIs and predicts missing state based on the state of the best neighbor (i.e., the PoI that has the largest correlation with the one being predicted). It is an example of exploiting local spatial correlations, where state of an individual node is predicted from state of another (well-chosen) *individual node*.
- *Majority*: It computes the majority state of all known PoIs and predicts all missing state to be the same as the majority state. This heuristic is another example of exploiting spatial correlations. It lies at the other end of the spectrum from *BestProxy*, in that it exploits a global notion of spatial correlations, where state of an individual node is predicted from *global state*.
- *LastKnownState*: It explores temporal correlations among PoI sites. Namely, the predicted state today is set equal to the last known state.
- *ARIMA*: This, in principle, is one of the most general forecasting methods for time series data that assumes an underlying non-stationary process [11].

The performance of the above baselines is discussed next.

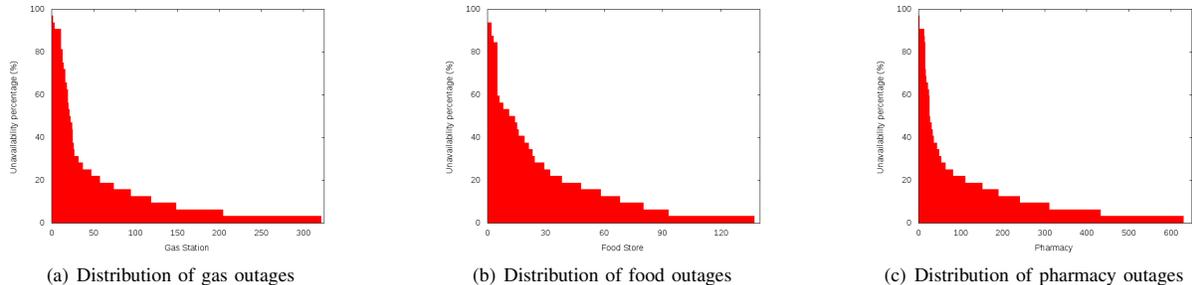


Fig. 1: Distribution of public services outages

B. Failure of Individual Baselines

We evaluate our baselines through a real-world disaster response application. In November 2012 [13], Hurricane Sandy made landfall in New York City. It was the second-costliest hurricane in United States history (surpassed only by hurricane Katrina) and the deadliest in 2012. The hurricane caused widespread shortage of gas, food, and medical supplies as gas stations, pharmacies and (grocery) retail shops were forced to close. The shortage lasted about a month. Recovery efforts were interrupted by subsequent events, hence triggering alternating relapse and recovery patterns.

The daily availability of gas, food, and medical supplies was documented by the All Hazard Consortium (AHC) [14], which is a state-sanctioned non-profit organization focused on homeland security, emergency management, and business continuity issues in the mid-Atlantic and northeast regions of the United States. Data traces¹ were collected in order to help identify locations of fuel, food, hotels and pharmacies that may be open in specific geographic areas to support government and/or private sector planning and response activities. The data covered states including West Virginia, Virginia, Pennsylvania, New York, New Jersey, Maryland, and District of Columbia. The information was updated daily (i.e., one observation per day for each gas station, pharmacy, or grocery shop). To give an example of the extent of damage, Figure 1(a) shows the distribution of the percentage of time that each of 300+ affected gas stations in the New York area was *unavailable* during the first *month* following the hurricane. We can see that 40 gas stations were not available for more than 1 week and some were out for the whole month. Similarly, Figure 1(b) shows the distribution of outage for affected food stores and Figure 1(c) shows the distribution of outage for affected pharmacies.

With these PoI sites and input data as ground truth, we evaluate the baselines described. The metrics we use are accuracy of inference and amount of data needed. We break time into cycles as discussed earlier. We set each cycle to a day to coincide with the AHC trace. We then plot the performance of the above baselines when a configurable amount of today’s data is available (in addition to all historic data since the beginning of the hurricane).

We evaluate the solutions on November 3rd, and November 8th. November 8th corresponds to a period of disruptive change due to a second snow storm that hit after Sandy, causing

massive temporary relapse of recovery efforts due to new power outages, followed by a quick state restoration to the previous recovery profile. November 3rd is an example of a period of little change, when damage was incurred but recovery efforts have not yet been effective. The same trend was observed for all datasets we have, namely, gas, pharmacy, and food.

Figure 2, Figure 3, and 4 plot the prediction error with standard deviation shown as error bars in availability of gas stations, food (grocery shops), and pharmacies, respectively. In each figure, sub-figures (a) and (b) refer to November 3rd and November 8th, respectively.

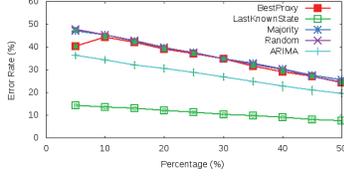
The reader is reminded that we assume that, on a given day, one knows the status of only a fraction of PoIs (where the status refers to whether they are open or closed). The purpose is to extrapolate this data and find out the status of the remaining ones. The horizontal axis in the aforementioned figures varies the percentage of PoIs whose status is known on the indicated day from 5% to 50%. To eliminate bias that may result from knowing the status of specific PoIs, each point (corresponding to a specific percentage of PoIs whose status is known) is an average of 50 different experiments. In each experiment, a different random set of PoIs is selected as known (adding up to the required percentage). The results shown are the average of the 50 experiments.

Consider Figure 2-a and Figure 2-b, that illustrate the overall prediction error rate for *gas availability* on November 3rd and 8th, respectively, as a function of the percentage of PoIs whose status is known that day. On the vertical axis, the performance of baselines is compared.

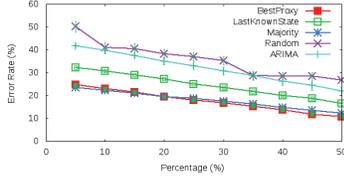
Figure 3 and Figure 4 compare the performance of baselines in predicting food and pharmacy availability.

It can be seen that no single baseline does consistently well in all figures. Specifically, LastKnownState does remarkably well on November 3rd, when the change was minimal from the day before. This is especially true for gas and food (grocery) availability prediction, where it beats the next heuristic by a wide margin. However, BestProxy does better on November 8th, when a second snow storm hits and its aftermath causes a lot of perturbation. More specifically, the error rate of BestProxy is around 8% lower than LastKnownState on November 8th. BestProxy clearly outperforms LastKnownState that day for gas and pharmacy availability prediction, and ties for food availability prediction. Majority does poorly on November 3rd and better (but not best) on November 8th. Random does

¹Available at: <http://www.ahcusa.org/hurricane-Sandy-assistance.htm>



(a) Error rate on November 3rd.



(b) Error rate on November 8th.

Fig. 2: Comparing baselines to predict gas availability after Sandy

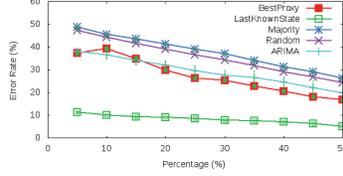
worse. Very interestingly, ARIMA does only marginally better than Random and much worse than the best heuristics on either day. This is attributed to the lack of sufficient training data, and the challenges caused by disruptive changes in the time-series. Also notice that, the standard deviations for all baseline methods are quite small compared to the error rates, which indicates that which PoIs are known does not have a significant effect on the performance.

The results confirm that algorithms that do spatial extrapolation (such as BestProxy) are better on days of more change, whereas algorithms that do temporal extrapolation (such as LastKnownState) are better on days of less change. The results also suggest that, due to lack of training data, complex prediction models that normally do well, such as ARIMA, are ineffective. We leverage these observations to guide the design of an algorithm that consistently offers the best performance. This algorithm appropriately adapts to periods of change versus periods of calm, and requires little training data. Note that, we do not aim to outperform any one heuristic at *all* times. Rather, our aim is to match consistently the best performing heuristic at any time, even though that heuristic changes, depending on circumstances. Such an algorithm is described next.

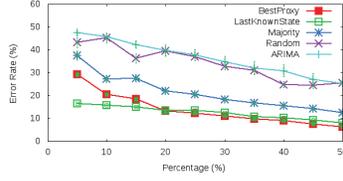
III. A HYBRID PREDICTION ALGORITHM

The above study leads to two insights that help develop an algorithm for data extrapolation in disaster response scenarios:

- *Insight #1:* The first insight is that our algorithm should be able to switch between spatial and temporal prediction modes. On days with little change, LastKnownState does really well and should be the default prediction. On days where change is abundant, spatial correlations are more appropriate to use for prediction.
- *Insight #2:* The second insight lies in refining the notion of spatial correlations to be used for prediction. Since our default prediction is LastKnownState (i.e., no change), we need spatial correlations only to predict *change*. Hence, rather than using Kendall’s Tau correlation to find a good proxy, we seek a proxy that helps predict change only. In

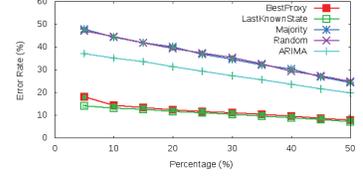


(a) Error rate on November 3rd

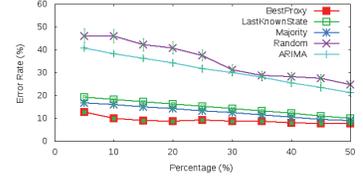


(b) Error rate on November 8th

Fig. 3: Comparing baselines to predict food availability after Sandy



(a) Error rate on November 3rd



(b) Error rate on November 8th

Fig. 4: Comparing baselines to predict pharmacy availability after Sandy

other words, we seek a proxy whose *state changes* (and not overall state) are most correlated with those of the target to be predicted.

The second insight is intuitive in retrospect. Just because two gas stations were out of gas or out of power for a long time, does not mean their state changes are correlated. What’s more indicative is whether or not they lost gas or power at the same time. The latter gives a better indication that if gas or power is restored to one, it may also be restored to the other.

More concretely, consider two PoIs, x and y , that have state (x_1, x_2, \dots, x_n) and (y_1, y_2, \dots, y_n) . Let x_n be unknown (i.e., it has not yet been delivered). Let us define the change time series as $(dx_1, dx_2, \dots, dx_n)$ and $(dy_1, dy_2, \dots, dy_n)$, where $dx_i = x_i - x_{i-1}$ and $dy_i = y_i - y_{i-1}$ (we assume that $x_0 = 1$ and $y_0 = 1$ (everything was working before the disaster)). To predict x_n (or equivalently predict the change dx_n), we would like to find a proxy y , whose current status is known and whose changes are maximally correlated with changes in x . We can then use dy_n to predict dx_n and hence predict x_n . To do so, we compute $P(\text{change in } x | \text{same change in } y)$ for all gas stations y whose current state is known. This probability can be approximated by:

$$P(\text{change in } x | \text{same change in } y) = \frac{\text{count}(dx_i = dy_i)}{\text{count}(dy_i \neq 0)} \quad (2)$$

where $\text{count}()$ is a function that counts the number of times the condition in its argument was true for $1 \leq i \leq n-1$. The best proxy for (predicting change in) x becomes the y that maximizes the above probability. Let us call such a y , y_{best} . Let the resulting probability, $P(\text{change in } x | \text{same change in } y_{best})$ be denoted P^{best} . Using insight #1 above, the sought algorithm is as follows:

Lines 1 to 4 indicate that the algorithm alternates between spatial and temporal prediction depending on whether the best found proxy for the target x is sufficiently good (i.e., better than a threshold, T). When spatial prediction is used, we predict that state of x will change (i) if it was the same as the state of the best proxy, and (ii) if the state of that proxy

Algorithm 1 ENHANCED BEST PROXY (x, n)

```
1: IF (  $P^{best} \geq \text{threshold } T$  )
2: use SpatialPrediction
3: ELSE
4: use LastKnownState (i.e.,  $x_n = x_{n-1}$ )
5:
6: SpatialPrediction
7: IF ( $(dy_n^{best}$  is not zero) AND ( $y_{n-1}^{best} = x_{n-1}$ ))
8: THEN  $x_n = y_n^{best}$ 
9: ELSE use LastKnownState (i.e.,  $x_n = x_{n-1}$ )
```

changed. Otherwise, we predict no change. Note that, it is possible that there is no best proxy for a certain PoI. When choosing the best proxy, we require one PoI to have at least a certain number of changes in its own history so far. To see why this is necessary, imagine we are now considering choosing PoI A as B's proxy, however, A has only 1 state change in its history and the change happened in the same cycle as B. In this case, A's P score will be 1, which is always larger than or equal to T and all other proxy candidates. Therefore, A will be selected as B's best proxy, although A is actually not a strong candidate.

It remains to derive the optimal value of the threshold, T . Let M denote the fraction of PoIs that had state = 1 in the last cycle. Hence, $1 - M$ is the fraction of PoIs with state = 0. Furthermore, let F denote the fraction of PoIs (that we are aware of so far) that change state in the current cycle. The optimal value of T is one that minimizes misprediction probability.

The above algorithm mispredicts either (i) when spatial prediction is used and it is wrong, or (ii) when temporal (LastKnownState) prediction is used and it is wrong. Hence, misprediction probability, P_m , is equal to the sum of spatial misprediction probability, P_{sm} , and temporal misprediction probability, P_{tm} . Below, we compute these probabilities.

Spatial Misprediction: From line 7 of Algorithm 1, spatial misprediction occurs when (i) P^{best} exceeds the threshold T and (ii) the best proxy has the same state as x in the last cycle, yet (iii) they have different states in the current cycle. Note that, the first two conditions are what invokes spatial prediction. The third condition causes that prediction to err.

Clearly, the probability of the first condition, $P(P^{best} > T)$, decreases with increasing threshold, T . Let us approximate $P(P^{best} > T) = 1 - T$. The probability of the second condition is simply $1 - 2M(1 - M)$. Since P^{best} is the probability of a correlated change in x (given a change in the proxy), the probability of the third condition (a misprediction) is approximately $1 - P^{best}$. We know that $P^{best} > T$. Assuming that P^{best} could be uniformly anywhere above T , we can replace $1 - P^{best}$ by $(1 - T)/2$. The spatial misprediction probability is then the product of probabilities of the three conditions above, leading to the expression:

$$P_{sm} = (1 - T)[1 - 2M(1 - M)](1 - T)/2 \quad (3)$$

Temporal misprediction occurs when the algorithm resorts to temporal prediction and is wrong. According to the algorithm, temporal (LastKnownState) prediction occurs when (i)

P^{best} exceeds the threshold T , but (ii) the best proxy does not have the same state as x in the last cycle, or when (iii) P^{best} is less than the threshold T . In either case, a misprediction occurs if the state of x changes (hence contradicting LastKnownState). The latter probability can be approximated by F , the fraction of nodes we know of that changed state today. Hence:

$$P_{tm} = (1 - T)[2M(1 - M)]F + [1 - (1 - T)]F \quad (4)$$

Recall that misprediction probability, P_m , is the sum of P_{sm} and P_{tm} . Hence, from Equation (3) and Equation (4), we get:

$$P_m = (1 - T)[1 - 2M(1 - M)](1 - T)/2 + (1 - T)[2M(1 - M)]F + [1 - (1 - T)]F \quad (5)$$

The optimal threshold, T , is one that minimizes the above probability. The equation is a quadratic function of T . Because the coefficient of T^2 is $[1 - 2M(1 - M)]$, which is always positive, the optimal threshold can be found by setting the derivative of the above function to zero and enforcing the natural constraints on values of probability (that they are between 0 and 1). In other words:

$$\frac{dP_m}{dT} = -(1 - T)[1 - 2M(1 - M)] - [2M(1 - M)]F + F = 0 \quad (6)$$

subject to the constraint $0 \leq T \leq 1$. After some rearranging and algebraic manipulation, we get:

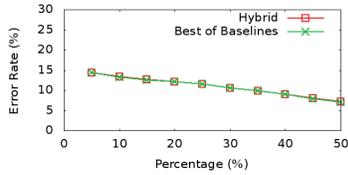
$$T = 1 - F \quad (7)$$

Unfortunately, we do not know the probability of change, F , in advance. In the absence of further knowledge, we can design for $F = 0.5$. In this case, $T = 0.5$.

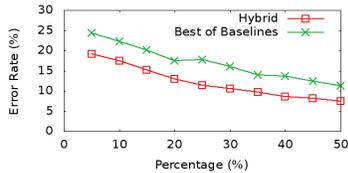
IV. EVALUATION

In this section, we evaluate the hybrid approach presented above versus the baselines described earlier in Section II-A (i.e., Random, LastKnownState, BestProxy, Majority, and ARIMA). For ground truth, we use the same data set, featuring the daily status of gas stations, pharmacies, and food stores in the aftermath of Hurricane Sandy. As before, we opt to predict the status of these PoIs on November 3rd and 8th, as examples of a day of relative calm and a day of significant change. We do so by varying the fraction of PoIs whose state is revealed to the predictor on a given day, and attempting to predict the rest using each of the compared approaches.

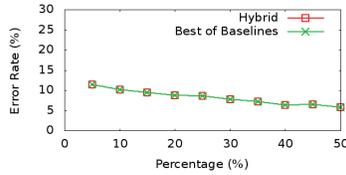
Figures 5-a and 5-b illustrate the accuracy of prediction of gas availability on November 3rd and 8th, respectively. The horizontal axis shows the percentage of PoIs whose state is known on the given day. As before, each point is the average of 50 experiments featuring different random selections of stations whose status is known. On the vertical axis, two curves are compared. One is the hybrid extrapolation algorithm developed in this paper. The second is the *best* of the predictions of the five baselines described in Section II-A. It can be seen that the new algorithm consistently matches or outperforms the best of all others.



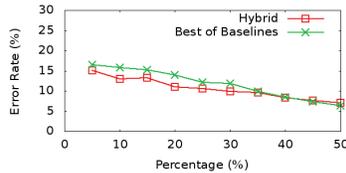
(a) Error rate on November 3rd.



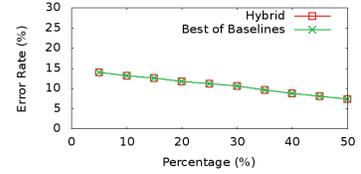
(b) Error rate on November 8th.



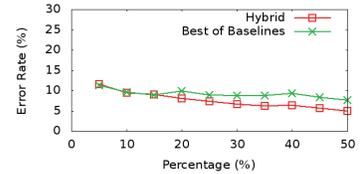
(a) Error rate on November 3rd



(b) Error rate on November 8th



(a) Error rate on November 3rd



(b) Error rate on November 8th

Fig. 5: Predicting gas availability after Sandy

Fig. 6: Predicting food availability after Sandy

Fig. 7: Predicting pharmacy availability after Sandy

Specifically, on November 3rd, the hybrid approach matches the best baseline. This is because it recognizes that change is small, and opts to use LastKnownState, which happens to be the best under the circumstances, as we have seen in Figure 2-a). On November 8th, it outperforms the best baseline, which tends to be BestProxy as we have seen in Figure 2-b. This is because of the new definition of correlation that it uses, which focuses only on changes, per *Insight #2* discussed earlier.

Figures 6-a and 6-b repeat the experiment on the food data set. They illustrate the accuracy of prediction of food availability on November 3rd and 8th, respectively. A similar trend is seen, where the hybrid matches the best baseline on November 3rd and outperforms the best baseline on November 8th. Figures 7-a and 7-b illustrate the same for pharmacies. Further experiments (not shown) demonstrated that the results are largely insensitive to the choice of threshold, T . The superior results presented above can therefore be robustly achieved.

The experimental results presented in this section show that the hybrid approach is as good as or better than the best of all compared algorithms on both November 3rd and November 8th. These two days were selected because of their representative nature, as they exemplified days of calm and days of change, respectively.

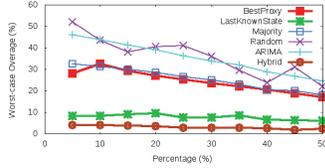
To show that the above results hold true for other days as well, we compute the *worst case* overage amount by which the prediction error of the hybrid approach, as well as the prediction error of each of the five individual baselines, exceeds the best of the five baselines. Hence, an algorithm that behaves as the best of the baselines under all circumstances will have a worst-case overage of zero. Algorithms that are not consistently the best will have a higher worst-case overage. The results are shown in Figure 8, where Figure 8-a, Figure 8-b, and Figure 8-c, are for the case of gas, food, and pharmacy availability prediction, respectively.

In Figure 8, the worst-case overage, for each algorithm, is computed by finding the maximum error overage computed over 10 days of the recovery phase (from November

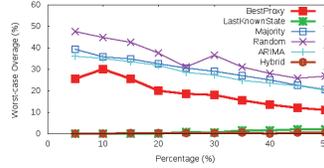
3rd through November 12th). For statistical significance, the performance of each heuristic on each day is first averaged over 50 experiments before the overage is calculated. Consistently with other figures, the horizontal axis shows the percentage of PoIs whose status is known. It is seen that the new Hybrid algorithm has a worst-case overage that is roughly zero. In other words, *it never does worse than the best solution over all days under consideration.*

The figure shows that the overage of other baselines is higher. Their relative prediction (in)accuracy follows roughly the same order in the three data sets. Specifically, LastKnownState is generally the next best algorithm to ours. In the aftermath of disasters, failures take long to fix, so the state changes gradually, making LastKnownState a good predictor most of the time. Errors occur when aftershocks hit or major repairs are made, and are related to the size of such perturbations. BestProxy comes next. Its accuracy depends on how spatially well-correlated the PoI states are. No significant difference is seen between its accuracy in gas and food availability prediction, but pharmacy prediction is better. This can be attributed to the size of the pharmacy data set, shown on the horizontal axis in Figure 1(c). Namely, the number of pharmacies is the largest. Hence, the odds of finding a good proxy are better than with the other data sets. Majority comes next after BestProxy. In scenarios where restoration is quicker, PoIs converge to the majority state faster, and the predictor becomes more accurate. Comparing Figure 1(a), 1(b), and 1(c), we can see that pharmacies and gas are restored the fastest, followed by food, which roughly corresponds to how well Majority works in the three cases. Finally, ARIMA and Random consistently do next-to-worst and worst, respectively, showing little variation across the data sets. This is because their worst-case behavior is random (for ARIMA, it occurs in the very early days), and hence not tightly related to the properties of input data.

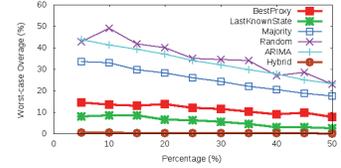
In conclusion, Figure 8 shows that while some prediction algorithms do best under some circumstances, no baseline does consistently well under all circumstances. The contribution of the new approach lies indeed in proposing a method that adapts intelligently between time-based extrapolation and spa-



(a) Worst-case coverage in gas availability prediction error.



(b) Worst-case coverage in food availability prediction error.



(c) Worst-case coverage in pharma availability prediction error.

Fig. 8: Distribution of public services outages

tial extrapolation, matching or outperforming the best baseline solution at all times.

V. RELATED WORK

The work reported in this paper complements a large body of sensor network literature that focused on monitoring and disaster alerts. For example, Werner-Allen et al. deployed three wireless sensor networks on active volcanoes [15]. Li et al. deployed a sensor network for monitoring and alerts in a coal mine [16]. Liu et al. present an automatic and reliable sensor network for firefighter applications [17], which allows a firefighter to carry a small dispenser filled with sensor nodes and deploy them one-by-one in a manner that guarantees reliable communication. The SensorFly project [18] develops a sensor cloud, which consists of many low cost and individually limited mobile sensing devices that only when functioning together can produce an intelligent cloud, in disaster situations such as an earthquake and fire. This paper is different from the above work in leveraging a participatory sensing framework, and considering first responders and volunteers as front-end sensors for data collection.

More importantly, our work focuses on a new problem in participatory sensing. Namely, the problem of automatically filling in the “blind spots” in reported observations. Past research on participatory sensing describes how to aggregate and clean-up collected data. A survey on analytic challenges in the field recently appeared [19]. For instance, CenWits [2] proposes a participatory sensor network to rescue hikers in emergency situations. BikeNet [3] presents a bikers sensor network for sharing cycling related data and mapping the cyclist experience. The Nericell project [4] presents a system that performs rich sensing using smartphones that users carry with them in normal course, to monitor road and traffic conditions. The GreenGPS system [5] provides a service that computes fuel-efficient routes for vehicles between arbitrary end-points, by exploiting vehicular sensor measurements available through the On Board Diagnostic (OBD-II) interface of the car and GPS sensors on smart phones. This paper complements that past work by looking at the important problem of how to fill in the data gaps. This unique challenge comes from the timeliness constraints in disaster response applications. In the absence of urgency, one can eventually fill in the data gaps by sending (or waiting for) more observers. Hence, there is less need to “guess” them. However, in disaster recovery scenarios, there is no time to wait, so the service provider needs to fill in the gaps immediately as best one can.

Also, our work is related to the large body of literature focusing on prediction-based data collection in sensor networks. Le Borgne et al. [20] apply time-series prediction technology

to reduce the communication effort while guaranteeing user-specified accuracy requirements on each sensor nodes in wireless sensor networks. Tulone et al. [21] propose a sensor network comprising sensor nodes and sink nodes. Sensor nodes transmit their local autoregressive models to sink node, and sink node uses the models to predict sensor values without communicating with sensors directly. Krause et al. [22] develop an algorithm called pSPIEL that is capable of measuring the predictive quality of sensor locations and then selecting sensor placements at informative and communication-efficient locations. All those researches utilize similar prediction technology to ours but focus on improving the communication efficiency while maximizing the quality of collected data.

Thanks to the fast development of smartphones and social networks, participatory sensing receives more attention in disaster response applications in recent years. People share their information about the disaster region to social networks and special-purpose services, to help each other beat the disaster together. For instance, popular social networks such as Facebook and Twitter, played an important role after natural disasters such as Japan Tsunami in 2011 [23] and US Hurricane Sandy in 2012 [24]. Many service providers, some notable names including Waze [25] and GasBuddy [26], set up special-purposes services to allow individuals to participate and report the availability of various resources (e.g., gas stations) after Sandy via the web or smartphones. Ushahidi [27] is another notable disaster and crisis management mapping tool. It can be used to collect and visualize data from multiple data streams including text messages, email, twitter and web-forms. However, due to the opportunistic nature of participatory sensing, there are typically “blind points” in the obtained PoI map at any given time point. Our work takes advantage of these services, aiming to complete the estimation of missing world state.

Finally, our system design is related to state of the art sensor selection algorithms that are paired with inference approaches for missing or incomplete data. For example, Aggarwal *et al.* formulate the problem of sensor selection, when redundancy relationships between sensors can be expressed through an information network by using external linkage information. They present methods for efficient sensor selection by using regression models to estimate predictability and redundancy [7]. The problem is extended to dynamic sensor selection in data streams [28]. Similarly, PhotoNet [29] provides a picture-collection service for disaster response applications that maximizes situation-awareness. Kobayashi et al. propose a sensor selection method with fuzzy inference for sensor fusion in robot applications [9]. However, this existing work assumes that correlations between data items are known in advance.

To the best of our knowledge, no previous work has been applied to real-world disaster response scenarios where inference algorithms were investigated that (i) specifically address the bimodal nature of damage propagation and that (ii) require very little training data. Our paper fills in this gap by analyzing the example of New York City gas crisis in the aftermath of Hurricane Sandy via real data traces.

VI. CONCLUSIONS

We presented the design, implementation, and evaluation of an inference-based algorithms for data extrapolation in participatory sensing systems for disaster response applications. It was shown to be capable of accurately predicting the status of PoI sites, when collected data is incomplete. The algorithm exploits correlations among state changes in PoI sites and changes adaptively between temporal and spatial extrapolation. Our experimental results via a real-world disaster response application demonstrate that our algorithm is consistently the best of all compared in terms of prediction accuracy, whereas others may suffer non-trivial degradation. The new algorithm is currently being adapted to more complex prediction tasks (e.g., non-binary variables) and evaluated on new data sets.

ACKNOWLEDGEMENTS

Research reported in this paper was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement W911NF-09-2-0053, DTRA grant HDTRA1-10-1-0120, and NSF grant CNS 13-29886. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," in *In: Workshop on World-Sensor-Web (WSW06): Mobile Device Centric Sensor Networks and Applications*, 2006, pp. 117–134.
- [2] J.-H. Huang, S. Amjad, and S. Mishra, "Cenwits: a sensor-based loosely coupled search and rescue system using witnesses," in *Proceedings of the 3rd international conference on Embedded networked sensor systems*. ACM, 2005, pp. 180–191.
- [3] S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn, and A. T. Campbell, "Bikenet: A mobile sensing system for cyclist experience mapping," *ACM Transactions on Sensor Networks (TOSN)*, vol. 6, no. 1, p. 6, 2009.
- [4] P. Mohan, V. N. Padmanabhan, and R. Ramjee, "Nericell: rich monitoring of road and traffic conditions using mobile smartphones," in *Proceedings of the 6th ACM conference on Embedded network sensor systems*. ACM, 2008, pp. 323–336.
- [5] R. K. Ganti, N. Pham, H. Ahmadi, S. Nangia, and T. F. Abdelzaher, "Greengps: A participatory sensing fuel-efficient maps application," in *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 2010, pp. 151–164.
- [6] R. K. Balan, K. X. Nguyen, and L. Jiang, "Real-time trip information service for a large taxi fleet," in *Proceedings of the 9th international conference on Mobile systems, applications, and services*. ACM, 2011, pp. 99–112.
- [7] C. Aggarwal, A. Bar-Noy, and S. Shamoun, "On sensor selection in linked information networks," in *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conference on*, 2011, pp. 1–8.
- [8] R. Biswas, L. Guibas, and S. Thrun, "A probabilistic approach to inference with limited information in sensor networks," in *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks*, 2004.
- [9] F. Kobayashi, D. Masumoto, and F. Kojima, "Sensor selection based on fuzzy inference for sensor fusion," in *Fuzzy Systems, 2004. Proceedings. 2004 IEEE International Conference on*, vol. 1. IEEE, 2004, pp. 305–310.
- [10] M. G. Kendall, "A new measure of rank correlation," *Biometrika*, vol. 30, no. 1/2, pp. 81–93, 1938.
- [11] R. Shumway and D. Stoffer, *Time series analysis and its applications*, 3rd ed. Springer, 2013.
- [12] G. Elliott, C. Granger, and A. Timmermann, *Handbook of Economic Forecasting*, 1st ed. North-Holland, 2006, ch. 13.
- [13] "Gasoline runs short, adding woes to storm recovery, ny times." http://www.nytimes.com/2012/11/02/nyregion/gasoline-shortages-disrupting-recovery-from-hurricane.html?pagewanted=all&_r=0.
- [14] "All hazards consortium." <http://www.ahcusa.org/>.
- [15] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh, "Fidelity and yield in a volcano monitoring sensor network," in *Proceedings of the 7th symposium on Operating systems design and implementation*. USENIX Association, 2006, pp. 381–396.
- [16] M. Li and Y. Liu, "Underground coal mine monitoring with wireless sensor networks," *ACM Transactions on Sensor Networks (TOSN)*, vol. 5, no. 2, p. 10, 2009.
- [17] H. Liu, J. Li, Z. Xie, S. Lin, K. Whitehouse, J. A. Stankovic, and D. Siu, "Automatic and robust breadcrumb system deployment for indoor firefighter applications," in *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 2010, pp. 21–34.
- [18] A. Purohit, Z. Sun, F. Mokaya, and P. Zhang, "Sensorfly: Controlled-mobile sensing platform for indoor emergency response applications," in *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*. IEEE, 2011, pp. 223–234.
- [19] C. Aggarwal and T. Abdelzaher, *Managing and Mining Sensor Data*. Springer, 2013, ch. Social Sensing.
- [20] Y.-A. Le Borgne, S. Santini, and G. Bontempi, "Adaptive model selection for time series prediction in wireless sensor networks," *Signal Processing*, vol. 87, no. 12, pp. 3010–3020, 2007.
- [21] D. Tulone and S. Madden, "Paq: Time series forecasting for approximate query answering in sensor networks," in *Wireless Sensor Networks*. Springer, 2006, pp. 21–37.
- [22] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg, "Robust sensor placements at informative and communication-efficient locations," *ACM Transactions on Sensor Networks (TOSN)*, vol. 7, no. 4, p. 31, 2011.
- [23] "Japan considers using social networks in disaster situations." <http://www.engadget.com/2012/08/30/japan-considers-using-social-networks-in-disaster-situations>.
- [24] "Twitter / search - #hurricanesandy." <https://twitter.com/search?q=%23hurricanesandy>.
- [25] "Waze: Free gps navigation with turn by turn." <http://www.waze.com/>.
- [26] "Gasbuddy: Find low gas prices in the usa and canada." <http://www.gasbuddy.com/>.
- [27] "Ushahidi platform." <http://ushahidi.com/products/ushahidi-platform/>.
- [28] C. Aggarwal, Y. Xie, and P. Yu., "On dynamic data driven selection of sensor streams," in *ACM Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 2011.
- [29] M. Y. S. Uddin, H. Wang, F. Saremi, G.-J. Qi, T. Abdelzaher, and T. Huang, "Photonet: a similarity-aware picture delivery service for situation awareness," in *Real-Time Systems Symposium (RTSS), 2011 IEEE 32nd*. IEEE, 2011, pp. 317–326.