OPTIMAL SENSOR PLACEMENT FOR URBAN HEAT RISK RESPONSE

Carl Malings[1], Matteo Pozzi[2], Kelly Klima[3], Elie Bou-Zeid[4], Prathap Ramamurthy[5], Mario Bergés[6]

[1] *Carnegie Mellon University: Porter Hall 115, Frew Street, Pittsburgh, PA, 15213, cmalings@andrew.cmu.edu*
[2] *Carnegie Mellon University: Porter Hall 107B, Frew Street, Pittsburgh, PA, 15213, mpozzi@cmu.edu*
[3] *Carnegie Mellon University: Posner Hall 254C, Frew Street, Pittsburgh, PA, 15213, kklima@andrew.cmu.edu*
[4] *Princeton University: E414 Engineering Quad, 59 Olden Street, Princeton, NJ, 08544, ebouzeid@princeton.edu*
[5] *City College of New York: Steinman 229, 160 Convent Av., New York, NY, 10031, pramamurthy@ccny.cuny.edu*
[6] *Carnegie Mellon University: Porter Hall 113, Frew Street, Pittsburgh, PA, 15213, marioberges@cmu.edu*

Urban areas are more at risk from heat waves than their rural surroundings for two main reasons. First, the urban heat island effect causes urban temperatures to be warmer than those of the rural surroundings, worsening the heat wave hazard. Second, the larger number of people in densely populated urban leads to greater heat wave exposure. To mitigate the negative consequences of extreme heat, city officials can conduct response activities such as the issuance of heat advisories. The issuance of targeted advisories to specific areas of a city can be guided by accurate spatio-temporal models of the heat hazard in the area, coupled with assessments of the spatial distribution of vulnerable populations. To support optimal heat advisory issuance in the urban area of Pittsburgh, Pennsylvania, this paper uses a) Gaussian process spatio-temporal probabilistic models of urban temperatures, b) temperature sensor placements informed by the value of information metric, and c) coarse resolution weather models and prior analyses of the daily temperature pattern in the urban area. Results indicate that a high value of information is provided by a few strategically selected measurements in supporting short-term decision-making for heat advisory issuance.

I. INTRODUCTION

High temperatures are well known to have adverse health impacts, which have been documented for recent heat waves in America, Europe, and Australia[1–3]. These impacts are felt with particular intensity in urban areas, due to a combination of factors. Firstly, the urban heat island (UHI) effect tends to increase daytime temperatures in urban areas by on average between 1 and 3°C with respect to surrounding rural areas[4–6]. Secondly, the heat-vulnerable populations in urban areas tend to be larger than in rural areas, due to higher population concentrations. Risk factors such as age, socio-economic status, and whether individuals live alone contribute to high temperature vulnerability[7–9]. Urban areas are more densely populated than rural areas and thus, assuming comparable demographic factors, will have more people exposed to high temperatures. This high hazard and vulnerability combine to create a large risk from high temperatures in urban areas.

To mitigate the consequences of extreme heat in urban areas, heat advisories are sometimes issued to warn individuals to remain indoors and/or seek shelter in air conditioned spaces. Accurate location-specific predictions of temperature are a critical factor to ensuring that these warnings are issued to maximum effect[10]. Recent advances in fine resolution temperature modeling in urban areas, incorporating better representations of urban surface heterogeneity and hydrology, have led to more accurate simulation of the UHI effect with respect to remotely sensed temperatures[11–13]. However, the computational complexity of these models precludes their efficient use for near-term fine resolution temperature forecasting[14–16].

Previous work by the authors has developed a Gaussian process model for urban temperatures, allowing for efficient probabilistic forecasting[17]. This model makes use of a variety of information sources for conditioning the posterior temperature model, including forecasts from coarse resolution weather models for the region in question, prior analysis of fine resolution weather models to determine spatial patterns of heat distribution in urban areas, and local measurements of temperature obtained from weather stations in the region. This model was used by the authors, together with various sensor placement metrics, to optimize the location of sensors for supporting decision-making for the issuance of heat advisories[18]. One of these metrics is the *value of information* (VoI), a decision-theoretic measure assessing how additional information can reduce costs for managing an uncertain system, with respect to how the system would be managed without this additional information[19]. In general, evaluation of the VoI is computationally expensive; however, prior work by the authors has

identified certain assumptions on the problem structure under which this metric can be tractably computed, even in larger systems[20]. These assumptions were then used by the authors, together with a vulnerability assessment of the city of Pittsburgh, Pennsylvania conducted by an author[21], to define a decision-making problem for heat advisory issuance, with associated consequences. This decision-making problem was used to evaluate the VoI metric for temperature measurements in the city supporting heat advisory issuance, proposing an optimal set of temperature sensor placements[18].

This paper extends the results of these previous papers in three ways. First, the optimal sensor placement set of the previous work was developed under the assumption that accurate assessments of the average temperature in the city (obtained using a coarse resolution weather model) and assessments of the spatial temperature patterns in the city (obtained from fine resolution weather simulations) were available in addition to the measurements of the local weather stations. In general, this additional information may not be readily available. In this paper, we analyze how the VoI of the optimized sensing set is affected if these data are not available, or only partially available. Furthermore, another sensor placement set is proposed, optimized under the assumption that none of this additional information is available.

Second, in previous work, the VoI of the measurements was assessed assuming there was no lead-time for the forecasting, i.e. that data available up to and including a certain time were used to guide the issuance of heat advisories at that time. However, in general, these advisories may be issued ahead of time, and therefore a certain lead-time must be considered for model forecasting. In this paper, various lead-times are considered, and the effect of the prediction lead time on the VoI of a set of measurements is investigated.

Third, the metric of weighted prediction error is used in previous work as an alternative to the VoI metric for optimizing sensor placement. In this paper, the relationship between these two metrics is further investigated. The VoI for sensor placement sets optimized by the weighted prediction error metric is computed, and vice versa, in order to determine how well the computationally efficient weighted prediction error metric might be used to approximately optimize sensors for supporting a decision-making problem in the context of urban temperature prediction and heat advisory issuance.

II. SUMMARY OF URBAN TEMPERATURE MODELING AND OPTIMAL MONITORING

This section summarizes the results of previous work[17,18] in developing a Gaussian process model for urban temperatures and making use of this model to optimize the placement of sensors for temperature monitoring and model updating. We denote by $T(\mathbf{x}, t)$ the surface temperature in degrees Celsius at location $\mathbf{x}$ and time t within an urban area. This temperature is decomposed into three components as follows:

$$T(\mathbf{x}, t) = T_0(t) + T_1(\mathbf{x}, t) + T''(\mathbf{x}, t) \qquad (1)$$

where $T_0(t)$ is the average temperature in the region at time t, $T_1(\mathbf{x}, t)$ represents the spatially and temporally varying pattern of temperature that recurs following a daily cycle, and $T''(\mathbf{x}, t)$ is the temperature residual. In the context of urban temperatures, $T_0(t)$ can be interpreted as the average city-wide surface temperature, and $T_1(\mathbf{x}, t)$ as the expected systematic difference in temperature between this average and the local temperature at a given time of day due to the UHI effect (although other recurring patterns at finer scales, not only those related to the UHI, will also be incorporated into this term).

Based on fine resolution simulations of the historical temperature patterns of Pittsburgh, obtained using the National Center for Atmospheric Research's Weather Research and Forecasting – Advanced Research model[22] together with the Princeton Urban Canopy Model[13], Gaussian process models are created and calibrated for each component of the temperature listed in Eq. (1). A Gaussian process model generalizes the multivariate Gaussian distribution to a continuous spatial and/or temporal domain, and for any finite set of coordinates within this domain, defines a multivariate Gaussian distribution over the variables associated with these coordinates for a spatio-temporal random field[23]. The Gaussian process model for temperature is denoted as:

$$T(\mathbf{x}, t) \sim \mathcal{GP}\left(\mathrm{M}_T(\mathbf{x}, t), \mathrm{K}_T\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right)\right) \qquad (2)$$

where $\mathrm{M}_T(\mathbf{x}, t)$ is the mean function of spatio-temporal coordinate $\{\mathbf{x}, t\}$ and $\mathrm{K}_T\left(\mathbf{x}_i, t_i, \mathbf{x}_j, t_j\right)$ is the covariance function of the pair of spatio-temporal coordinates $\{\mathbf{x}_i, t_i\}$ and $\{\mathbf{x}_j, t_j\}$. As the sum of Gaussian processes is itself a Gaussian process, the models calibrated for each of the components of temperature are summed to define this overall model.

For a finite spatial domain $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_{n_X}\}$ and temporal domain $\tau = \{t_1, \ldots, t_{n_\tau}\}$, this Gaussian process model defines a multivariate Gaussian distribution for $\mathbf{T}$, the vector of temperatures at each space-time coordinate in the domain:

$$\mathbf{T} \sim \mathcal{N}(\boldsymbol{\mu}_T, \mathbf{K}_T) \qquad (3)$$

where $\boldsymbol{\mu}_T$ is the mean vector and $\mathbf{K}_T$ the covariance matrix of the distribution, obtained by evaluating the mean and covariance functions respectively over the spatial and temporal domain. This distribution constitutes the prior temperature forecasting model for the region in question.

II.A. Measurements and Model Updating

The prior model for temperature may be updated to a posterior model following measurements of the temperature field or any of its component parts. Three potential types of measurements are considered here. First, a measure of the average temperature over the domain at a certain time, i.e. a measure of $T_0(t)$, is considered. This measurement is modeled as:

$$Y_{T_0}(t) = \frac{1}{n_X}\sum_{\mathbf{x}\in \mathrm{X}} T(\mathbf{x}, t) + \epsilon_{T_0}(t) \qquad (4)$$

That is, the measurement can be expressed as the average of the temperature field at time t over the spatial domain, plus a Gaussian random error $\epsilon_{T_0}(t)$. This error is modeled as a Gaussian random variable at each time. These measurements can be interpreted as estimates of the regional average temperature, obtained from coarse resolution weather simulation models, such as those used by the National Weather Service for forecasting. The estimated quality of these forecasts, along with any systematic bias due to a difference in spatial domain between the coarse resolution model and the region being modeled with this Gaussian process method, can be accounted for through proper definition of the distribution on the Gaussian random error; the reader is referred elsewhere for further details[23,24].

Second, measurements of the cyclic temperature pattern, $T_1(\mathbf{x}, t)$, are possible. These measurements may be obtained, for example, from previous simulations of temperatures by fine resolution models for the region in question, such as the simulations used to calibrate the Gaussian process model of the temperature field[17]. Such measurements can be expressed as:

$$Y_{T_1}(\mathbf{x}, t_{\text{cycle}}) = \frac{1}{n_{\text{cycle}}}\sum_{\{t:\, t \bmod \Delta t = t_{\text{cycle}}\}\in\tau}[T(\mathbf{x}, t) - T_0(t)] + \epsilon_{T_1}(\mathbf{x}, t_{\text{cycle}}) \qquad (5)$$

That is, these measurements are obtained by averaging over n_{cycle} days, at a specific time t_{cycle} within the diurnal cycle of $\Delta t = 24$ hours, the differences of the temperature field and the spatial average of the temperature at this time. This is then an estimate from data of the average (over multiple diurnal cycles) difference between the temperature at location $\mathbf{x}$ and the average (over space) temperature, at a given time of day, where n_{cycle} is the number of days in the data. Due to the definition of the Gaussian process model for the cyclic temperature pattern[17], it is assumed that there is no long-term change in these temperature patterns. However, in practice, simulation results will become outdated as changes in the urban area or surrounding land use affect the spatial distribution of temperatures[25]. Error in these estimates resulting from insufficient data or lack of recent simulation results can be accounted for with appropriate definition of Gaussian error $\epsilon_{T_1}(\mathbf{x}, t_{\text{cycle}})$ in these pattern measurements.

Third, local measurements of the temperature field can be incorporated as noisy measurements of $T(\mathbf{x}, t)$:

$$Y_T(\mathbf{x}, t) = T(\mathbf{x}, t) + \epsilon_T(\mathbf{x}, t) \qquad (6)$$

where $\epsilon_T(\mathbf{x}, t)$ is the Gaussian noise associated with the measurement of temperature at this location and time. Such data are, for example, available from temperature gauges and weather stations, which provide estimates of the local temperature at their locations.

All of the above measurements can be expressed as linear combinations of elements of the temperature field $T(\mathbf{x}, t)$ together with the noise of the measurements themselves. Let Y denote a set of measurements of the forms described above, with vector $\mathbf{y}$ as the values of these measurements. This vector is expressed as:

$$\mathbf{y} = \boldsymbol{\beta}_Y \mathbf{T} + \boldsymbol{\epsilon} \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}_\epsilon, \mathbf{K}_\epsilon) \qquad (7)$$

where, through appropriate definition of matrix $\boldsymbol{\beta}_Y$ and of the distribution for $\boldsymbol{\epsilon}$, any of the above measurements can be expressed and included in set Y. Note that each measurement error is assumed to be Gaussian, therefore the joint distribution of these errors is a multivariate Gaussian. As a linear combination of Gaussians, the distribution for $\mathbf{y}$ is itself a Gaussian, with mean $\boldsymbol{\mu}_Y = \boldsymbol{\beta}_Y \boldsymbol{\mu}_T + \boldsymbol{\mu}_\epsilon$ and covariance $\mathbf{K}_Y = \boldsymbol{\beta}_Y \mathbf{K}_T \boldsymbol{\beta}_Y^{\mathrm{T}} + \mathbf{K}_\epsilon$. Conditional to such a measurement, the prior distribution of the temperature field is updated to a posterior distribution:

$$\mathbf{T}|\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_{T|\mathbf{y}}, \mathbf{K}_{T|Y}) \qquad (8)$$

The posterior mean and covariance are:

$$\boldsymbol{\mu}_{T|\mathbf{y}} = \boldsymbol{\mu}_T + \mathbf{K}_{TY}\mathbf{K}_Y{}^{-1}(\mathbf{y} - \boldsymbol{\mu}_Y) \quad \mathbf{K}_{T|Y} = \mathbf{K}_T - \mathbf{K}_{TY}\mathbf{K}_Y{}^{-1}\mathbf{K}_{TY}^{\mathrm{T}} \tag{9}$$

where $\mathbf{K}_{TY} = \mathbf{K}_T\boldsymbol{\beta}_Y^{\mathrm{T}}$. Note that $\mathbf{K}_{T|Y}$ is a function of measurement set Y only, while $\boldsymbol{\mu}_{T|\mathbf{y}}$ is a function of the measurements $\mathbf{y}$ as well. In this manner, any measurement of the field described above can be combined with the prior model to yield an updated posterior model of the temperature field conditioned on measurements $\mathbf{y}$.

II.B. Sensor Placement

Using the prior model and updating procedure described, pre-posterior optimization of sensing can be performed. Such optimization allows a set Y of measures to be selected that will best improve the model predictions under a specified performance metric. Two such metrics are considered in this paper: the weighted prediction error and the VoI.

The weighed prediction error metric is defined based on the expected square error of using the posterior mean of the model, updated with set of measurements Y, as an estimate of the temperature field. In a Gaussian process model, this error is a function of the posterior covariance of the model only, and so can be evaluated efficiently[26]. A weight matrix $\boldsymbol{\Omega}$ is also used, such that the prediction errors might be weighted, in order to more heavily penalize incorrect prediction of temperatures in certain areas. The optimal set of measurements, by this metric, is the set which minimizes the expected weighted posterior prediction error[18]:

$$Y^*_{\mathrm{WPE}} = \mathrm{argmin}_Y\ \mathrm{tr}\left(\boldsymbol{\Omega}\mathbf{K}_{T|Y}\boldsymbol{\Omega}^{\mathrm{T}}\right) \tag{10}$$

where $\mathrm{tr}(\cdot)$ denotes the matrix trace. For the purposes of this paper, the weighting for this metric is defined using an analysis of the spatially varying vulnerability of populations in the city of Pittsburgh to extreme heat[21]. This choice of a weighting function is based on the reasoning that accurate prediction of temperatures is more important where populations are more vulnerable, and so prediction errors should be more heavily penalized by the metric in more vulnerable areas.

A second metric used for optimizing senor placements is the VoI, which quantifies the benefits that additional information can provide in terms of reducing costs or losses in a decision-making problem in an uncertain environment[19]. In the context of urban temperatures, an example decision-making problem is analyzed for the issuance of location-specific heat advisories; for location $\mathbf{x}$ and time t, an action $A(\mathbf{x}, t)$ is undertaken as to whether or not to issue an advisory. If no advisory is issued, and the local temperature exceeds an upper limit of T_{limit} (set to 30°C in this problem), a penalty of $C_f(\mathbf{x})$ is incurred for that location, corresponding to the consequences of failing to warn the population to take necessary precautions. If, on the other hand, an advisory is issued, a cost of $C_r(\mathbf{x})$ is incurred regardless of temperature; this reflects the consequences of the warning itself, in terms of costs for opening heat shelters, loss of productivity, increased energy usage, etc. These costs are calibrated using the same vulnerability analysis used for the weighed prediction error metric[21], with costs $C_f(\mathbf{x})$ assumed to be proportional to the vulnerability, and $C_r(\mathbf{x})$ selected arbitrarily to be half the maximum value of $C_f(\mathbf{x})$ for all locations. Based on this problem statement, the local loss function is:

$$\mathrm{L}_{\mathbf{x},t}\big(T(\mathbf{x},t), A(\mathbf{x},t)\big) = C_f(\mathbf{x})\mathbb{I}\big[T(\mathbf{x},t)\big(1 - A(\mathbf{x},t)\big) > T_{limit}\big] + C_r(\mathbf{x})A(\mathbf{x},t) \tag{11}$$

where $\mathbb{I}[\cdot]$ is the indicator function, taking on value 1 when its argument is true and 0 otherwise (it is assumed that issuing a heat advisory corresponds to $A(\mathbf{x}, t) = 1$, while not issuing one corresponds to $A(\mathbf{x}, t) = 0$).

Based on this loss function, the VoI metric is defined as[18,19]:

$$\mathrm{VoI}(Y) = \sum_{t\in\tau}\sum_{\mathbf{x}\in X}\big[\min_A \mathbb{E}_T \mathrm{L}_{\mathbf{x},t}\big(T(\mathbf{x},t), A(\mathbf{x},t)\big) - \mathbb{E}_Y \min_A \mathbb{E}_{T|\mathbf{y}} \mathrm{L}_{\mathbf{x},t}\big(T(\mathbf{x},t), A(\mathbf{x},t)\big)\big] \tag{12}$$

where $\mathbb{E}_T$ denotes the statistical expectation with respect to the temperature field, and $\mathbb{E}_Y$ the expectation with respect to measurements of the set Y. Namely, the VoI is the sum, across space and time, of the difference in expected value of the loss when actions are taken without and with knowledge gained through the measurement set Y. Under the assumption that the loss can be expressed locally, as in Eq. (11), this metric can be evaluated in a relatively efficient manner[27]. Under this metric, the optimal measurement set is that which maximizes the VoI minus the cost of acquiring that information, denoted $\mathrm{C}(Y)$:

$$Y^*_{\mathrm{VoI}} = \mathrm{argmax}_Y\ \mathrm{VoI}(Y) - \mathrm{C}(Y) \tag{13}$$

The optimization problems of Eqs. (10, 13) are examples of combinatorial problems, which are generally intractable to solve exactly in all but the smallest systems, due to the exponential growth in the number of possible observation sets Y

which must be evaluated[28]. Because of this, approximate heuristics for optimization are often used. In this work, where these optimization problems are solved, we make use of the greedy optimization algorithm, defined as building the optimal set of measurements one element at a time by, at each step of the algorithm, adding the measurement to the set which most improves the objective. In general, there are no guarantees on the performance of this approach; however for the weighted prediction error metric, the property of submodularity satisfied by this metric provides a guarantee that the solution obtained by greedy optimization will be at least 63% of the optimal solution value[26]. Furthermore, for the VoI metric, although no such guarantees exist, prior work suggests this algorithm performs well for the type of loss function considered here[27].

III. ANALYSIS AND RESULTS

This section expands on previous results for the optimal placement of sensors to support temperature prediction and decision-making for the city of Pittsburgh[18]. First, a comparison is made between the VoI provided by sets of local sensors with and without information on the average regional temperature and cyclic temperature pattern. Second, the effect of prediction lead-times on VoI is investigated. Third, a comparison is made between the VoI and weighted prediction error metrics. For the purposes of the investigation, error terms in Eqs. (4, 5, 6) are assumed to be negligibly small (on the order of 0.02°C) and independent for different times and locations, except in the case of Section III.B, where the standard deviation of the error term in Eq. (4) for predictions of the future regional average temperature is assumed to increase with the prediction lead time.

III.A. Optimal Sensor Placement With and Without Additional Information

An optimal set of sensor placements in Pittsburgh is proposed based on the VoI metric[18]. This set is optimized under the assumption that additional information on the average temperature and cyclic temperature patterns and the data collected from local sensors would be available to support decision-making. All of these data sources may not always be available. Therefore, in this paper a second sensor set is proposed, following optimization using the VoI metric under the same decision-making problem, but without the assumption of information on the average temperature or cyclic temperature patterns being available.

We denote by Y_{T_0} measurements of the regional average temperature, as in Eq. (4), and we denote by Y_{T_1} measures of the cyclic temperature pattern, as in Eq. (5), from the training temperature data simulated for this region[17]. The set of local temperature measurements, as in Eq. (6), optimized following Eq. (13) under the assumption that both Y_{T_0} and Y_{T_1} are also available, is denoted as $Y^*_{T,A}$. The alternative local measurement set, optimized assuming no other information is available, is denoted as $Y^*_{T,B}$. Both of these sets are indicated in Fig. 1. Finally, for comparative purposes, we denote by Y_T the set of all local measurements, i.e. the set of measurements taken at each of the discrete spatial locations in X. The VoI provided by this set of measurements, together with measures of the regional average and cyclic pattern, constitute the set of all possible measurements considered in this work and thus denote the value of complete information, an upper limit to the VoI provided by any of the other sets considered here.

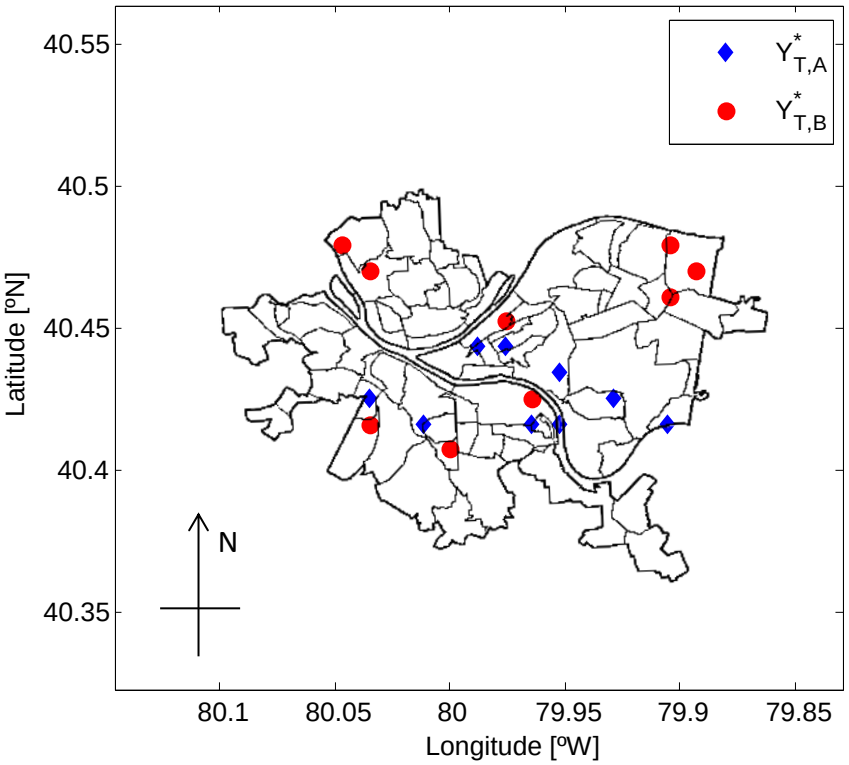

Fig. 1. Optimized locations of local temperature measurements in Pittsburgh, PA.

For the various measurement sets, VoI is computed (as listed in the second column of Table 1) and compared with the value of complete information (listed in the third column). Note that an accurate prediction of the regional average temperature alone provides 82% of the value of complete information, indicating that even this information is quite valuable by itself in supporting decision-making. This is useful since such regional-averaged temperatures are easily obtainable from regular weather forecasts. A prior assessment of the cyclic temperature pattern provides no value by itself, but when combined with the regional average temperature provides 95% of the value of complete information. These cyclic temperature patterns, unlike regional-averaged temperatures, require fine-scale modeling for that particular city, and so are more difficult to obtain.

As expected, $Y^*_{T,B}$ slightly outperforms $Y^*_{T,A}$ in the absence of additional information on the regional average temperature and cyclic temperature pattern. Note that the set of local measurements $Y^*_{T,A}$ can still be used by itself, even though it was optimized under the assumption that additional information from Y_{T_0} and Y_{T_1}would also be available. Because $Y^*_{T,A}$ is not being used as intended in this case, its performance is suboptimal; however, the set $\{Y_{T_0}, Y_{T_1}, Y^*_{T,A}\}$ outperforms $\{Y_{T_0}, Y_{T_1}, Y^*_{T,B}\}$ when the additional information from Y_{T_0} and Y_{T_1} is available. This is likely because some information gathered as part of $Y^*_{T,B}$ becomes redundant when measures Y_{T_0} and Y_{T_1} are also available. Note in Fig. 1 that the spatial spread of $Y^*_{T,B}$ is greater than that of $Y^*_{T,1}$. In the absence of data on the regional average temperature provided by Y_{T_0}, this spread of local measurements allows for an estimate of this average which is more robust against local perturbations than the more tightly grouped measures of $Y^*_{T,A}$. Furthermore, in the absence of prior information on the temperature pattern provided by Y_{T_1}, the spread of $Y^*_{T,B}$ allows this pattern to be learned from local data. However, when Y_{T_1} is available, this learning is unnecessary. Note, however, that the VoI provided by both measurement sets is nearly the same; this suggests that the VoI metric is not very sensitive to the placements of local sensors under the assumptions made here in defining the heat advisory decision-making problem for the city of Pittsburgh.

It should also be noted that both $Y^*_{T,A}$ and $Y^*_{T,B}$ provide 96% of the value of complete information, even without Y_{T_0} and Y_{T_1}. When these additional measures are added, VoI only increases slightly (up to 97% of the value of complete information). This shows that, on the one hand, the prior model updated with only local information can account for cyclic temperature

patterns and regional temperature trends without direct measurements of these. On the other hand, the direct measurements (i.e. Y_{T_0} and Y_{T_1}) together provide 95% of the value of complete information, without any local temperature measures. These results demonstrate that the value provided by local measurements alone is comparable to that provided by an accurate predictive model of the regional temperature combined with a prior fine resolution analysis of the region.

TABLE I. Comparative results for VoI and weighted prediction error of various measurement sets

Measurement Set	VoI	Percentage of Complete Information	Weighted Prediction Error [°C]
$\{Y_{T_0}\}$	21.31	82%	1.94
$\{Y_{T_1}\}$	0.00	0%	162.49
$\{Y_{T_0}, Y_{T_1}\}$	24.67	95%	0.71
$\{Y^*_{T,A}\}$	25.02	96%	0.85
$\{Y^*_{T,B}\}$	25.03	96%	0.67
$\{Y_{T_0}, Y^*_{T,A}\}$	25.06	96%	0.80
$\{Y_{T_0}, Y^*_{T,B}\}$	25.04	96%	0.66
$\{Y_{T_0}, Y_{T_1}, Y^*_{T,A}\}$	25.27	97%	0.31
$\{Y_{T_0}, Y_{T_1}, Y^*_{T,B}\}$	25.13	96%	0.26
$\{Y_{T_0}, Y_{T_1}, Y_T\}$	26.10	100%	0.07
$\{Y^*_{T,E}\}$	24.86	95%	0.54
$\{Y_{T_0}, Y^*_{T,E}\}$	24.90	95%	0.53
$\{Y_{T_0}, Y_{T_1}, Y^*_{T,C}\}$	24.91	95%	0.21

III.B. Value of Information with Prediction Lead Time

In the previous section, VoI is assessed for same-time temperature prediction and warning issuance. That is, decisions about heat advisory issuance are made using information collected up to and including the time of the advisory. If decisions are instead made ahead of time, using a certain lead time, predictions will be less accurate, and therefore provide a lower VoI. Fig. 2 displays results for how the VoI is affected by the prediction lead time, up to 48 hours. Values are shown for two measurement sets; the optimal set without additional information, $\{Y^*_{T,B}\}$, and the optimal set including additional information, $\{Y_{T_0}, Y_{T_1}, Y^*_{T,A}\}$. Here, Y_{T_0} represents a forecast for the future average temperature, the standard error of which is assumed to increase linearly with the lead time, such that for 12 hours ahead, the standard deviation of ϵ_{T_0} in Eq. (4) is 1℃.

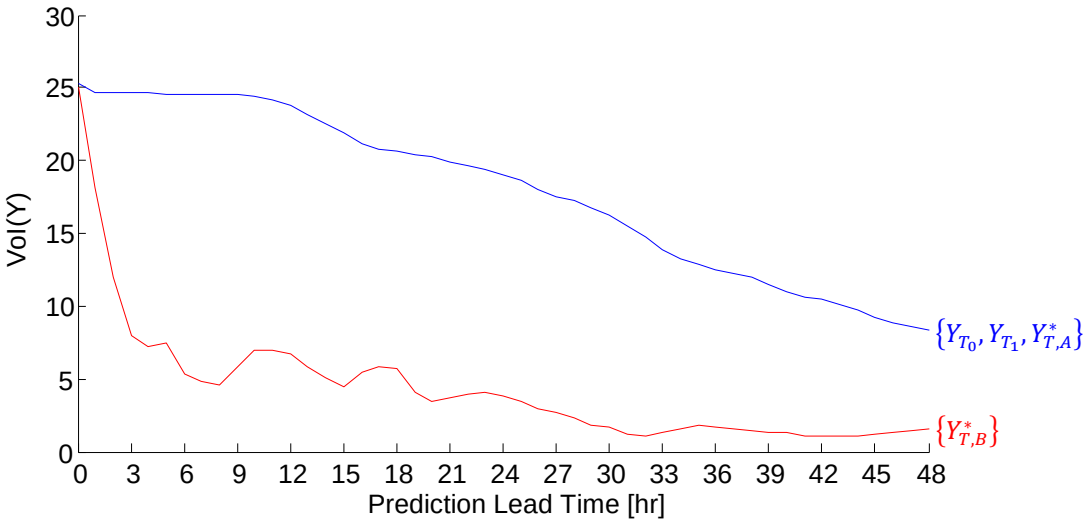

Fig. 2. VoI versus prediction lead time for two measurement sets.

As expected, VoI tends to decrease with prediction lead time. For measurement set $\{Y_{T_0}, Y_{T_1}, Y^*_{T,A}\}$, this decrease is gradual, as relatively accurate predictions of the regional average temperature are available to the model, even up to two days ahead. However, without this future forecasting information provided by Y_{T_0}, the VoI for measurement set $\{Y^*_{T,B}\}$ decreases more rapidly with time, dropping below half its initial value at 2 hours ahead, and is about a fifth of its initial value at one day ahead. Interestingly, while there is a high degree of correlation between temperature values at 24 hour intervals, due to cyclic temperature patterns and the daily temperature cycle throughout the region, the VoI does not increase at 24 hours prediction lead time. This is likely because of the temperature variance due to the residual $T''(\mathbf{x}, t)$, which exhibits very low correlation at time differences greater than about 12 hour. These results indicate how important accurate regional temperature predictions can be to ahead-of-time decision-making, as the value provided by local temperature measurements alone decays rather quickly with time.

III.C. Comparison of Value of Information and Weighted Prediction Error Metrics

To compare the VoI metric with the less computationally intensive weighted prediction error metric, a set of sensors of the same size as $Y^*_{T,A}$ or $Y^*_{T,B}$ is optimized by this latter metric, and denoted $Y^*_{T,C}$. Note that, while the VoI metric can be traded off against sensing cost as in Eq. (13) to determine the optimal number of measurements to include, the weighted prediction error metric will only decrease as more measures are included, and cannot be readily traded off against cost. Therefore, the number of measures to include in the set is arbitrary (in reality, it would be dictated by the available funds and resources for deploying more sensors). This set is indicated in Fig. 3.

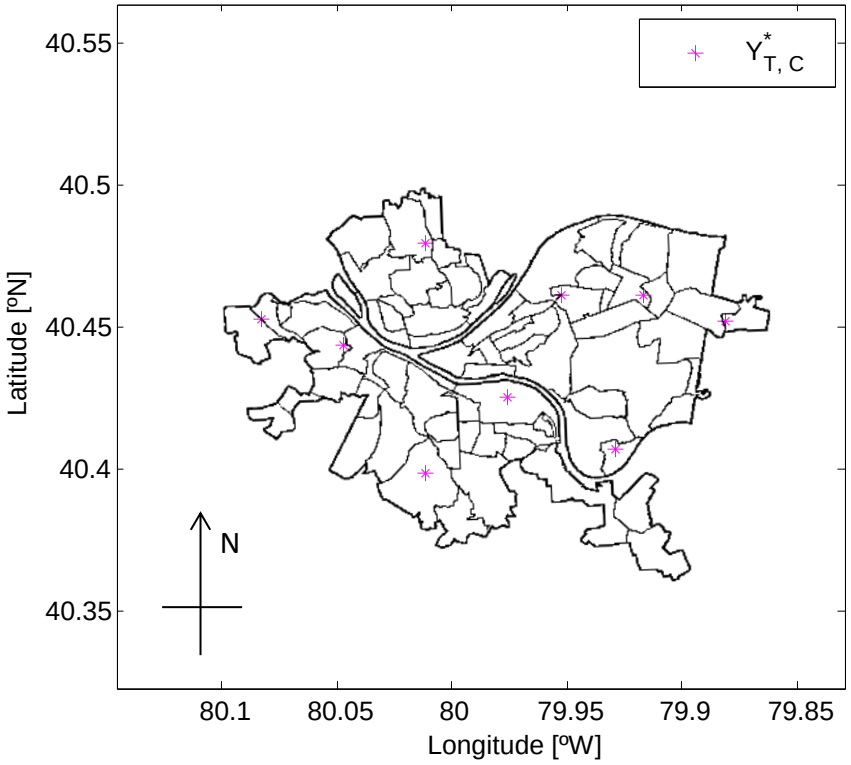

Fig. 3. Local temperature measurement locations optimized by the weighted prediction error metric in Pittsburgh, PA.

At the bottom of Table 1, the VoI for this set of measurements is listed, with and without additional information about the average temperature and cyclic temperature patterns. Additionally, the weighted prediction error metric is evaluated for these sets, and is listed in the last column of Table 1. This weighted prediction error metric is presented in units of degrees Celsius and represents the average (in space and time) of the root mean square error of the temperature prediction multiplied by the local vulnerability index[21] (ranging from 1 to 6). In this way, prediction errors in more vulnerable areas are weighted more heavily when computing the average.

These results show that, even though sensor placements optimized using the weighted prediction error metric provide less VoI that those optimized using that metric directly, their value is still comparable, corresponding to 95% of the value of complete information. This further supports the supposition that the VoI metric is not very sensitive to sensor placements in the case investigated here. Furthermore, the spread in prediction error for different sensor placements is much greater than that of the VoI metric. For example, while the difference between values of information provided by measurements $Y^*_{T,A}$ and $Y^*_{T,C}$ is 0.6% (of the higher value), the difference in the weighted prediction error between these sets is 37%. Overall, these results indicate that, while sensor placements optimized under the weighted prediction error metric also perform fairly well under the VoI metric, the reverse does not hold. However, as the weighted prediction error metric is the more computationally efficient of the two, this result provides some support for the use of this metric as a proxy for VoI in supporting sensor placement over large areas.

IV. CONCLUSIONS

This paper investigates the VoI provided by various temperature measurements in the city of Pittsburgh for supporting decision-making concerning the issuance of extreme heat advisories. Results of this analysis indicate that a relatively high VoI is provided by accurate forecasts of the regional average temperature alone. However, strategically selected local measurements can also provide a comparably high VoI, even without additional information on the regional average temperature and cyclic temperature patterns. When decision-making is done ahead of the latest data collection, VoI decreases with increased lead times, and this decrease can be drastic if accurate predictions of the regional average temperature are not available. Finally, the VoI provided by measurements optimized by the computationally efficient weighted prediction error metric is comparable to that provided by measurements optimized by the VoI metric directly. It should be noted, however, that all these results have been obtained for the analysis of the single urban area of Pittsburgh, under certain assumptions on the structure of the heat advisory decision-making problem considered, and may not generalize to other urban areas. Characteristics of certain urban areas, such as higher heterogeneity in temperatures and/or vulnerability of the population, may make the VoI in these areas more sensitive to different sources of information or to the placement of sensors for local temperature measurement.

ACKNOWLEDGMENTS

This work was funded in part by the Metro21 Initiative at Carnegie Mellon University. Elie Bou-Zeid is funded by the Helen Shipley Hunt Fund of Princeton University and by US National Science Foundation's Sustainability Research Network Cooperative Agreement # 1444758. The authors would also like to thank S. Harris for contribution of the project title (Surface Heat Assessment for Developed Environments) and acronym (SHADE).

REFERENCES

1. J.-M. ROBINE, S. L. CHEUNG, S. LE ROY, H. VAN OY, C. GRIFFITHS, J.-P. MICHEL, and F. R. HERRMANN, "Death toll exceeded 70,000 in Europe during the summer of 2003," *C. R. Biol,*. **331**, 171–178 (2008).
2. K. ANDREWS, "The consequences of heatwaves in Australia," School of Earth Sciences, Macquarie University (1994).
3. R. DAVIS, P. KNAPPENBERG, P. MICHAELS, and W. NOVICOFF, "Changing Heat-Related Mortality in the United States," *Environ. Health Perspect.*, **111**, 1712–1718 (2003).
4. T. R. OKE, *Applied Climatology: Principles & Practices*, pp. 273–287, Routledge (1997).
5. Sailor, D. J. Urban Heat Islands, Opportunities and Challenges for Mitigation and Adaptation. (2002).
6. US EPA, *Reducing Urban Heat Island: Compendium of Strategies - Urban Heat Island Basics*, U.S. Environmental Protection Agency (2012).
7. S. HAJAT, R. S. KOVATS, and K. LACHOWYCZ, "Heat-related and cold-related deaths in England and Wales: who is at risk?" *Occup. Environ. Med.*, **64**, 93–100 (2006).
8. J. K. ROSENTHAL, P. L. KINNEY, and K. B. METZGER, "Intra-urban vulnerability to heat-related mortality in New York City, 1997–2006," *Health Place*, **30,** 45–60 (2014).
9. M. STAFOGGIA, F. FORASTIERE, D. AGNOSTINI, N. CARANCI, F. DE'DONATO, M. DEMARIA, P. MICHELOZZI, R. MIGLIO, M. ROGNONI, A. RUSSO, and C. A. PERUCCI, "Factors affecting in-hospital heat-related mortality: a multi-city case-crossover analysis," *J. Epidemiol. Community Health*, **62**, 209–215 (2008).
10. K. ZHANG, Y.-H. CHEN, J. D. SCHWARTZ, R. B. ROOD, and M. S. O'NEILL, "Using Forecast and Observed Weather Data to Assess Performance of Forecast Products in Identifying Heat Waves and Estimating Heat Wave Effects on Mortality," *Environ. Health Perspect.* (2014).

11. D. LI and E. BOU-ZEID, "Quality and sensitivity of high-resolution numerical simulation of urban heat islands," *Environ. Res. Lett.*, **9**, 55001 (2014).
12. D. LI and E. BOU-ZEID, "Synergistic Interactions between Urban Heat Islands and Heat Waves: The Impact in Cities Is Larger than the Sum of Its Parts," *J. Appl. Meteorol. Climatol.*, **52**, 2051–2064 (2013).
13. Z.-H. WANG, E. BOU-ZEID, and J. A. SMITH, "A coupled energy transport and hydrological model for urban canopies evaluated using a wireless sensor network," *Q. J. R. Meteorol. Soc.*, **139**, 1643–1657 (2013).
14. P. H. WORLEY, A. A. MIRIN, A. P. CRAIG, M. A. TAYLOR, J. M. DENNIS, and M. VERTENSTEIN, *Performance of the community earth system model*, ACM Press (2011).
15. Z. LIN-JIONG, L. YI-MIN, B. QING, Y. HAI-YANG, and W. GUO-XIONG, "Computational Performance of the High-Resolution Atmospheric Model FAMIL," *Atmospheric Ocean. Sci. Lett.*, **5**, 355–359 (2012).
16. E. SPER DE ALMEIDA and M. BAUER, "Reducing time delays in computing numerical weather models at regional and local levels: a grid-based approach," *Int. J. Grid Comput. Appl.*, **3**, 1–17 (2012).
17. C. MALINGS, M. POZZI, K. KLIMA, E. BOU-ZEID, P. RAMAMURTHY, and M. BERGÉS, "Surface Heat Assessment for Developed Environments: Probabilistic Urban Temperature Modeling" to be submitted to *Environ. Res. Lett.* (2016).
18. C. MALINGS, M. POZZI, K. KLIMA, E. BOU-ZEID, P. RAMAMURTHY, and M. BERGÉS, "Surface Heat Assessment for Developed Environments: Optimizing Urban Temperature Monitoring" to be submitted to *Environ. Res. Lett.* (2016).
19. H. RAIFFA and R. SCHLAIFER, *Applied Statistical Decision Theory*, Harvard University Press (1961).
20. C. MALINGS & M. POZZI, "Conditional entropy and value of information based metrics for sensor placement in infrastructure systems" *Proceedings of the 6th World Conference on Strucutral Control and Monitoring* (2014).
21. K. BRADFORD, L. ABRAHAMS, M. HEGGLIN, and K. KLIMA, "A Heat Vulnerability Index and Adaptation Solutions for Pittsburgh, Pennsylvania," *Environ. Sci. Technol.*, **49**, 11303–11311 (2015).
22. W. C. SKAMAROCK, J. B. KLEMP, J. DUDHIA, D. O. Gill, D. M. BARKER, W. WANG, and J. G. POWERS, *A Description of the Advanced Research WRF Version 2*, National Center for Atmospheric Research (2005).
23. C. E. RASMUSSEN and C. K. I. WILLIAMS, *Gaussian processes for machine learning*, MIT Press (2006).
24. N. CRESSIE and C. K. WIKLE, *Statistics for spatio-temporal data*, Wiley (2011).
25. D. LI, E. BOU-ZEID, and M. OPPENHEIMER, "The effectiveness of cool and green roofs as urban heat island mitigation strategies," *Environ. Res. Lett.*, **9**, 55002 (2014).
26. A. KRAUSE, B. McMAHAN, C. GUESTRIN, and A GUPTA, "Robust Submodular Observation Selection," *J. Mach. Learn. Res.*, **9**, 2761–2801 (2008).
27. C. MALINGS and M. POZZI, "Conditional Entropy and Value of Information Metrics for Optimal Sensing in Infrastructure Systems," *Struct. Saf.*, **60**, 77–90 (2016).
28. A. KRAUSE, *Optimizing Sensing: Theory and Applications*, Carnegie Mellon University (2008).