Environmental Sensing using Land-based Spectrally-selective Cameras and a Quadcopter

Jnaneshwar Das, William C. Evans, Michael Minnig, Alexander Bahr, Gaurav S. Sukhatme, Alcherio Martinoli

Abstract We investigate the reconstruction of an environmental scalar field using robotic mobility and heterogeneous sensing. Using two land-based, immobile, co-located spectrally selective cameras, and a non-contact infrared-based temperature sensor on a quadcopter, we study the problem of reconstructing the surface temperature of the ground under survey. Both land units — a thermographic camera for low-resolution thermal images and a commercial digital camera for high resolution truecolor images — are mounted on an elevated camera rig. We explore methods for field reconstruction using a combination of the three imaging sensors. First, we show that the quadcopter data is correlated with the synoptic snapshots obtained by the thermal imaging camera. Next, we demonstrate upsampling of the low-resolution thermal camera data with truecolor images. This results in high-resolution reconstruction of the temperature field. Finally, we discuss adaptive sampling techniques that utilize the mobility of the quadcopter to 'fill the gaps' in data acquired by the thermal imaging camera. Our work experimentally demonstrates the feasibility of heterogeneous sensing and mobility to effectively reconstruct environmental fields.

1 Introduction

Fast sampling of terrestrial environmental fields is of importance for various studies. In this work, we address rapid sampling of environmental fields using static and mobile imaging sensors. Specifically, we demonstrate effective reconstruction of the surface temperature field of a patch of vegetation. Accurate monitoring of surface temperature is desirable for atmospheric bound-

W. C. Evans, M. Minnig, A. Bahr, and A. Martinoli are with the Distributed Intelligent Systems and Algorithms Laboratory, École Polytechnique Fédérale de Lausanne e-mail: {william.evans, michael.minnig, alexander.bahr,alcherio.martinoli}@epfl.ch · J. Das, and G. S. Sukhatme are with the Robotic Embedded Systems Laboratory and the Dept. of Computer Science, University of Southern California, Los Angeles, CA 90089 USA e-mail: {jnaneshd,gaurav}@usc.edu

ary layer studies over complex terrain [1], to cite one example. We describe an experiment where the surface temperature of a region is observed with images from a land-based thermal imaging camera augmented with true-color images from a parallel-mounted commercial digital camera. In addition, a downward-facing non-contact infrared temperature sensor mounted on a quadcopter serves as a mobile sensing platform. The land-based thermal imaging camera is mounted at an elevation along with the truecolor camera, providing snapshots of the surface temperature and high resolution true-color images respectively. The quadcopter serves as a fast aerial observation platform, allowing rapid sampling of surface temperature using its downward-facing temperature sensor. This provides both speed and flexibility compared to land-based observation platforms such as robotic rovers. This work has three goals: a) investigate upsampling of the thermal camera data using high resolution truecolor images from the digital camera, b) compare the temperature data acquired by the quadcopter with the synoptic thermal image captured from the thermal camera, and c) to explore adaptive sampling strategies that use the synoptic data from the thermal camera to guide the quadcopter to regions of high prediction uncertainty. Our goal is to demonstrate synergistic use of mobile and static sensors for rapid characterization of environmental phenomena. Such a capability is necessary when there are constraints on the use of land-based imaging sensors resulting in sparse data. This can happen due to long distance between test site and the land-based camera, or insufficient elevation of the camera rig. By using *mixed sensing*, we can reconstruct the temperature field at a resolution higher than that provided by the individual sensors.

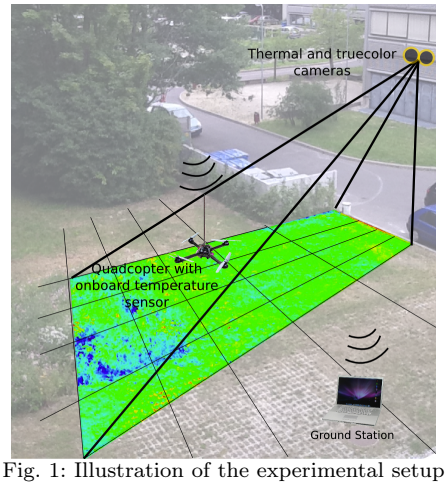

Fig. 1: Illustration of the experimental setup to sample the surface temperature of a patch of land.

The paper is organized as follows. In Section 2 we briefly describe related work. In Section 3, we lay the groundwork for the analysis of the experimental data by describing our technical approach. In Section 4, we describe our field setup followed by analysis of the data from the field trials. We conclude with a summary and discussion of future work in Section 5.

2 Related Work

Ecological monitoring of large farmlands using UAVs has been studied for rapid mapping and classification of vegetation [2]. Adaptive sampling for environmental monitoring has been investigated in the context of intelligent placement of static sensor nodes [3], and informative paths for mobile aquatic platforms [4, 5]. Upsampling of multimodal remote sensing images has been explored to fuse low-resolution hyperspectral images with high-resolution truecolor images [6]. Our work presents an agile setup that provides a quick reconstruction of the environmental field in a region by use of selective spectral-cameras operating at different resolutions, aided by the mobility of a quadcopter.

3 Technical Approach

Our goal is to investigate the use of mixed sensing in the form of static land-based cameras and a quadcopter to rapidly sample the surface temperature of a terrestrial patch of vegetation. We will first describe the land-based camera rig and the quadcopter, followed by a discussion of unwarping and correction of the thermal and truecolor images. We then describe the three contributions of this work for mixed sensing field reconstruction: a) upsampling of thermal images using high-resolution truecolor images (Subsection 3.3), b) comparison of quadcopter data with land-based camera images (Subsection 3.4), and c) an adaptive sampling scheme for the quadcopter to augment land-based sensors (Subsection 3.5).

3.1 Sensing apparatus

The sensing apparatus consists of three imaging sensors, two mounted on a land-based camera rig, and one mounted on a quadcopter. The camera rig consisted of an FLIR A320 thermographic camera with a resolution of 640x480 and a CANON 300D digital camera, both mounted on a tripod head, triggered by a computer to simultaneously capture truecolor and thermal images of the survey site. The aerial platform was an Ascending Technologies Hummingbird

Fig. 2: The camera rig consisting of a thermographic camera and a digital truecolor camera mounted on a pan-tilt head.

quadcopter. It used its onboard computer to log data from a non-contact IR temperature sensor at a rate of 5 Hz along with GPS data at 1 Hz. We chose an IR-based temperature sensor so that the quadcopter would share a similar modality with the thermographic camera. A human operator maneuvered the quadcopter via remote control. The IR sensor captured the emitted IR radiance from a 10 degree view of field. Hence the temperature captured by the sensor was dependent on the height from the target. However, in this work, we did not model the sensor properties, and ignored the height of the quadcopter in flight.

3.2 Image unwarping and correction

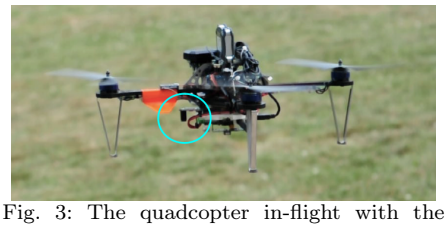

Fig. 3: The quadcopter in-flight with the downward looking infra red temperature sensor.

We first explore fusion of the thermal and truecolor images for quick inspection of the scene. Since the thermal and truecolor cameras are mounted at an elevation, generating a perspective view of the scene, we compute perspective transforms to unwarp the truecolor and thermal images. First, we manually marked four landmarks in the thermal image and the truecolor image. We selected corners of man-made structures such as metal electric poles and concrete slabs because these were easily recognizable in both the thermal and truecolor images.

Next, we must ensure that all data shares a common frame of reference. We proceed by transforming all images to the Earth's coordinate frame, with an approach similar to that used during the image unwarping step. Landmarks in the thermal and truecolor images are used along with landmarks in a satellite truecolor image of the scene obtained from commercial map servers (e.g., Google Earth). The transformation was computed using the OpenCV library. Once a perspective transformation matrix is computed, we obtain the unwarped data points $z = [x, y, t]^T$, where x is the longitude, y is the latitude, and t is the color value of the pixel that was unwarped.

3.3 Upsampling

To demonstrate upsampling of sparse thermal data using dense truecolor data, we subsample a sparse set of points from the unwarped thermal camera

image along with the corresponding truecolor pixel values[1]. Our goal is to learn a model that predicts surface temperature at unobserved locations using the truecolor data. The underlying assumption is that surface patches with similar color will have similar temperature.

Gaussian process regression

We use Gaussian process regression (GPR) [7], a nonlinear Bayesian regression technique commonly used in geostatistics under the name 'Kriging'. It assumes that the samples from the function to be estimated are normally distributed with the covariance between samples given by a 'kernel' or covariance function. We consider the case where the observations are unbiased, that is, the mean of the joint Gaussian distribution is zero. This can be simply satisfied by 'demeaning' the observed data. As a result of its formulation, GPR automatically achieves model regularization from data only, without having to choose model complexity parameters a priori. Additionally, GPR is defined completely by a kernel function that controls how quickly the input space becomes decorrelated. This enforces smoothness constraints in the trained function, ideal for spatial models where usually the observed values for nearby input samples are more correlated than the ones farther apart.

Assume we have training data given by $D = < x_1, y_1 >, ..., < x_n, y_n >$, drawn from the noisy process,

$$y_i = f(x_i) + \epsilon \tag{1}$$

where ϵ is a Gaussian noise term.

Given the training data, posterior mean and covariance for a test data point x_* is given by the following equations,

$$GP_\mu = k_*(K + \sigma_n^2 I)^{-1} y \tag{2}$$

$$GP_\Sigma = k(x_*, x_*) + k_*(K + \sigma_n^2 I)^{-1} k_* \tag{3}$$

The kernel function k is usually chosen to be a squared-exponential function given by,

$$k(x_p, x_q) = e^{-\frac{1}{2\lambda^2}|x_p - x_q|^2} \tag{4}$$

where λ is the decorrelation length scale. The hyperparameters for the kernel function can be learned using iterative methods such as conjugate gradient descent. K is the Gram matrix with its elements given by $K_{pq} = k(x_p, x_q)$, I is the identity matrix, and k_* is the vector of covariances between the test data point and the training data points.

To apply the GPR model to upsample thermal image data, let us consider the unwarped pixels from the truecolor camera given by the vector $X_{tc} =<$

[1] computed using a nearest neighbor search with the thermal camera data points

$Lon_{tc}, Lat_{tc}, R_{tc}, G_{tc}, B_{tc} >$. We use nearest neighbor search to obtain a training dataset of thermal camera image data points and their corresponding truecolor pixel data. This is given by $X_{train} =< Lon, Lat, R, G, B >$, and $Y_{train} = T$. Now, we use GPR to learn a function f that maps prediction points $X_{test} =< Lon_{test}, Lat_{test}, R_{test}, G_{test}, B_{test} >$ where Lon_{test} and Lat_{test} are the longitude and latitude of unobserved locations, and R_{test}, G_{test},and B_{test} are the RGB pixel values from the high-resolution images captured by the truecolor camera corresponding to the query points.

3.4 Comparison of quadcopter and thermal camera data

To investigate the feasibility of using a quadcopter along with a land-based thermal camera, we need to compare the data obtained by the two sensors. Linear interpolation on the quadcopter sensor data on a regular grid immediately reveals visual similarity with the thermal imaging camera data. However, for a quantitative comparison of the two datasets, we compute the Pearson correlation coefficient for quadcopter data and the co-located thermal camera data calculated using nearest-neighbor search on each quadcopter data point. The Pearson correlation coefficient is given by,

$$R = \frac{\sum_{i=0}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=0}^{n}(x_i - \bar{x})^2 \sum_{i=0}^{n}(y_i - \bar{y})^2}} \tag{5}$$

where x and y are the two sensor data streams being compared. Higher values of R indicate a stronger correlation.

The imaging sensors on the thermographic camera and the IR sensor on the quadcopter are not cross-calibrated at the outset. For calibration, we choose a 440 sample data window (90 seconds) of the quadcopter data that is highly correlated with the thermal camera data (R>0.8) and use it to learn a linear mapping from raw quadcopter data to corrected quacopter data, given by, $t_{corr} = a_1 t + a_2$, where t is raw quadcopter data point, t_{corr} is the corrected data point, and a_1 and a_2 are regression coefficients.

3.5 Adaptive sampling with quadcopter

Fig. 4 illustrates the sparsity of data away from the thermal camera once the acquired image is unwarped. This effect is more pronounced when the camera rig is farther away from the test site, or not highly elevated. This scenario will be common in unstructured environments. Also, there are sections of the thermal image (the vertical corners) without any data points. Our goal is to investigate field reconstruction that takes into account the uncertainty of esti-

mates from the thermal camera data as a result of data sparsity. We propose greedily collecting data from regions with high variance. We use a sparser

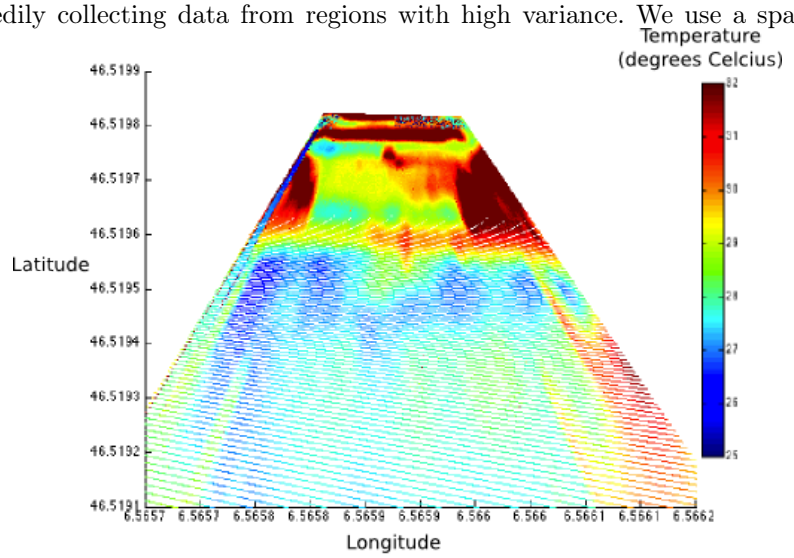

Fig. 4: Perspective corrected (top-view) thermal camera data highlighting sparsity of measurements farther away from the camera (bottom of the image), and complete lack of data in the top right and left corners.

version of the thermal camera data to build a probabilistic spatial model of the temperature field using GPR as described earlier for the upsampling task. Then, we greedily add data points and analyze how many points are needed from the quadcopter data to reduce uncertainty in the reconstructed thermal image.

4 Experiments and Results

We carried out a series of field experiments on the EPFL campus in a 25m x 22m area with a significant portion being vegetation (grass), and a section of gravel path (Fig. 5). This provided a natural environmental field for the measurement of surface temperature. The land based camera-rig was mounted on a tripod at an elevation of 5m and at a distance of 20m from the experiment site. Images were captured simultaneously by both cameras every minute during the course of the experiment. For our analysis, we use one such concurrent snapshot. The quadcopter was operated manually at a mean height of 3.8m by a human pilot for a period of 10 minutes to capture the surface temperature over a lawnmower pattern as shown in Fig. 6. We show results from one of the field trials.

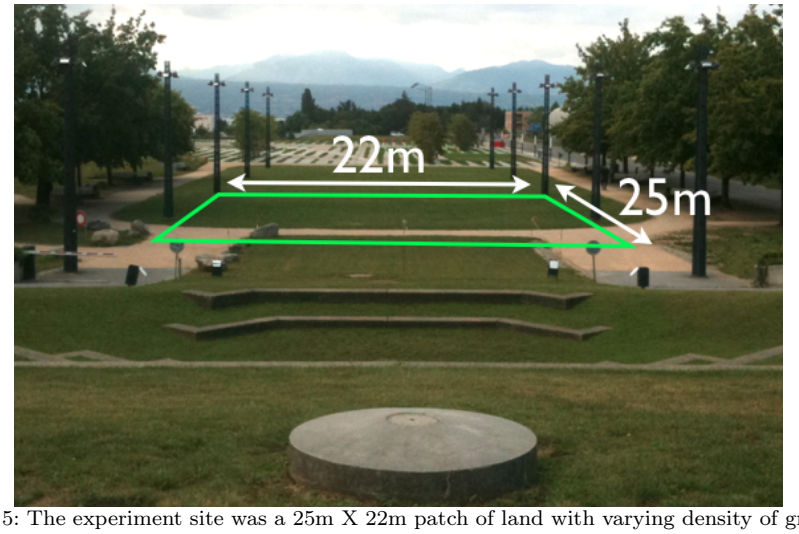

Fig. 5: The experiment site was a 25m X 22m patch of land with varying density of grass, and a gravel path. The camera-rig was at a height of 9m from the test site, at a distance 30 m from the nearest edge of the test site.

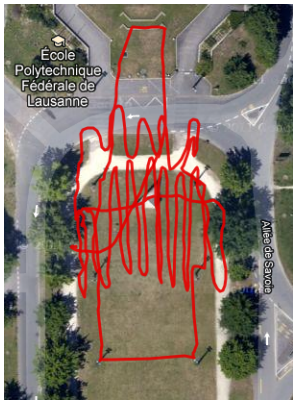

(a) The quadcopter path (solid red path) from one of the experiment runs. A human operator maneuvered the quadcopter to carry out a 'lawnmower' pattern.

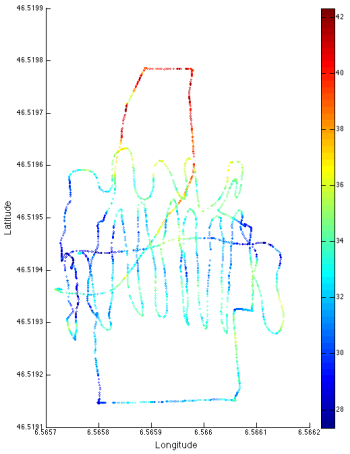

(b) Quadcopter temperature data points from the field trial.

Fig. 6: Data from quadcopter field trial

4.1 Image unwarping and correction

Fig. 7 shows the result of using visual landmarks in the truecolor and thermal images to find the transformation between the two cameras, which was then used to generate an overlaid image of the scene showing both truecolor and thermal images. This is useful as an initial overview of the scene and can be obtained in realtime. Next, we obtained a remotely-sensed truecolor image of

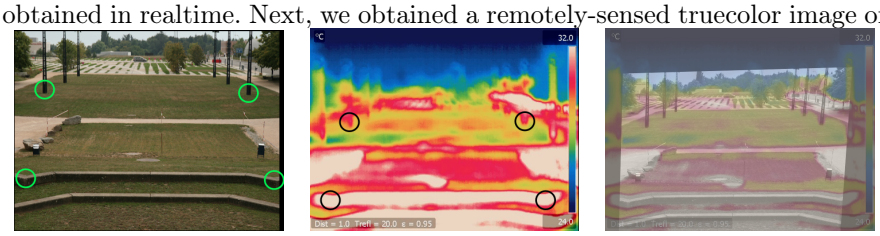

(a) Truecolor image from Canon 300D digital camera.

(b) Thermal image from FLIR A320 thermographic camera.

(c) Truecolor and thermal image superimposed after image registration.

Fig. 7: Landmarks in the truecolor and thermal images used to generate an overlaid image for quick survey of the scene.

the scene from a commercial map-server. From this, we used visual landmarks to compute a perspective transform to unwarp the images from the thermal camera and the digital truecolor camera. Fig. 8 shows the unwarped top view of the truecolor and thermal images. We performed interpolation on the GPS-tagged surface temperature data collected by the quadcopter for initial visual comparison of the two datasets. Fig. 8b shows this image.

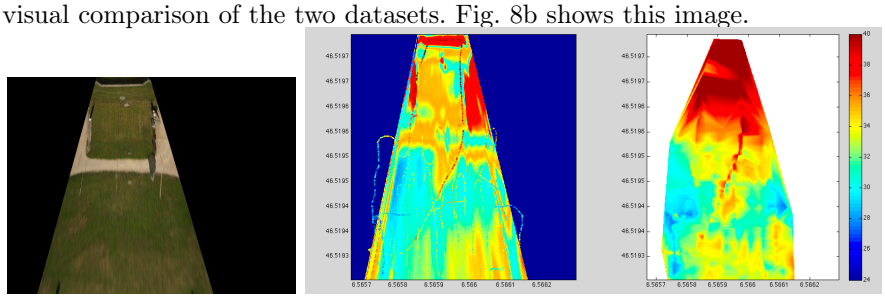

(a) Unwarped truecolor image.

(b) Unwarped thermal camera image and linear interpolated quadcopter temperature data.

Fig. 8: Visual comparision of corrected truecolor image, thermal image, and interpolated quadcopter data.

4.2 Upsampling

As described in Section 3, we use GPR on thermal camera data augmented with truecolor data to predict temperature at unobserved locations where

truecolor data are available. The result of the upsampling analysis is shown in Fig. 9.

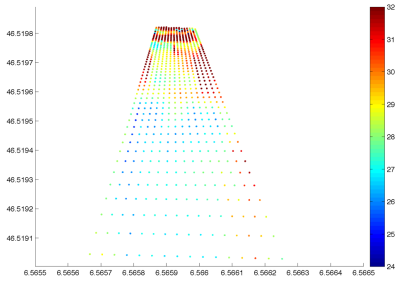

(a) Sparse thermal camera data.

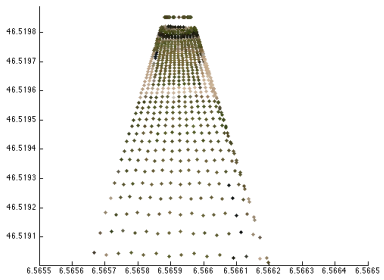

(b) Sparse truecolor data corresponding to thermal camera data.

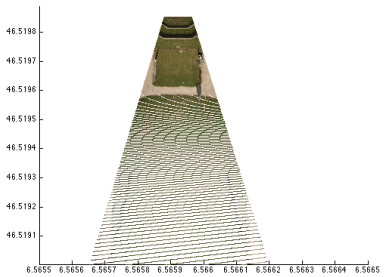

(c) Dense prediction points, with truecolor data.

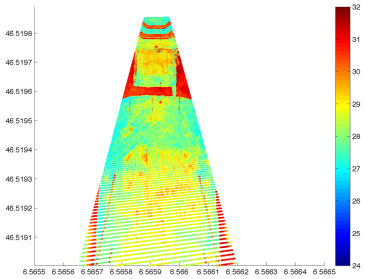

(d) Estimated thermal field.

Fig. 9: Upsampling thermal camera data with high-resolution truecolor data using Gaussian process regression with input augmented with RGB data.

4.3 Comparision of quadcopter and thermal camera data

We computed the Pearson correlation coefficient between the quadcopter data points and the perspective corrected data points from the thermal imaging camera. Since the thermal camera data density is much higher than the quadcopter data, we found the Euclidean nearest-neighbor thermal data points to the quadcopter data. Fig. 10 shows the quadcopter data alongside nearest-neighbor thermal camera data. The two vectors were of length 3607 datapoints each, and showed $R = 0.477$, demonstrating a statistically significant correlation between quadcopter data and thermal camera data. Additionally, to analyze the effect of outliers in the quadcopter data (for example, due to unfavorable altitude), we computed the Pearson coefficient on a sliding win-

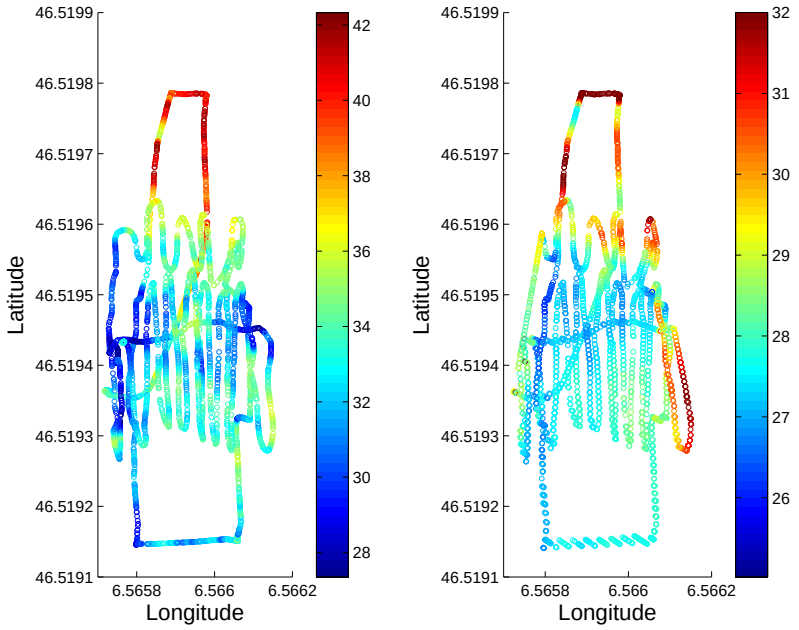

Fig. 10: Quadcopter data points (left) and nearest-neighbor thermal camera data points (right).

dow of 440 data points (corresponding to 90 second of quadcopter flight time). The resulting distribution of correlation coefficient is shown in Fig. 11a.

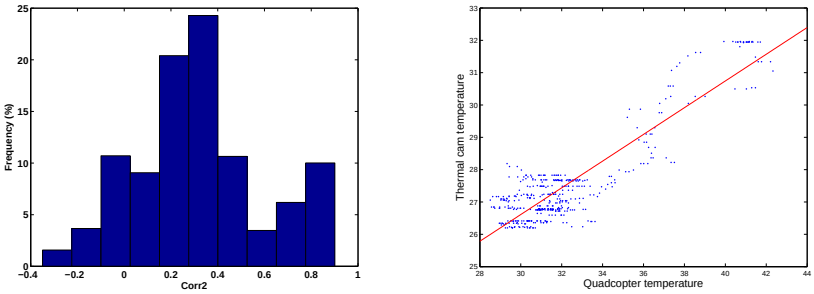

(a) Pearson correlation coefficient distribution for a sliding window of 440 datapoints over quadcopter and nearest-neighbor thermal camera data.

(b) Cross-calibration fit between quadcopter and thermal camera data.

Fig. 11: Correlation between quadcopter and thermal camera data, and plot of linear cross-calibration fit.

Next, we cross-calibrate the quadcopter data with the thermal camera data by finding a linear fit to a section of the quadcopter and thermal camera data with R>0.8. Fig. 11b shows the the data points and the resulting linear fit.

4.4 Adaptive sampling with quadcopter

We flew the quadcopter remotely to carry out lawnmower surveys of the test area capturing samples of temperature with the downward looking IR temperature sensor. From this, we can choose points to emulate adaptive collection of data. This approach allows us to try various techniques adaptively without having to perform multiple experiments. We use a subsampled version of the thermal camera image as the pilot data to learn a GPR-based probabilistic regression model of the temperature field. We then use the variance of the field to greedily choose new sample points. As a reference, we show the reconstructed field from the quadcopter data in Fig. 12a. Fig. 12b shows the

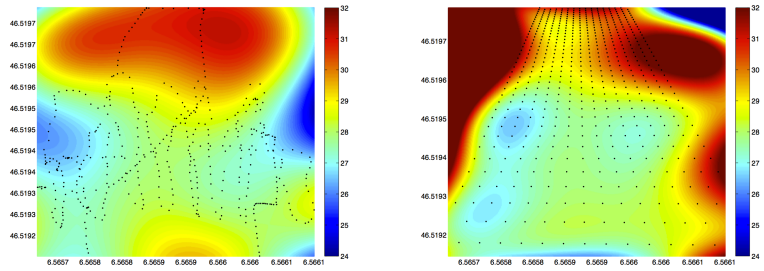

(a) GPR reconstruction on a regular grid for quadcopter data. Black dots show training data locations.

(b) GPR reconstruction on a regular grid for thermal data. Black dots show training data locations.

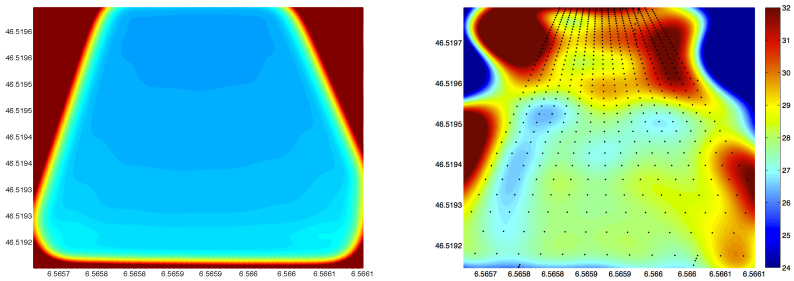

(c) Prediction variance for reconstructed thermal image. Red regions show high uncertainty.

(d) GPR reconstruction on a regular grid for thermal data after 20 adaptive samples from quadrocopter dataset.

Fig. 12: GPR reconstructed thermal and quadcopter data.

initial reconstruction of the temperature field from the thermal camera data. The top left and right corner of the reconstructed field are regions that exhibit extrapolation, with high associated uncertainty, as showed in Fig. 12c. Data from the quadcopter is used to fill these gaps, and in Fig. 12d, we see twenty additional data points added greedily to the reconstruction from the cross-calibrated quadcopter dataset. Each addition of a data point is followed

by relearning of the temperature field. As seen in this figure, as a result of addition of the new samples, the top left corner of the field now exhibits moderate temperature.

5 Discussion and Conclusions

In this paper, we presented an experiment using mixed sensing to reconstruct the surface temperature of a patch of land. A land-based camera-rig consisting of a thermographic camera and a high-resolution truecolor camera was used to generate upsampled thermal images of the experiment site. A quadcopter equipped with a downward looking IR temperature sensor measured surface temperature during flight. We compared the data from quadcopter with the thermal camera and found they are correlated (Pearson coefficient of 0.47). Finally, we investigated adaptive sampling strategies to fill gaps in thermal camera data using the quadcopter. We achieved this by using data from a quadcopter run offline.

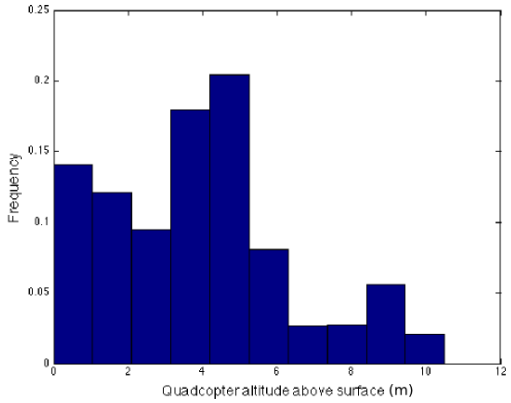

Fig. 13: The distribution of quadcopter altitude during the field trial.

This work has limitations that merit future work. First, we have not included the quadcopter altitude in the estimation of the temperature field. Since the IR temperature sensor has a relatively large field of view, the altitude has an impact on the measured data. This likely this has an impact on the correlation between the quadcopter data gathered during the field trial and the corresponding thermal camera data. Fig. 13 shows the distribution of altitude for the 10 minute flight at the experiment site. In the future, we plan to model the IR-based temperature sensor as a pixel-average sensor and take into account the effect of height while reconstructing the temperature field. Second, we have used RGB data from the truecolor camera to augment the thermal camera image to perform upsampling. We plan to explore other characteristics of the land patch in addition to RGB for this task. Finally, we have not carried out online adaptive sampling experiments with the quadcopter. Instead, we collected data using a lawnmower survey and used

the data offline to emulate new samples. In future, we will carry out online experiments to validate our approach.

Acknowledgments

This work was supported in part by the National Science Foundation under award CCF-0120778 and IIS-1107011. We thank the organizers of the Twenty-second International Joint Conference on Artificial Intelligence (IJCAI) Doctoral Consortium, and the Swiss National Center of Competence in Research in Robotics (NCCR Robotics) for facilitating the extended research visit by the first author to EPFL that made this work possible. William C. Evans and Alexander Bahr were partially supported by "The Swiss Experiment" of the Competence Center Environment and Sustainability of the ETH Domain (CCES), and by the NCCR Transfer project "Tamperproof Monitoring Solution for Weather Risk Management" sponsored by the Swiss National Science Foundation and managed by the National Center of Competence in Research in Mobile Information and Communication Systems (NCCR-MICS). We thank Holly Oldroyd and Daniel Nadeau of the Environmental Fluid Mechanics Laboratory at EPFL for lending us the thermographic camera and helping us with the science motivation.

References

1. D. Nadeau and M. Parlange, "Atmospheric Boundary Layer Dynamics of Transitional Flows over Complex Terrain," Ph.D. dissertation, EPFL, 2011.
2. M. Bryson, A. Reid, C. Hung, F. Ramos, and S. Sukkarieh, "Cost-Effective Mapping using Unmanned Aerial Vehicles in Ecology Monitoring Applications," in *12th International Symposium on Experimental Robotics, 2010*, Dec 2010.
3. A. Krause, A. Singh, and C. Guestrin, "Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies," *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, 2008.
4. B. Zhang and G. S. Sukhatme, "Adaptive Sampling for Estimating a Scalar Field using a Robotic Boat and a Sensor Network," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 3673 – 3680.
5. J. Binney, A. Krause, and G. S. Sukhatme, "Informative Path Planning for an Autonomous Underwater Vehicle," in *IEEE International Conference on Robotics and Automation*, 2010, pp. 4791 – 4796.
6. C. Thomas, T. Ranchin, L. Wald, and J. Chanussot, "Synthesis of Multispectral Images to High Spatial Resolution: A Critical Review of Fusion Methods Based on Remote Sensing Physics," *IEEE T. Geoscience and Remote Sensing*, vol. 46, no. 5, pp. 1301–1312, 2008.
7. C. E. Rasmussen, "Gaussian Processes for Machine Learning." MIT Press, 2006.