Improving the Reliability of Crowdsourced Radiation Monitoring using Heteroscedastic Gaussian Process Regression

Matteo Venanzi, Alex Rogers, and Nick Jennings

Agents, Interaction, and Complexity Research Group – Electronics and Computer Science

University of Southampton

EPSRC Engineering and Physical Sciences Research Council

ORCHID

Introduction

Crowdsourcing information of a spatial function is becoming a key application in disaster response.

Example: 2011 earthquake in Japan
A number of nuclear power plants were severely damaged causing a nuclear post-earthquake emergency.

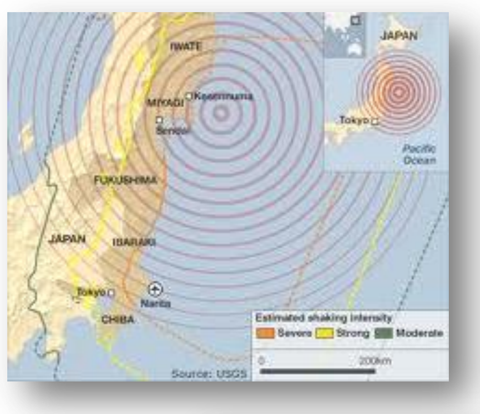

Local responders (sensors) submit geo-tagged radioactivity readings using Geiger counters connected to their mobile phones (e.g. Wikisensor iphone app).

Reports include:

- The GPS position of the reporter.
- Precision of the GPS fix.
- Measured radiation value.

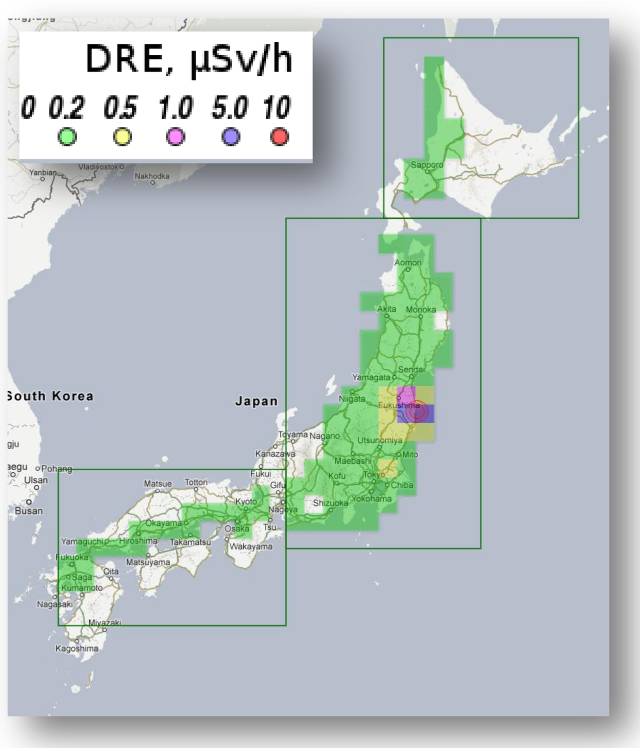

Problem:
Reports can be inaccurate due to the **noise** of the GPS and the **untrustworthiness** of the human reporter.

Objective:
Make accurate spatial predictions of radioactivity.

A Gaussian Process Approach

We used Gaussian process (GP) regression to estimate the spatial function from the reported noisy observations.

- Individual noise terms on the inputs (**Heteroscedastic GP**).

$$y_{i,j} = f(\boldsymbol{x}_{i,j}) + \epsilon_{i,j}$$

Trust model:
Tthe reported precisions are parametrised with an individual **trustworthiness parameter** for each user: $t_i \in [0,1]$

$$\epsilon_{i,j} \sim N(0, t_i \theta_{i,j})$$

GP prior:

$$f \sim GP(0, K(\boldsymbol{x}, \boldsymbol{x}'))$$

Predictive distribution:

$$\Sigma_{\boldsymbol{x}} = \mathrm{diag}\left((t_i \theta_{i,j})^{-1}\right)$$

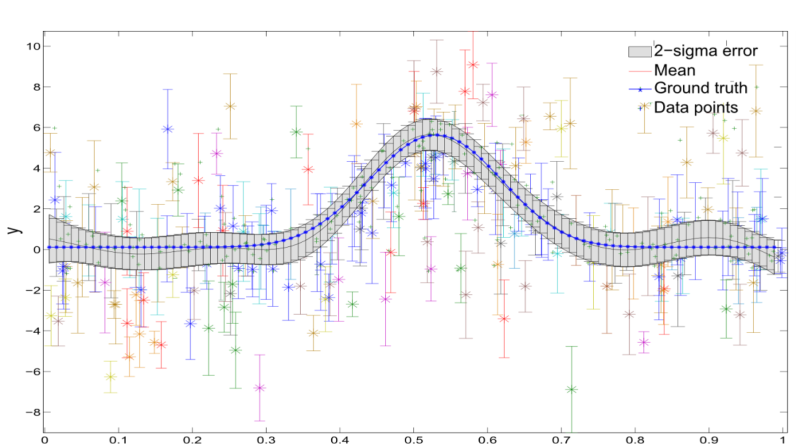

Example of heteroscedastic GP in 1D.

$$E[f(\boldsymbol{x}_*)] = K(\boldsymbol{x}_*, \boldsymbol{x})[K(\boldsymbol{x}, \boldsymbol{x}) + \Sigma_{\boldsymbol{x}}]^{-1}\boldsymbol{y}$$

$$\sigma^2[f(\boldsymbol{x}_*)] = K(\boldsymbol{x}_*, \boldsymbol{x}_*) + K(\boldsymbol{x}_*, \boldsymbol{x})[K(\boldsymbol{x}, \boldsymbol{x}) + \Sigma_{\boldsymbol{x}}]^{-1}K(\boldsymbol{x}, \boldsymbol{x}_*)$$

Learn User Trustworthiness

We can learn the trustworthiness of each user by optimising the GP marginal likelihood.

Marginal log-likelihood:

$$\log p(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{t}) = -\frac{1}{2}\boldsymbol{y}^{\boldsymbol{T}}[K(\boldsymbol{x}, \boldsymbol{x}) + \Sigma_{\boldsymbol{x}}]^{-1}\boldsymbol{y} - \log(K(\boldsymbol{x}, \boldsymbol{x}) + \Sigma_{\boldsymbol{x}}) - \frac{n}{2}\log(2\pi)$$

The TrustGP algorithm:

$$t_{max} = \arg\max_{t}(\log p(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{t}))$$

- **The computational cost of the algorithm is $O(n^3)$ to invert the covariance matrix.**
- **We can compute the analytical gradient of the parameters and use efficient optimisation algorithms (e.g. conjugate gradient).**

Radiation Monitoring in Japan

We used the TrustGP algorithm to estimate the radioactivity map of Japan from the data provided by a network of crowdsourced sensors (Cosm sensors) one year after the earthquake.

Datasets

- **Cosm**: 557 sensors provided by internet users with avg. frequency of 48 readings per day (cosm.com).
- **Speedi:** Sensor network provided by the MEXT (Ministry of science of Japan) with 2122 sensors with frequency of 144 readings per day.
- We computed the mean and variance of the sensor readings taken over a single day (1 April 2012).
- There are 188 overlapping sensors between the two networks.

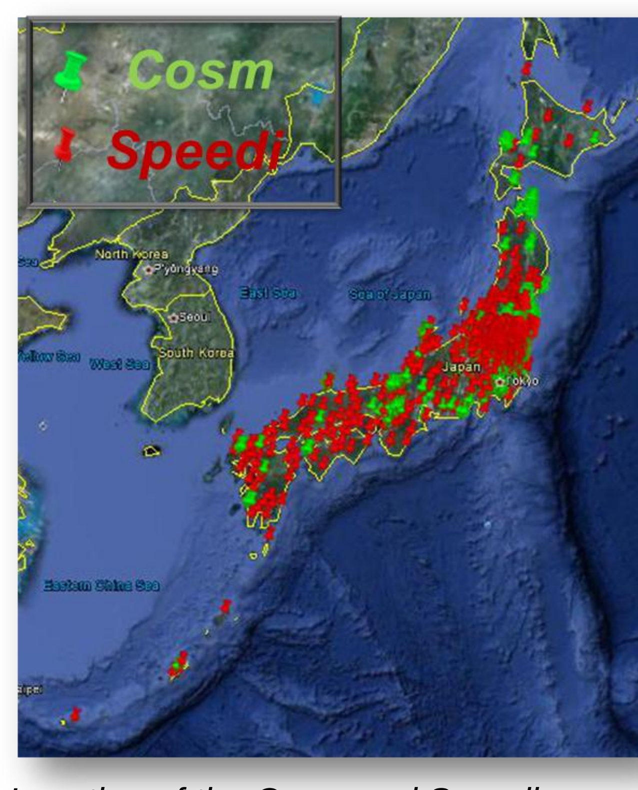

Location of the Cosm and Speedi sensors.

Results

- **The prediction of the standard GP run on the Speedi dataset is taken as the ground truth.**
- **We compared the standard GP vs Trust GP on the Cosm dataset against the ground truth.**
- **The TrustGP recognised 1% of the Speedi sensors and 17% of the Cosm sensors as untrustworthy respectively (t<0.5).**

The TrustGP makes 88% better in prediction, scored by the Continuous Rank Probability Scoring rule (CRPS).

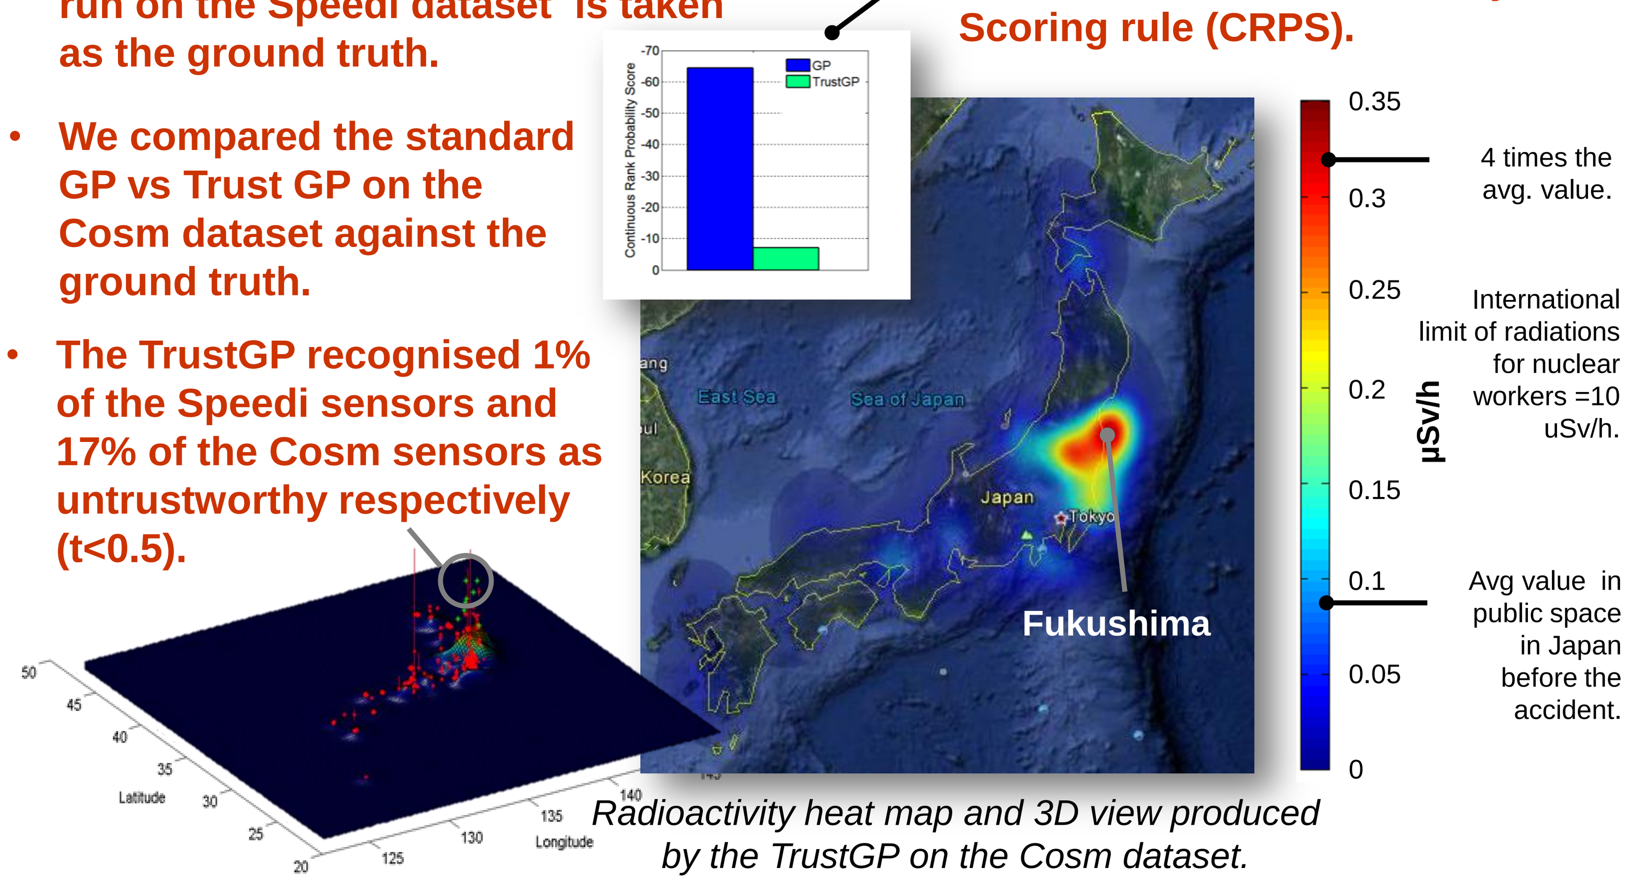

Radioactivity heat map and 3D view produced by the TrustGP on the Cosm dataset.

Conclusion: Considering trustworthiness of information sources improves the quality of spatial predictions from unreliable crowdsourced data.

Future work: (i) Model the temporal dimension of trustworthiness, (ii) integrate mobile sensors in the predictor.

BAE SYSTEMS · ACFR Australian Centre for Field Robotics · SECURE · University of Southampton