

Bayesian optimisation for decision making in spatial-temporal processes

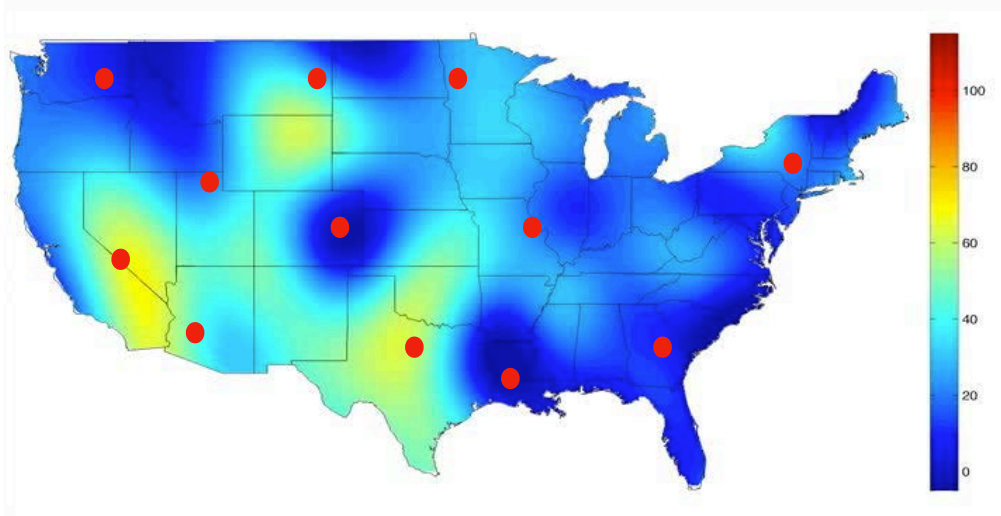
Fabio Ramos

fabio.ramos@sydney.edu.au

Joint work with Roman Marchant, Gilad Francis, Vitor Guizilini, Jefferson Souza, Lionel Ott, Rafael Oliveira, Philippe Morere, Scott Sanner, Amarjeet Singh, Sahill Garg



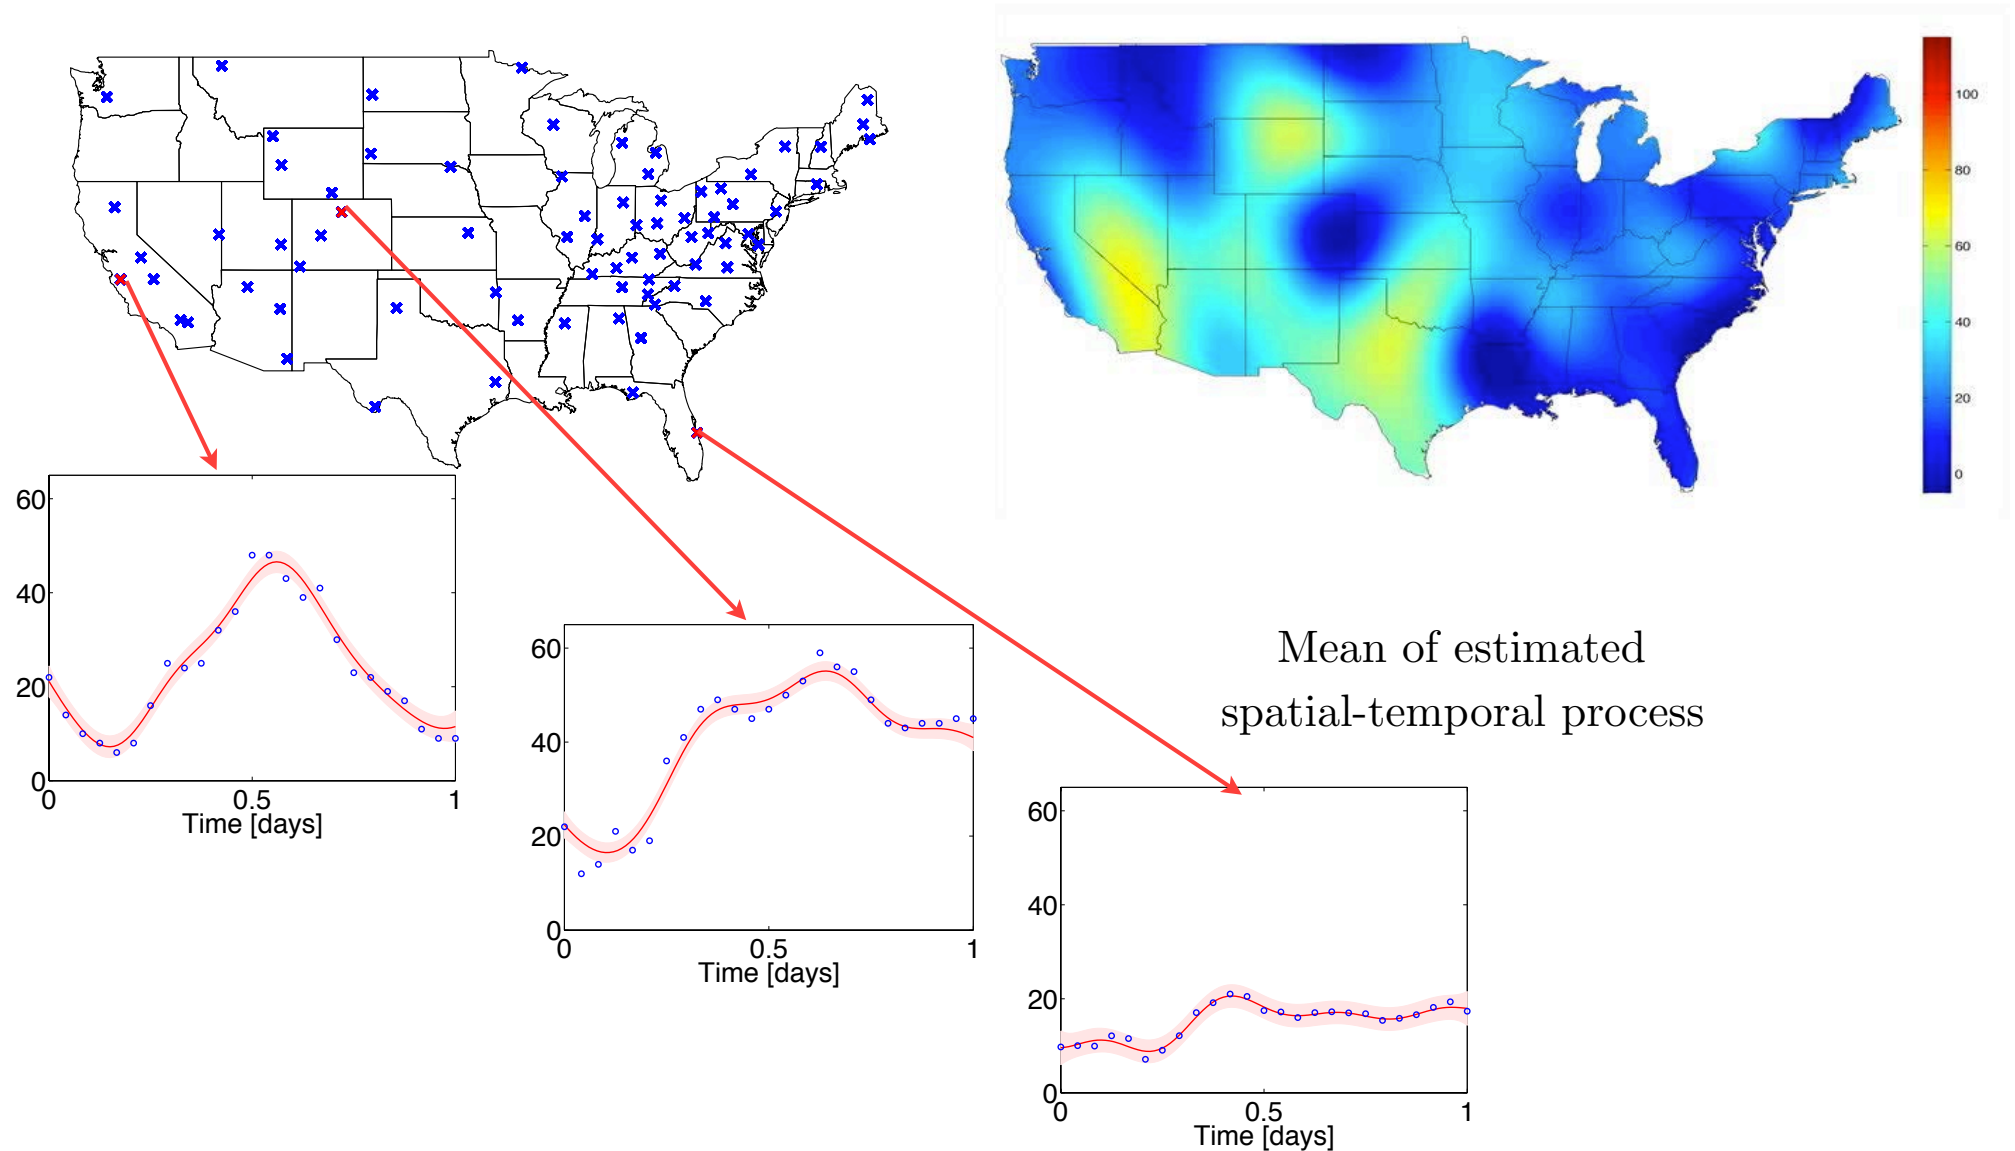
Joint monitoring and sampling



- How to interpolate?
- How to estimate predictive uncertainty?
- Where to place the sensors?
- Is the phenomenon dynamic?
- Can the sensors move? If so what is the best path?
- Are there constraints?

- Sensing radius: Hochbaum and Maas, 1985; Gonzalez Banos and Latombe, 2001
- Entropy based: Cressie, 1991; Shewry and Wynn, 1987
- Mutual Information: Caselton and Zidek, 1984; Krause et al, 2008

Dynamics of pollution monitoring



Challenges

- In many problems we would like to explore areas of high intensity (pollution)
- Complex space and time relationships introduce problems for information gain approaches (non-stationarity, non-separability)
- The need to model the dynamics is as important as fast planning
- Both modelling and planning need to be executed simultaneously

Modelling

Gaussian process regression

Given a set of samples

$$X = [\mathbf{x}_0, \dots, \mathbf{x}_{N-1}]^T$$

$$\mathbf{y} = [y_0, \dots, y_{N-1}]^T$$

Choose a covariance function



$$k(\mathbf{x}_i, \mathbf{x}_j \mid \theta)$$

Predict the value of $f(\mathbf{x}_*)$

Mean $\mu(f(\mathbf{x}_*))$

Variance $\sigma(f(\mathbf{x}_*))$


Regression



Log Marginal Likelihood

Priors over functions





$$\text{posterior} \propto \text{prior} \times \text{likelihood}$$



- A prior over the weights induces a prior over functions
 - e.g. smooth functions
 - Closeness in input space \rightarrow closeness in output space




Hilbert Spaces




Separable and non-separable covariance functions

[Singh, Ramos, ICRA 2010]



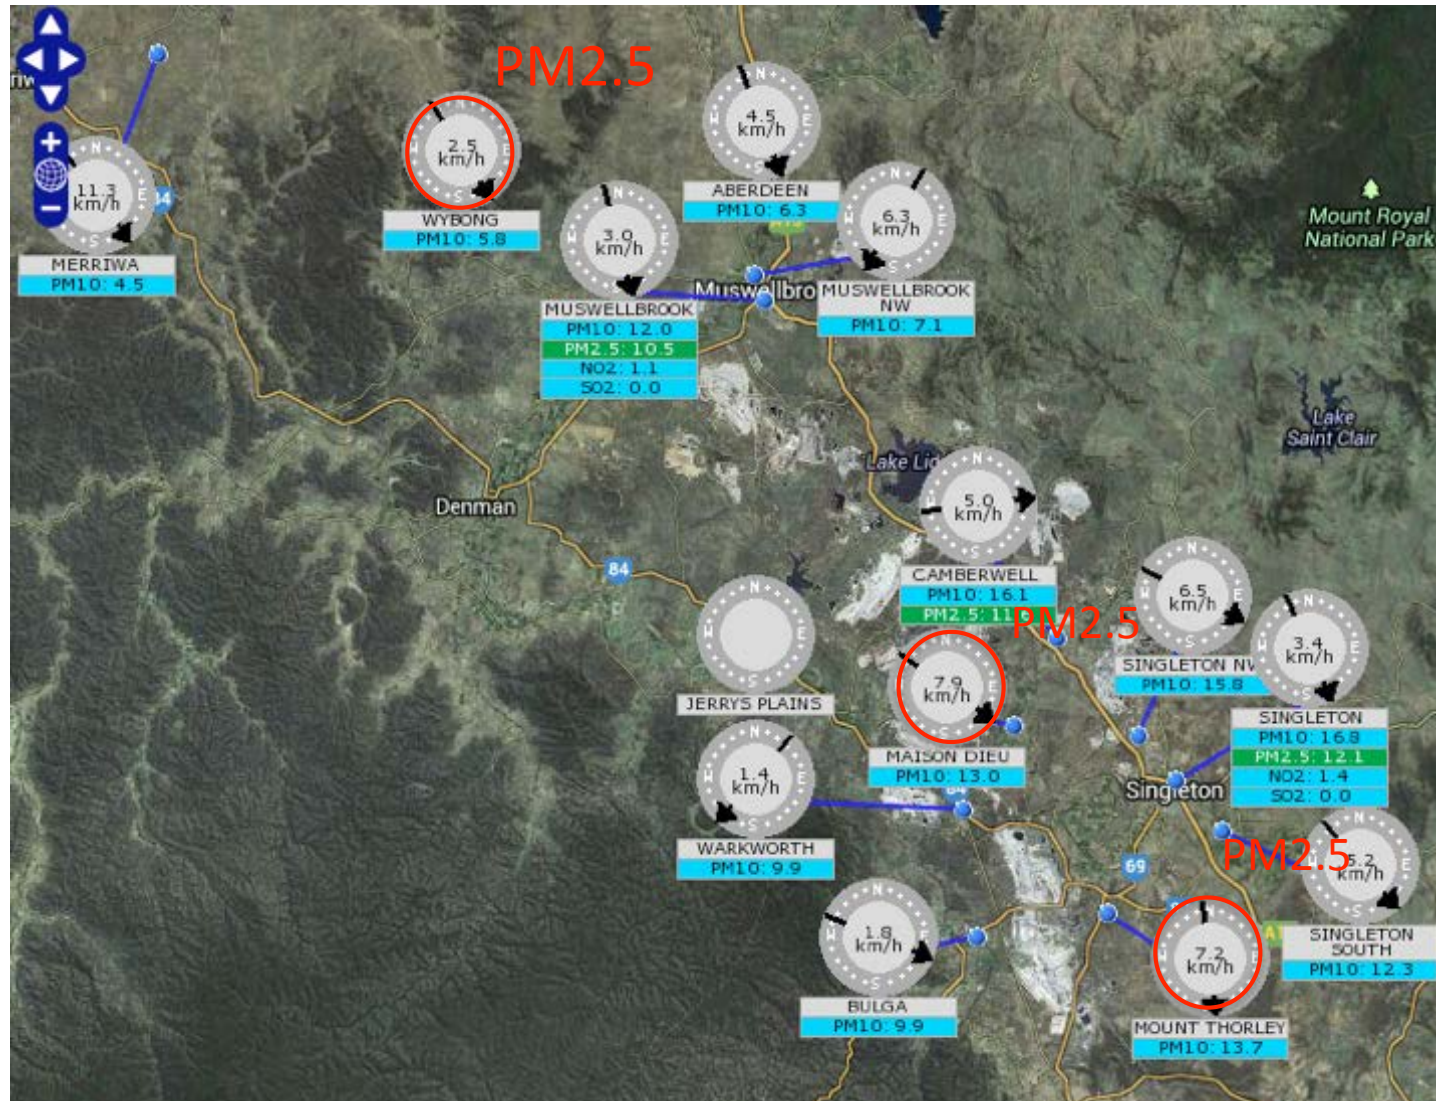
Non-separable, non-stationary



Separable, non-stationary




Air Pollution: Hunter Valley



Adding frequencies in the kernel

[Guizilini, Ramos, AAAI 2015]




Covariance function


$$k_{MT} = \theta_v^2 \left(1 + \frac{\sqrt{3}r}{\theta_l^2} \right) \exp \left(-\frac{\sqrt{3}r}{\theta_l^2} \right),$$

$$k_{PR} = \theta_v^2 \exp \left(\frac{-2 \sin^2 \left(\frac{r}{\theta_p} \right)}{\theta_l^2} \right).$$


$$k = \left(k_{MT}^{xy} + k_{MT}^{xy} \right) \times \left(k_{MT}^t + \sum_{p=1}^P k_{MT}^t k_{PR}^{t(p)} \right).$$




(a) Period of 12 hours




(b) Period of 123 hours




(c) Period of 79 hours




(d) Period of 45 hours




(e) Period of 29 hours



(f) Period of 24 hours

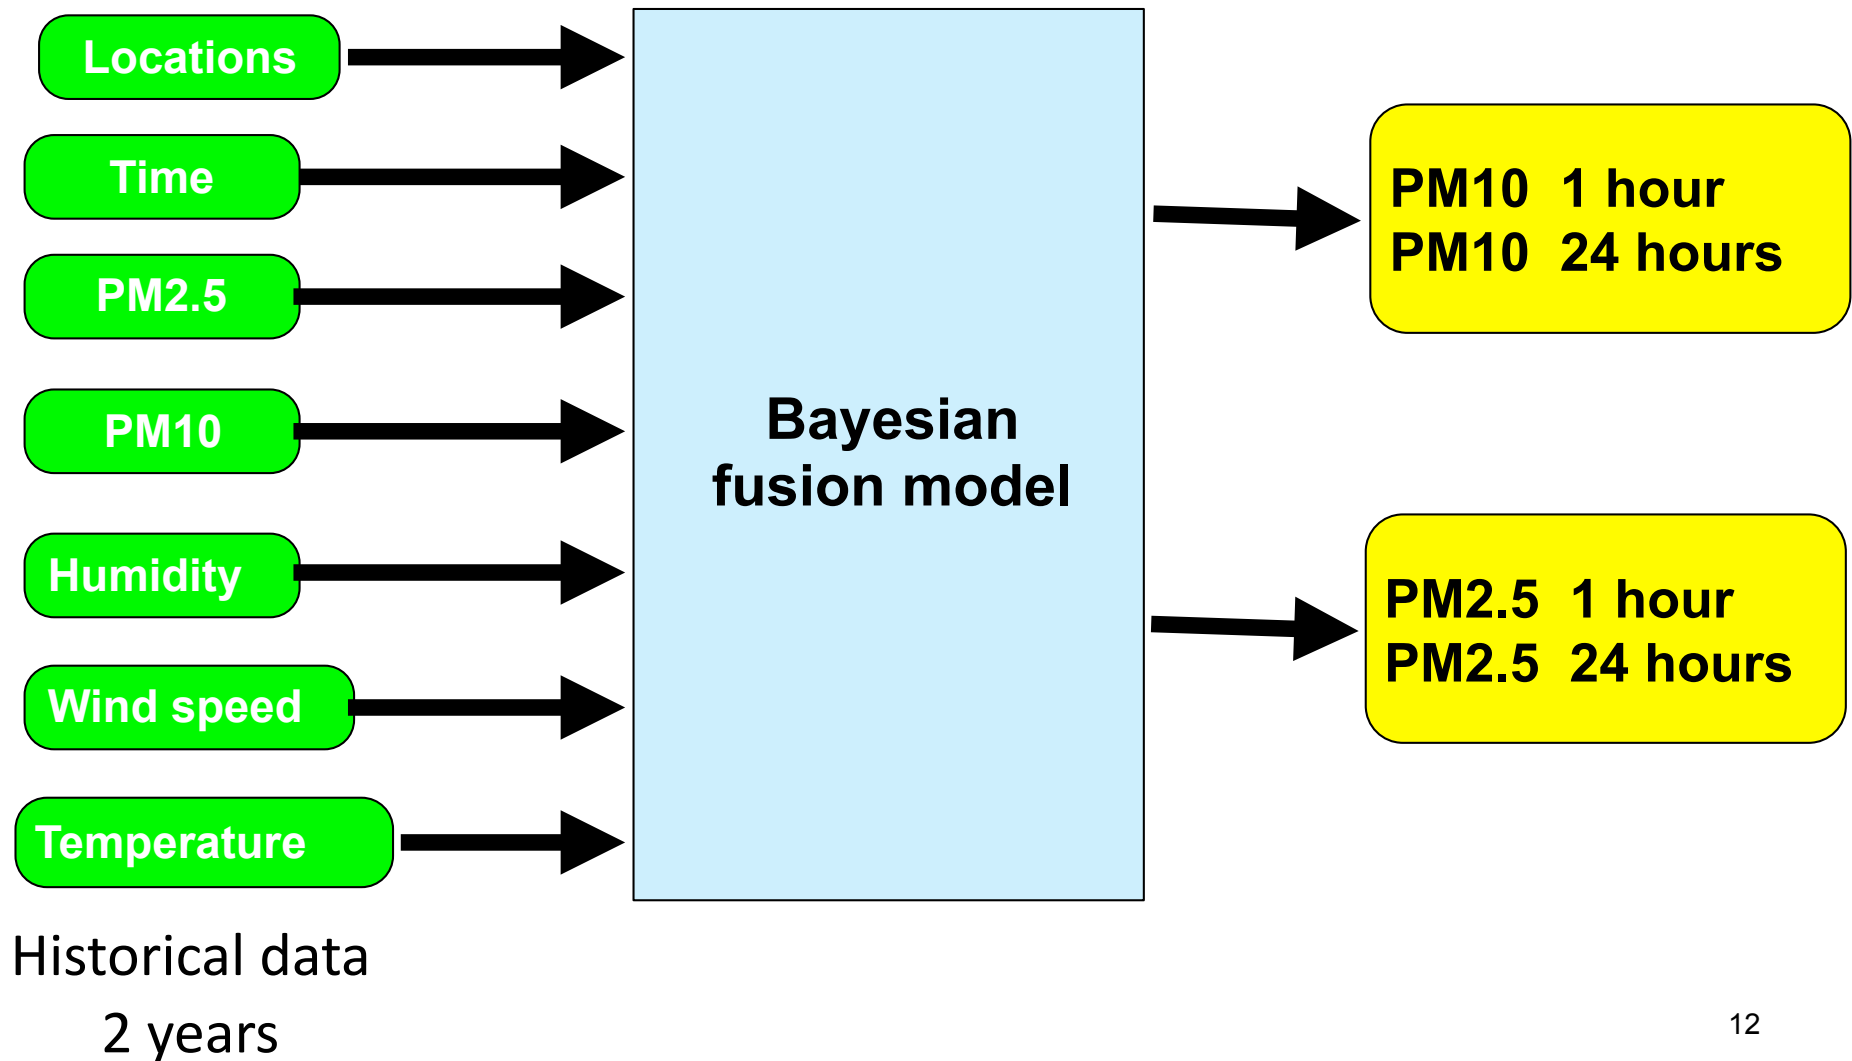


(g) Period of 52 hours



(h) Period of 17 hours

Fusion of meteorological data

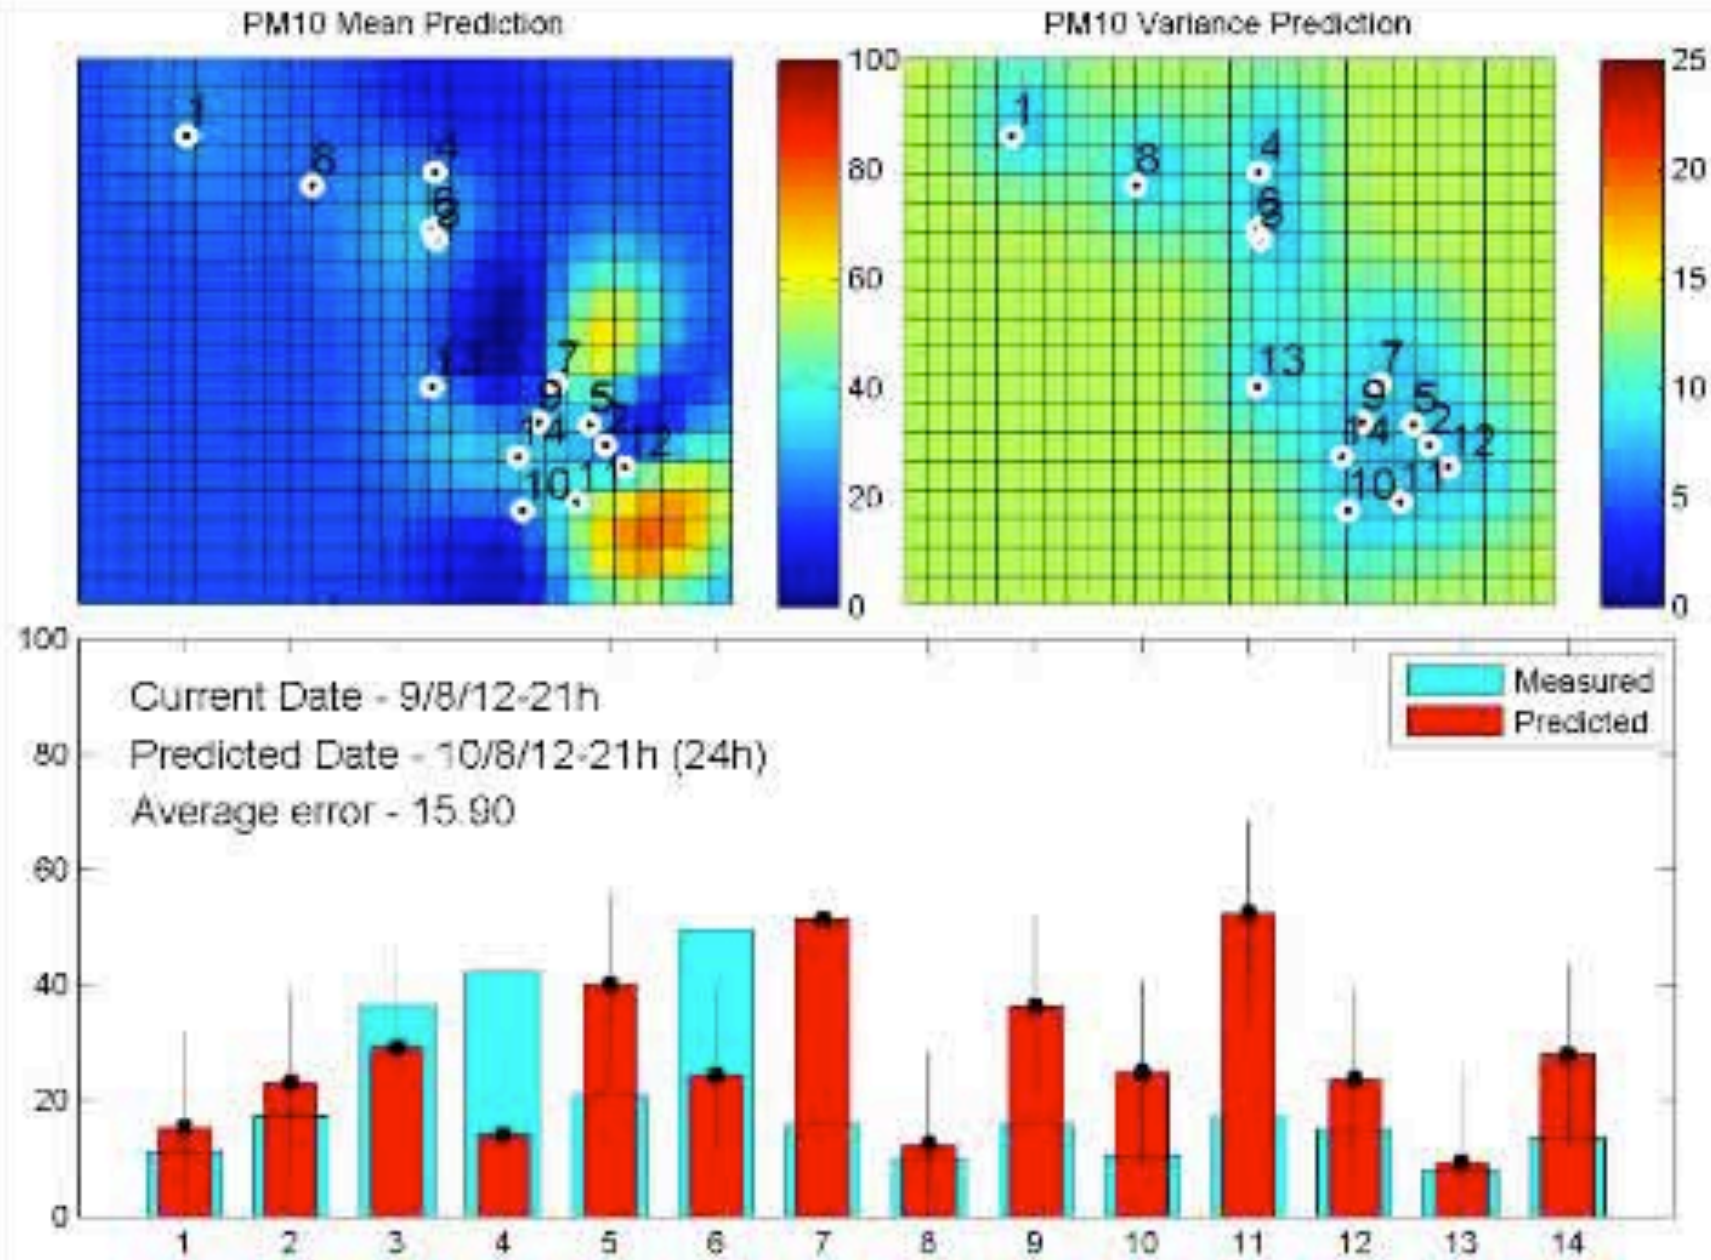


Future work: Fuse satellite images, traffic, street maps



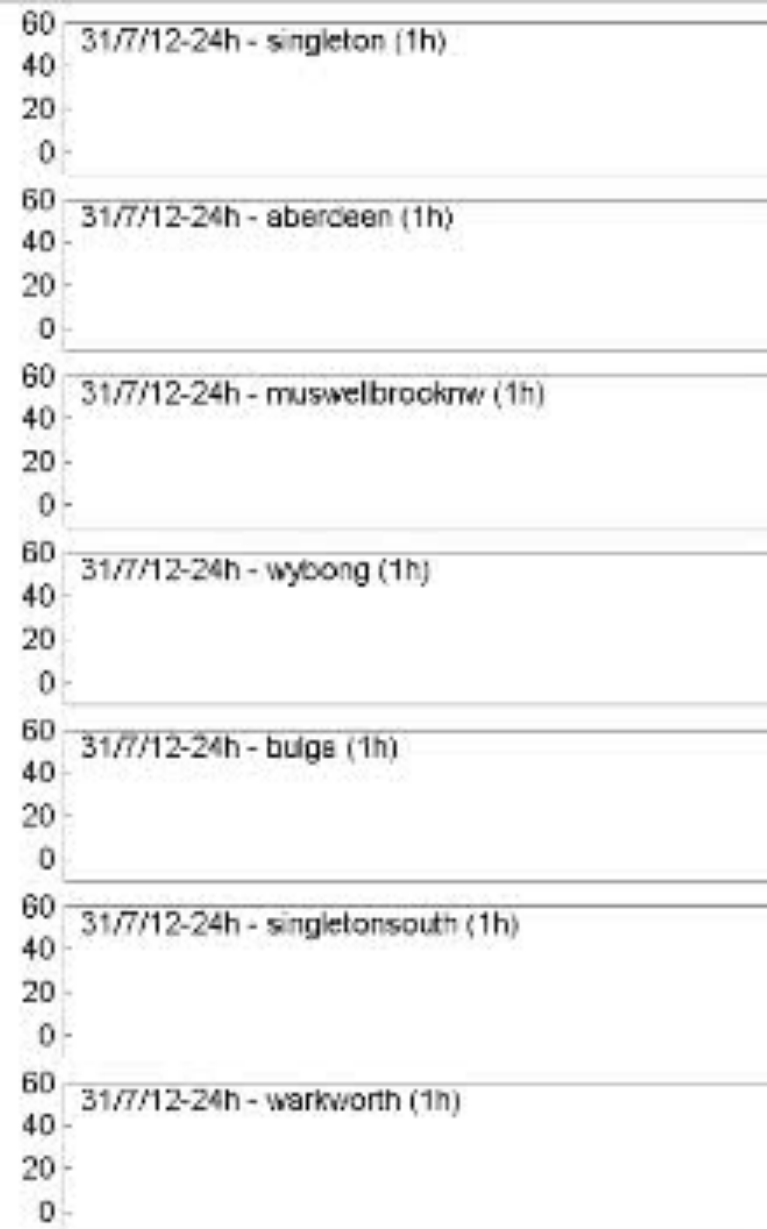
Predicting Spatio-Temporal Air Pollution

Guizilini, Ramos, AAAI 2015

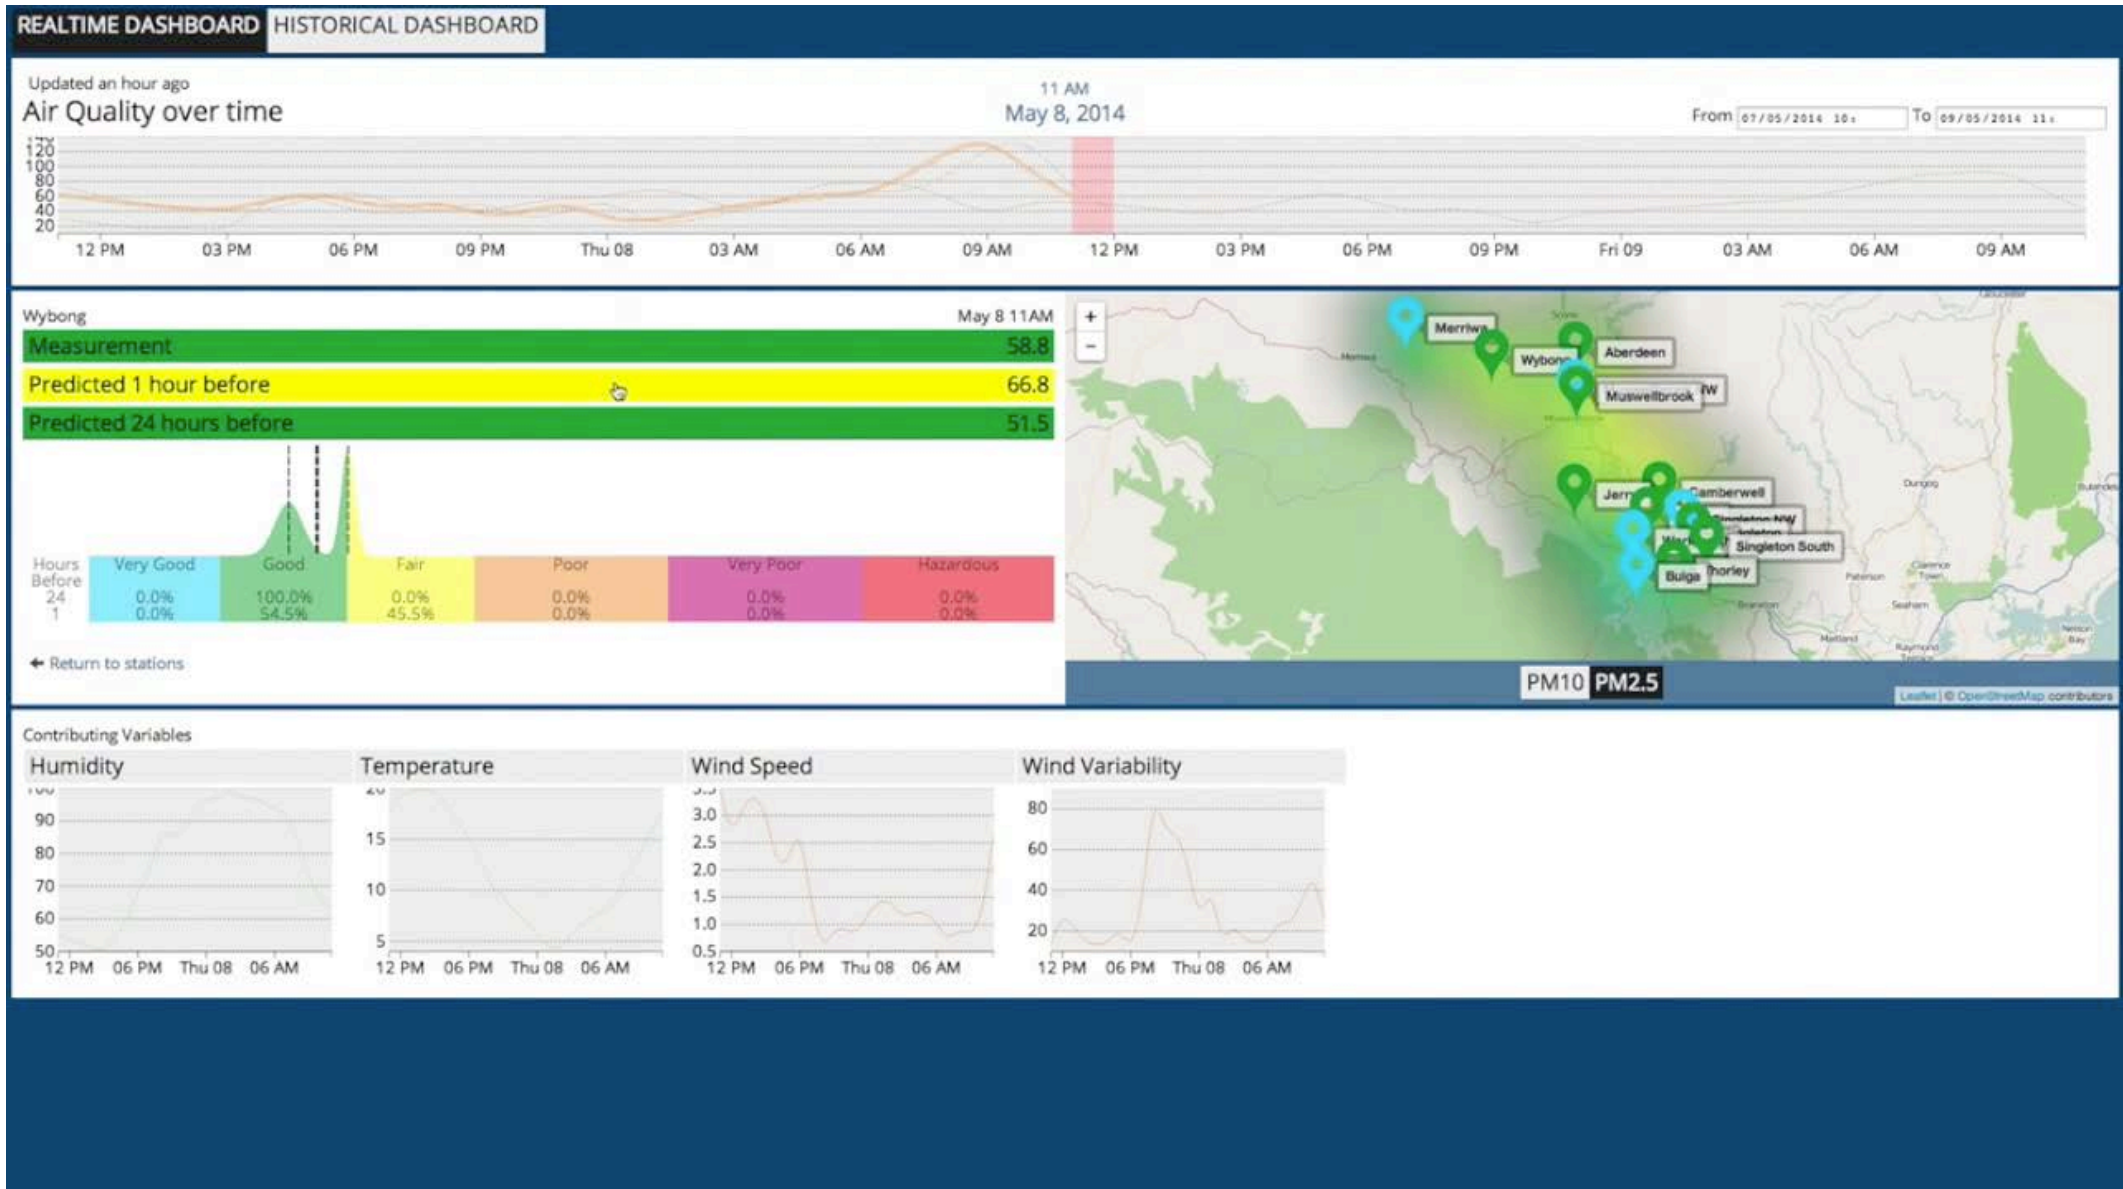




Forecasting Air Pollution



Web Interface



Predicting Spatio–Temporal Propagation of Seasonal Influenza


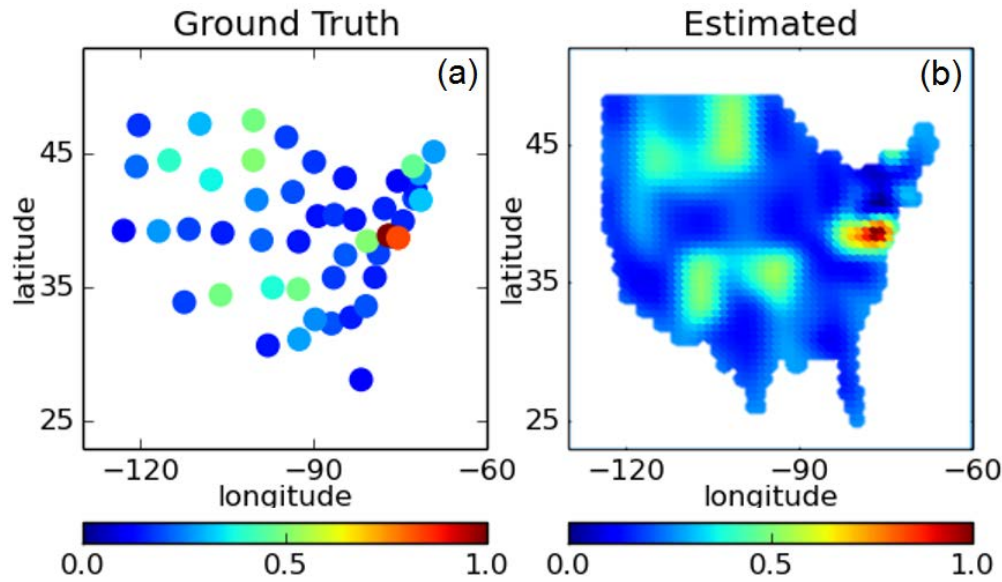


Table 1: Experiment 5. Mean squared error (MSE) averaged over states for spatio-temporal prediction of 150 weeks.

Method	1 week	2 weeks	1 year
Our approach	0.0043	0.0128	0.2900
GP - subset of data (SoD)	0.0156	0.0157	0.8642
k-NN regression	0.0068	0.0105	0.3586
LS polynomial regression	0.0169	0.0211	1.6557

[Senanayake, O’Callaghan, Ramos, AAAI 2016]

$$k_{sin}(t_i, t_j; \theta) = \sigma_{sin}^2 \exp \left(- \frac{2 \sin^2(\pi f \Delta_t)}{l_{sin}^2} \right)$$

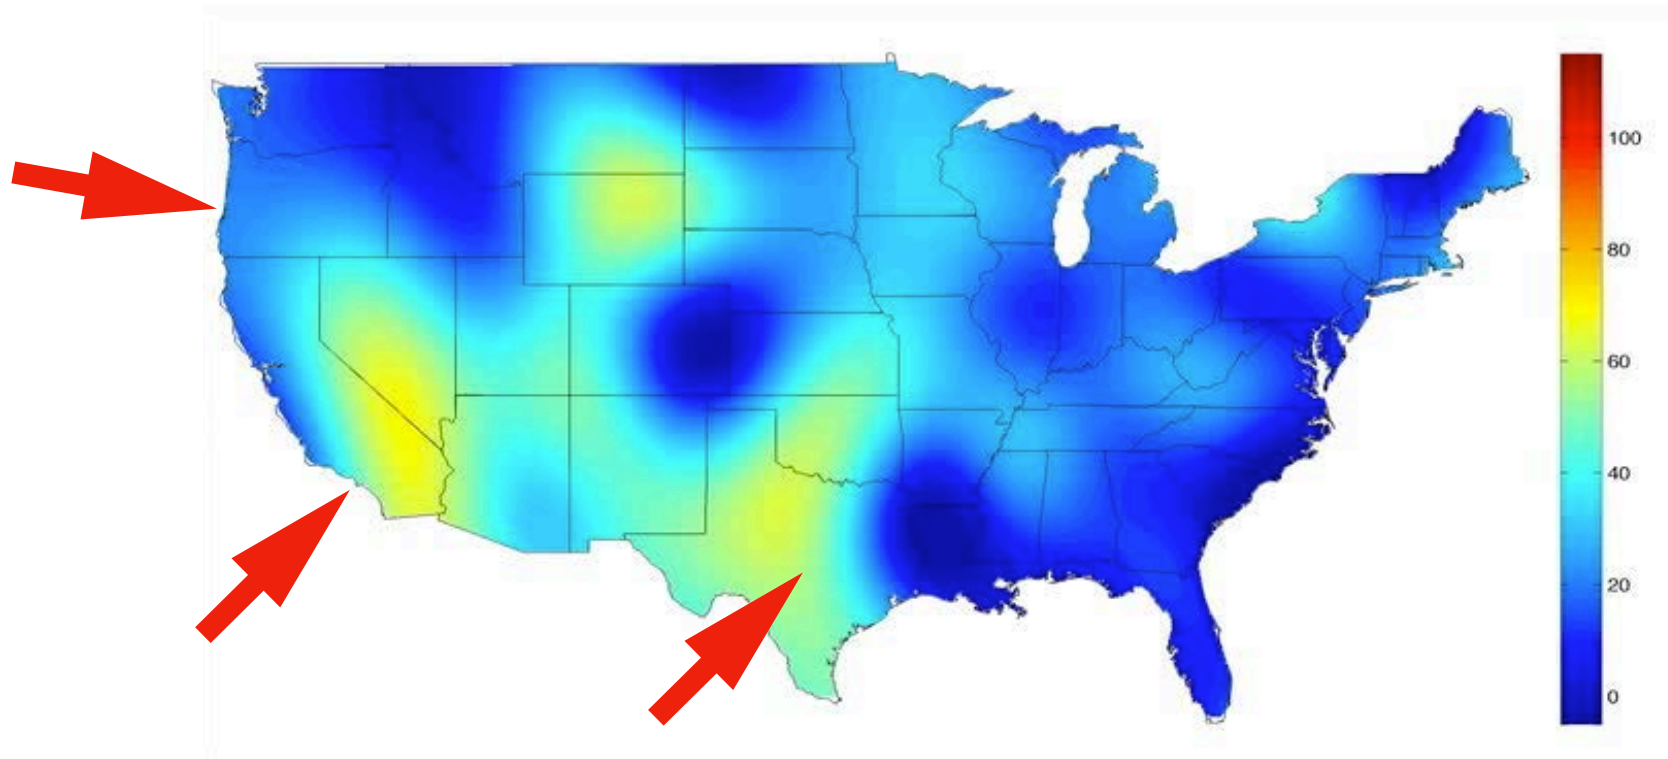
$$k_{pac}(\mathbf{z}_i, \mathbf{z}_j) = \sigma_{pac}^2 |\Sigma_i|^{\frac{1}{4}} |\Sigma_j|^{\frac{1}{4}} \left| \frac{\Sigma_i + \Sigma_j}{2} \right|^{-\frac{1}{2}}$$

$$\exp \left(- (\mathbf{z}_i - \mathbf{z}_j)^\top \left(\frac{\Sigma_i + \Sigma_j}{2} \right)^{-1} (\mathbf{z}_i - \mathbf{z}_j) \right)$$

Planning




High values are more important



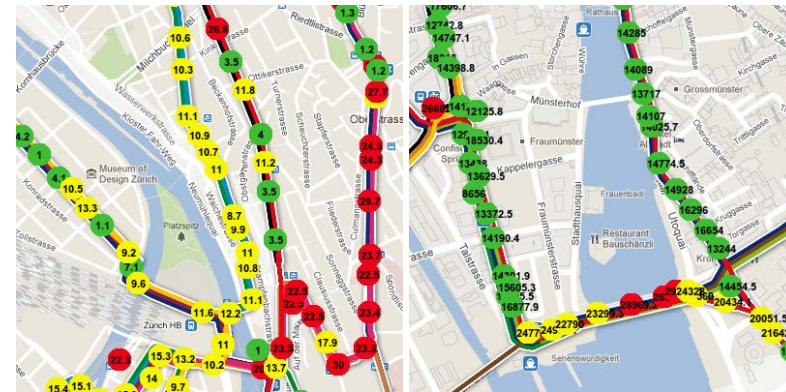
Mobile Monitoring: Where and When to Monitor

OpenSense Zurich Monitoring System



(a) Sensor box

(b) DiSCmini



(a) Ozone

(b) PM

Figure 4: Web-based interactive data browsing through the datasets of ozone and PM measurements in Zurich.



(a) Tram deployment



(b) Fixed station deployment



UAVs



USVs

Bayesian optimisation

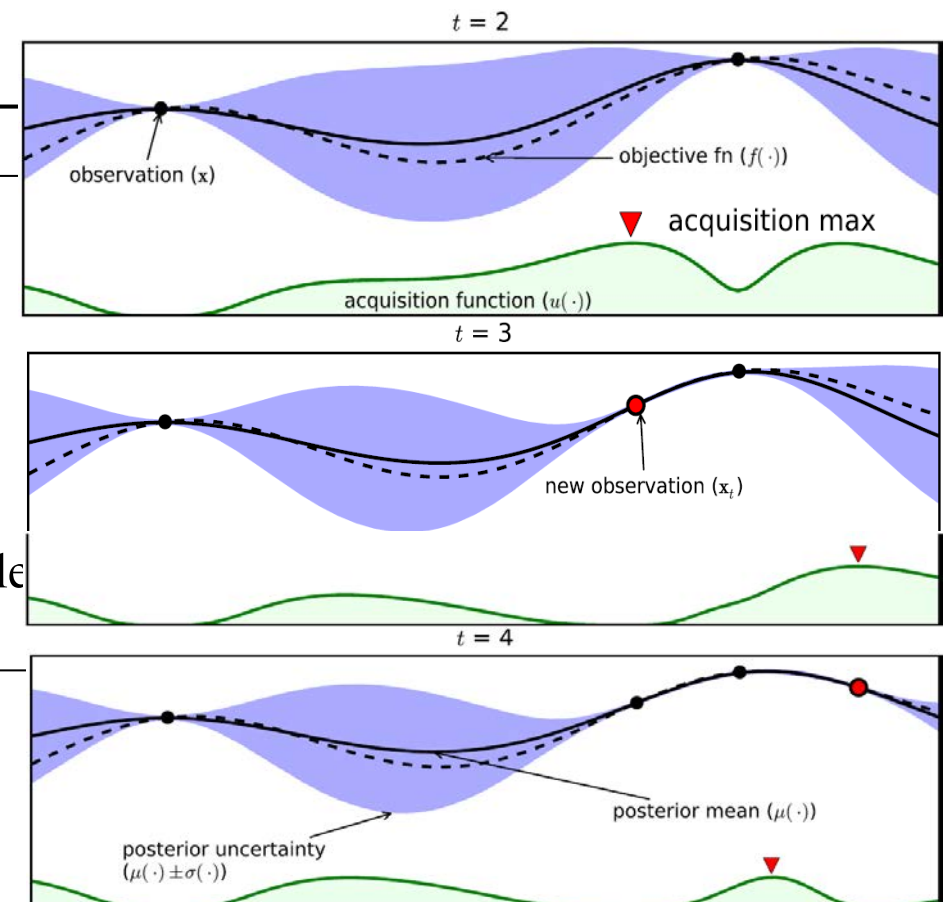
- Finds the maximum of an unknown objective function difficult to evaluate
- Learns the objective function by obtaining noisy samples from it
- Uses the *Gaussian process prior* to estimate the objective function
- Applies an *acquisition function* trading off exploration and exploitation

Decisions as Bayesian optimisation

Find the maximum of an unknown, noisy and costly to evaluate function f .
Idea: Choose next sampling location \mathbf{x} by maximising an acquisition function s over the domain of the GP model of the function.

Algorithm 1 Bayesian Optimisation

- 1: \mathbf{x}_i : chosen sampling point at iteration i .
- 2: s : acquisition function.
- 3: f : unknown function. (terrain roughness)
- 4: **for** $i = 1, 2, 3, \dots$ **do**
- 5: \rightarrow Find $\mathbf{x}_i = \arg \max_{\mathbf{x}} s(\mathbf{x})$
- 6: \rightarrow Acquire a sample from f at location \mathbf{x}_i .
- 7: \rightarrow Update the GP model of f with the new sample
- 8: **end for**



[Brochu, Cora, Freitas 2010]

Acquisition functions

Probability of Improvement

$$PI(\mathbf{x}) = \Phi \left(\frac{\mu(\mathbf{x}) - f(\mathbf{x}^+) - \xi}{\sigma(\mathbf{x})} \right)$$

$$DPI(\mathbf{x}|\mathbf{x}^-) = \Phi \left(\frac{\mu(\mathbf{x}) - f(\mathbf{x}^+) - \xi}{\sigma(\mathbf{x})} \right) + \gamma \cdot d(\mathbf{x}, \mathbf{x}^-)$$

Expected Improvement

$$EI(\mathbf{x}) = \sigma(\mathbf{x}) [Z\Phi(Z) + \phi(Z)]$$

$$DEI(\mathbf{x}|\mathbf{x}^-) \triangleq \sigma(\mathbf{x}) [Z\Phi(Z) + \phi(Z)] + \gamma \cdot d(\mathbf{x}, \mathbf{x}^-)$$

$$Z = \frac{\mu(\mathbf{x}) - f(\mathbf{x}^+) - \xi}{\sigma(\mathbf{x})}$$


Upper Confidence Bound

$$UCB(\mathbf{x}) \triangleq \alpha \cdot \mu(\mathbf{x}) + \kappa \cdot \sigma(\mathbf{x})$$

$$DUCB(\mathbf{x}|\mathbf{x}^-) \triangleq \alpha \cdot \mu(\mathbf{x}) + \kappa \cdot \sigma(\mathbf{x}) + \gamma \cdot d(\mathbf{x}, \mathbf{x}^-)$$


Introduce **penalty on distance** to last sampled location \mathbf{x}^-

With $\gamma < 0$




[Brochu, Cora, Freitas 2010]

Planning as Bayesian optimisation



Ozone concentration






Mean of Estimation at $t=0.62$

[Marchant, Ramos, IROS 2012]




IG (Information Gain)

UCB

DUCB



Trajectories



RMSE 13.78
Distance 717.18

13.04
718.25

12.33
639.06



Indoor ST-Luminosity model

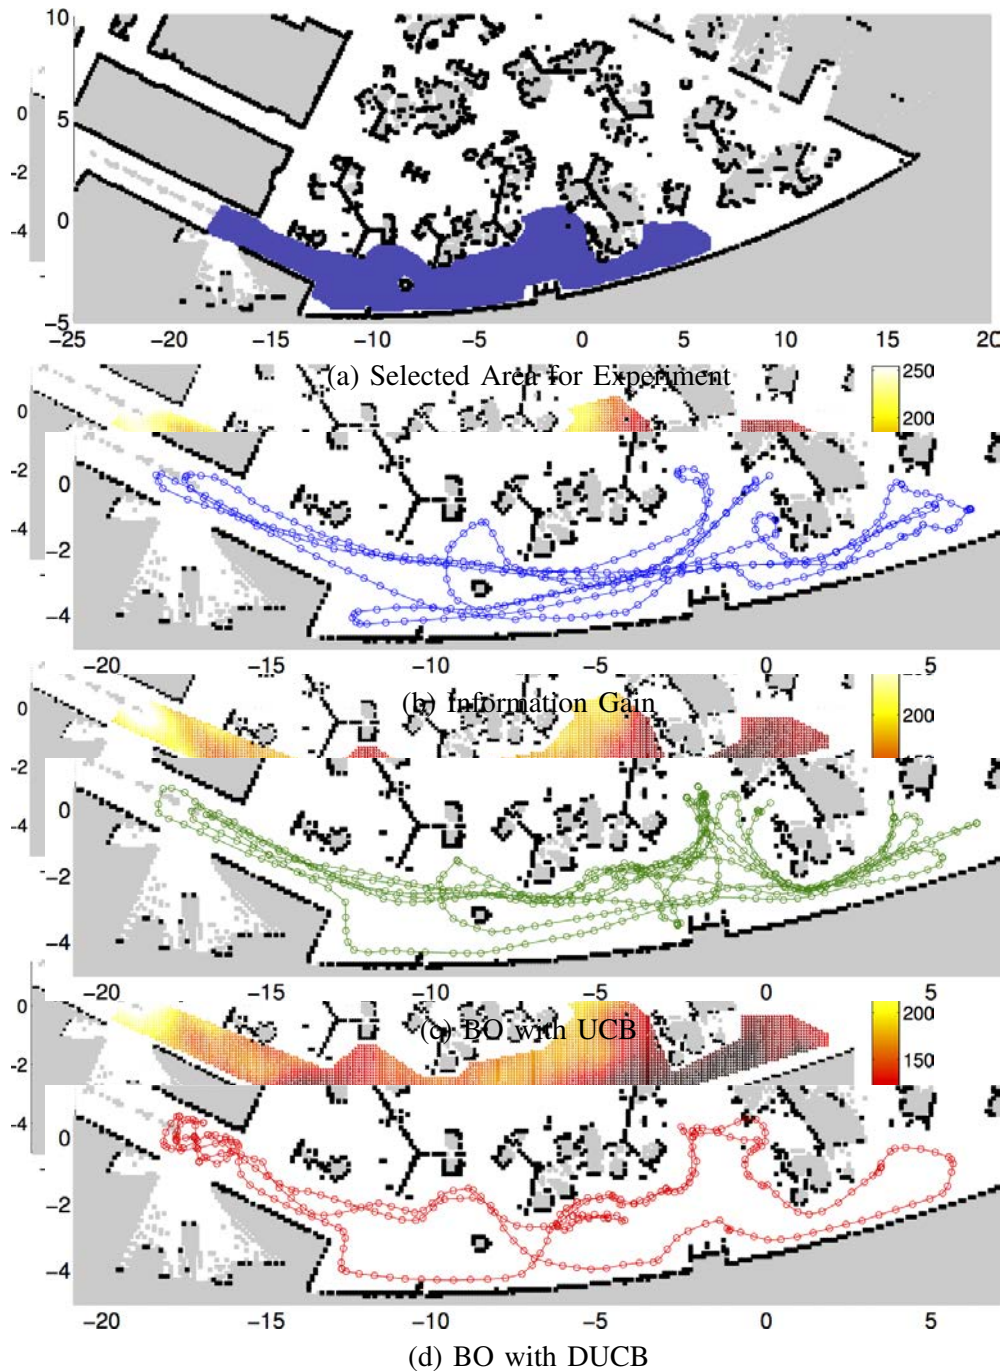



TABLE II
RESULTS FOR REAL EXPERIMENT

Indicator	Method	Mean	Distance [m]
RMSE	IG	16.44	157.01
RMSE	UCB	17.16	217.45
RMSE	DUCB	16.91	98.08
WRMSE	IG	5.41	157.01
WRMSE	UCB	4.60	217.45
WRMSE	DUCB	4.70	98.086




[Marchant, Ramos, IROS 2012]

Avoiding high vibration


[Souza, et al, ICRA 2014]



(a) Map matching



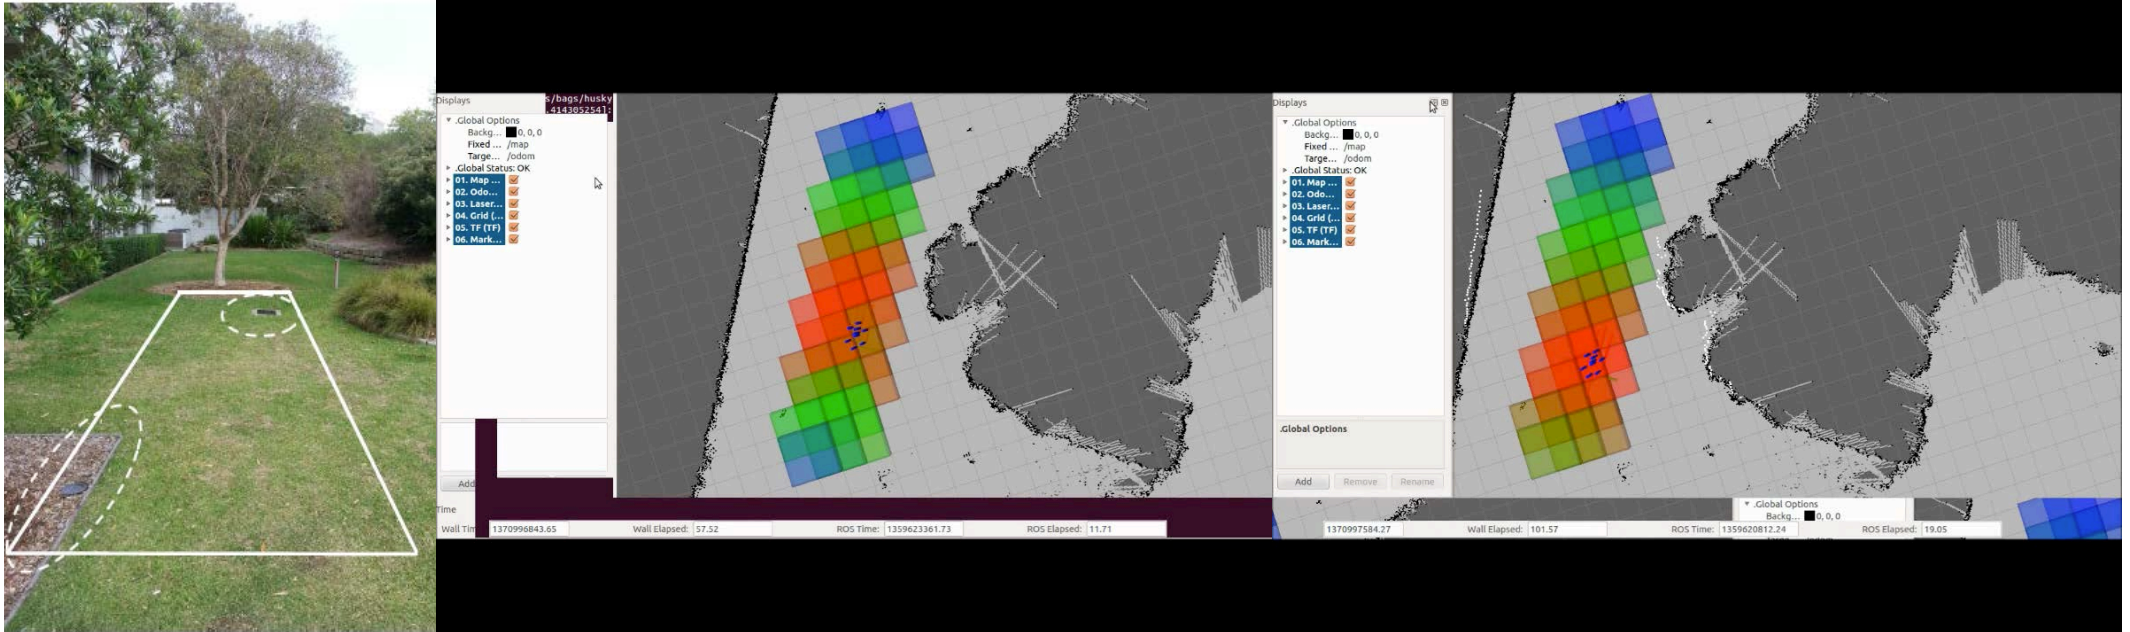
(b) Mean



(c) Variance

The mean of the vibration estimate shows a clear distinction between two explored terrains, grass and asphalt.

Grass scene




BO Planner

Information gain planner


[J. Souza, R. Marchant, L. Ott, D. Wolf, F. Ramos. Bayesian Optimisation or Active Perception and Smooth Navigation. In ICRA, 2014]

Experimental Results


(f) Asphalt Scene




(g) Vibration Map and Path BO




(h) Vibration Map and Path Entropy



(i) Position Heatmap BO



(j) Position Heatmap Entropy



(f) show experiment area as a rectangle and obstacles that generate vibration are marked with circles. (g) and (h) show a surface for the prediction of vertical acceleration in $[m/s^2]$ and axis in meters. (i) and (j) show a heat map of the experiment.

Average acceleration in $[m/s^2]$ for BO and Entropy planners.

Scenario	Method	Mean	Std
Grass	BO	0.096	0.068
Grass	Entropy	0.154	0.148
Asphalt	BO	0.103	0.122
Asphalt	Entropy	0.127	0.256

The BO planner clearly results in less vibration encountered by the robot, in terms of mean and standard deviation.

BO for continuous paths

We use parametrised curves

$$\mathcal{C}_1(u) = a_x u^3 + b_x u^2 + c_x u + d_x$$

$$\mathcal{C}_2(u) = a_y u^3 + b_y u^2 + c_y u + d_y$$


$$\mathcal{C}_1(u=0) = p_{x_i} = d_x,$$

$$\mathcal{C}_2(u=0) = p_{y_i} = d_y,$$

$$\left. \frac{\partial \mathcal{C}_2 / \partial u}{\partial \mathcal{C}_1 / \partial u} \right|_{u=0} = \frac{c_y}{c_x} = \tan \alpha_i.$$

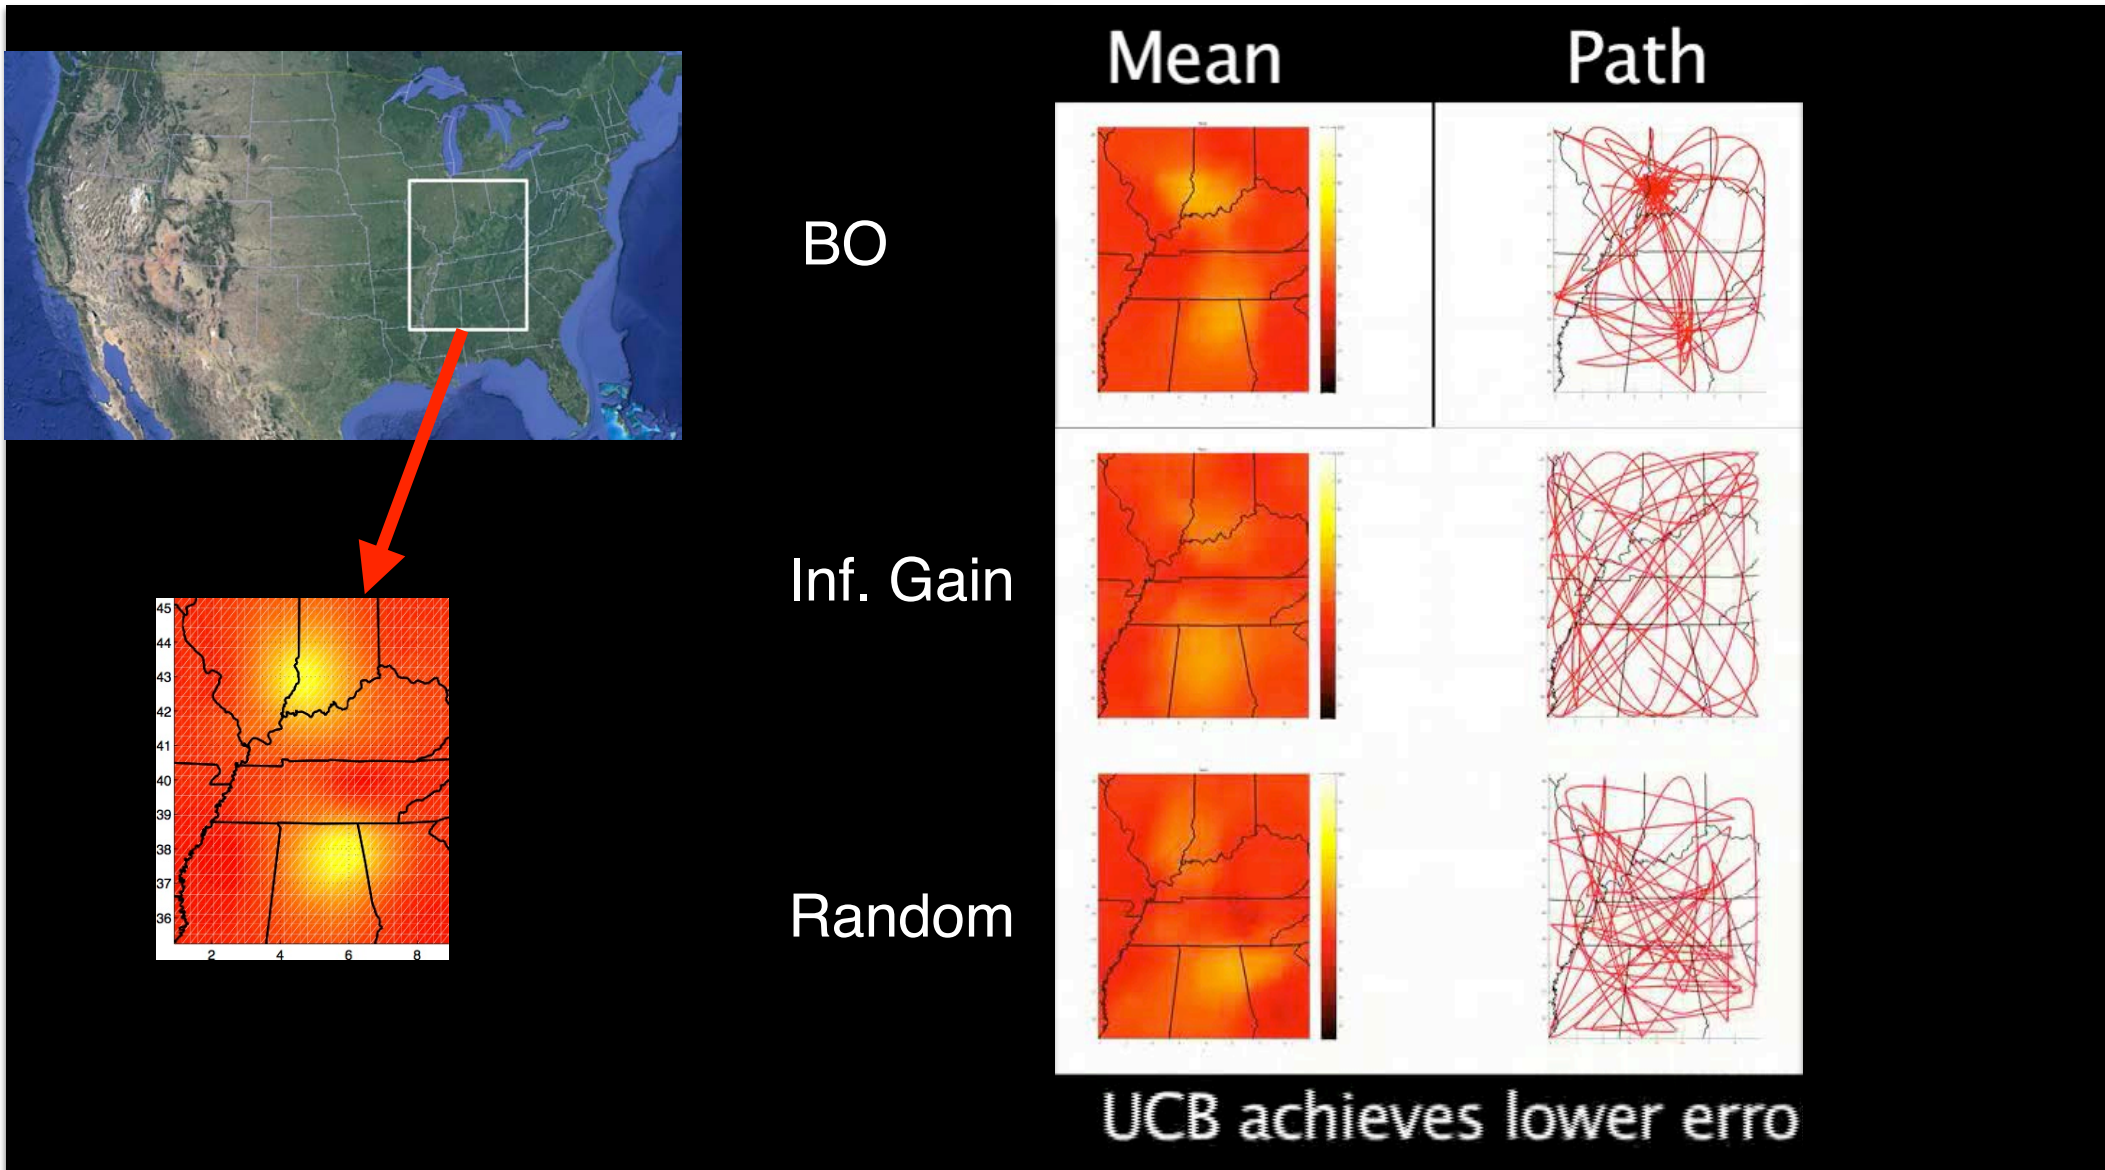
Modified cost for sampling over paths

$$\begin{aligned} r(\mathcal{C}(u, \beta) | h) &= \int_{\mathcal{C}(u, \beta)} h(v) dv \\ &= \int_0^1 \text{UCB}(\mathcal{C}(u, \beta)) \|\mathcal{C}'(u, \beta)\| du \\ &= \int_0^1 [\mu(\mathcal{C}(u, \beta)) + \kappa \sigma(\mathcal{C}(u, \beta))] \|\mathcal{C}'(u, \beta)\| du. \end{aligned}$$



[Marchant, Ramos, ICRA 2014]


Planning to improve predictions




R. Marchant, F. Ramos, Bayesian Optimisation for Informative Continuous Path Planning. In IEEE International Conference on Robotics and Automation (ICRA), 2014.

R. Marchant, F. Ramos, S. Sanner, Sequential Bayesian Optimisation for Spatial-Temporal Monitoring. In Uncertainty in Artificial Intelligence (UAI), 2014


BO Continuous vs discrete




(a) RD




(b) ED




(c) UCBD



(d) RC



(e) EC



(f) UCBC

Sequential Bayesian Optimisation (SBO)

[Marchant, Ramos, Sanner, UAI 2014]

At every step a generic reward can be obtained by sampling at \mathbf{x}^*

Myopic expectation of reward:

$$\begin{aligned} \text{ER}(\mathbf{x}^* | \mathcal{D}_0) &= \mathbb{E}_{f^*} [r(\mathbf{x}^*, f^* | \mathcal{D}_0)] \\ &= \int r(\mathbf{x}^*, f^* | \mathcal{D}_0) p(f^* | \mathbf{x}^*, \mathcal{D}_0) df^* \end{aligned} \quad [\text{Osborne, Garnett and Roberts 2009}]$$


n-step lookahead:

$$\begin{aligned} \text{ER}_n(\mathbf{x}^* | \mathcal{D}_0) &= \int r(\mathbf{x}^*, f^* | \mathcal{D}_0) p(f^* | \mathbf{x}^*, \mathcal{D}_0) \left(\sum_{i=2}^n \int r(\mathbf{x}_i, f_i | \mathcal{D}_{i-1}) \right) \\ &\quad \times \prod_{i=2}^n p(f_i | \mathbf{x}_i) p(\mathbf{x}_i | \mathcal{D}_{i-1}) \\ &\quad df_2 \cdots df_n d\mathbf{x}_2 \cdots d\mathbf{x}_n \end{aligned}$$

Wang and Sanner 2014: all possible function realisations.
Wang and Sanner 2014: all possible function realisations.


Sequential BO as POMDPs

[Marchant, Ramos, Sanner, UAI 2014]



$$\begin{aligned}
 &\text{Long-term expected reward} \quad \mathbb{E} R_n(\mathbf{x}^*, \mathcal{D}_0) \\
 &= \int \cdots \int \left(r(\mathbf{x}^*, f^* | \mathcal{D}_0) + \sum_{i=2}^n (r(\mathbf{x}_i, f_i | \mathcal{D}_{i-1})) \right) \\
 &\quad p(f^* | \mathbf{x}^*, \mathcal{D}_0) \times \prod_{i=2}^n p(f_i | \mathbf{x}_i) p(\mathbf{x}_i | \mathcal{D}_{i-1}) \\
 &\quad df^* df_1 \cdots df_n d\mathbf{x}_1 \cdots d\mathbf{x}_n,
 \end{aligned}$$

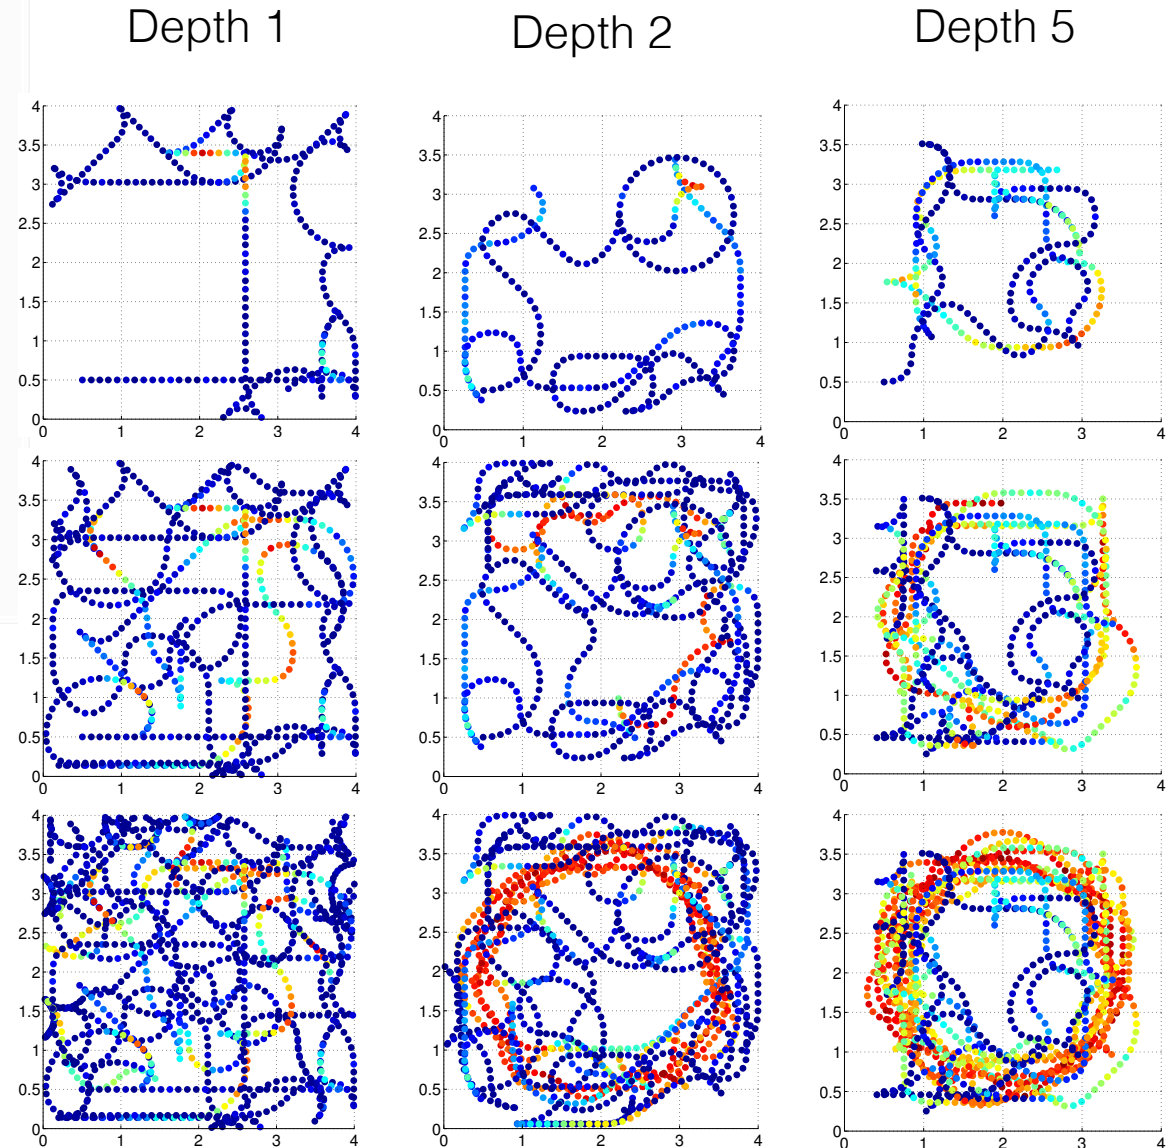
Experiments: Dynamic Function




Space Time Function
Ground Truth

$$r(\mathbf{x}) = \mu(\mathbf{x}) + 10\sigma(\mathbf{x})$$


Clear relation between lookahead
steps and ability to find maximum in
space and time.



Experiments: Dynamic Function



- MCTS prioritises search over promising part of the tree.




- Deeper search translates into greater accumulated reward.
- Full Tree search is an upper bound on MCTS.
- MCTS accumulated reward is close to Full Tree search, but is much faster.



Sequential planning with BO-POMDP

[Morere, Marchant, Ramos, ICRA 2017]



POMDP belief

POMDP reward function

New acquisition/reward function:

$$r(\Theta) = \int_0^1 \nabla \mu(\mathcal{T}(\Theta)|_t) + \kappa \sigma(\mathcal{T}(\Theta)|_t) dt$$

Trajectory parameters

Model gradient




Exploration-exploitation balance

Model uncertainty



Comparisons

[Morere, Marchant, Ramos, ICRA 2017]



Ground truth

Reconstructed

SBO [Marchant et al., 2014]

- Modelling dynamics makes the problem of sampling much more complex
- Planning for information gain or mutual information is not ideal to many environmental problems despite guarantees from sub-modularity
- Bayesian optimisation is a powerful alternative to more general planning problems
- The connections between sequential BO and POMDPs are intriguing

References

- Philippe Morere, Roman Marchant, Fabio Ramos. “Sequential Bayesian Optimisation for POMDPs and Environment Monitoring with UAVs”. In IEEE International Conference on Robotics and Automation (ICRA), 2017.
- Ransalu Senanayake, Simon O’Callaghan and Fabio Ramos. “Predicting Spatio-Temporal Propagation of Seasonal Influenza using Variational Gaussian Process Regression”. In Association for the Advancement of Artificial Intelligence (AAAI), 2016.
- Vitor Guizilini, Fabio Ramos, “A Nonparametric Online Platform for Air Quality Prediction”. In Association for the Advancements of Artificial Intelligence (AAAI), 2015.
- Roman Marchant, Fabio Ramos, Scott Sanner, “Sequential Bayesian Optimisation for Spatial-Temporal Monitoring”, In Uncertainty in Artificial Intelligence (UAI), 2014.
- Roman Marchant, Fabio Ramos, “Bayesian Optimisation for Informative Continuous Path Planning”, In IEEE International Conference on Robotics and Automation (ICRA), 2014.
- Jefferson Souza, Roman Marchant, Lionel Ott, Denis Wolf, Fabio Ramos, “Bayesian Optimisation for Active Perception and Smooth Navigation”, In IEEE International Conference on Robotics and Automation (ICRA), 2014.
- Roman Marchant, Fabio Ramos. “Bayesian Optimisation for Intelligent Environmental Monitoring”. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012.
- Sahill Garg, Amarjeet Singh, Fabio Ramos. “Learning Non-Stationary Space-Time Models for Environmental Monitoring”. In Association for the Advancements of Artificial Intelligence (AAAI), 2012.
- Amarjeet Singh, Fabio Ramos, Hugh Durrant-Whyte, and William J Kaiser. “Modeling and Decision Making in Spatio-Temporal Processes for Environmental Surveillance”. In IEEE International Conference on Robotics and Automation (ICRA), 2010.