Environmental Sensor Placement with Convolutional Gaussian Neural Processes

British Antarctic Survey
NATURAL ENVIRONMENT RESEARCH COUNCIL

Tom R. Andersson[1], Wessel P. Bruinsma[2], Stratis Markou[3], James Requeima[3,4], Alejandro Coca-Castro[5], Anna Vaughan[3], Anna-Louise Ellis[6], Matthew Lazzara[7], Daniel C. Jones[1], J. Scott Hosking[1,5], Richard E. Turner[3]

[1]British Antarctic Survey, [2]Microsoft Research, [3]University of Cambridge, [4]Invenia Labs, [5]The Alan Turing Institute, [6]Met Office, [7]University of Wisconsin-Madison

UNIVERSITY OF CAMBRIDGE

The Alan Turing Institute

Summary

- **Problem**: Sensor placement with non-stationary spatiotemporal data.
- **ConvGNP** *meta-learns* spatial covariance.
- **Active learning** with mutual information identifies highly informative placements.

Problem setup

- Goal: minimise uncertainty over a complicated, non-stationary spatiotemporal function, $f_\tau(x)$, by proposing new sensor locations
- Ground truth $f_\tau(x)$ = simulated Antarctic surface air temperature

ConvGNP: A meta-learned mapping from raw data to a GP predictive

- ConvGNP[1,2]: Learns a map from context data $C = \{(x_i^{(c)}, y_i^{(c)})\}_{i=1}^{N_c}$ and target inputs $\mathbf{x}^{(t)}$ to a Gaussian over target outputs $\mathbf{y}^{(t)}$
- Training: Minimise negative log-likelihood (NLL) using randomly sampled $(\mathbf{x}^{(c)}, \mathbf{y}^{(c)})$ and $(\mathbf{x}^{(t)}, \mathbf{y}^{(t)})$ over 1980-2013

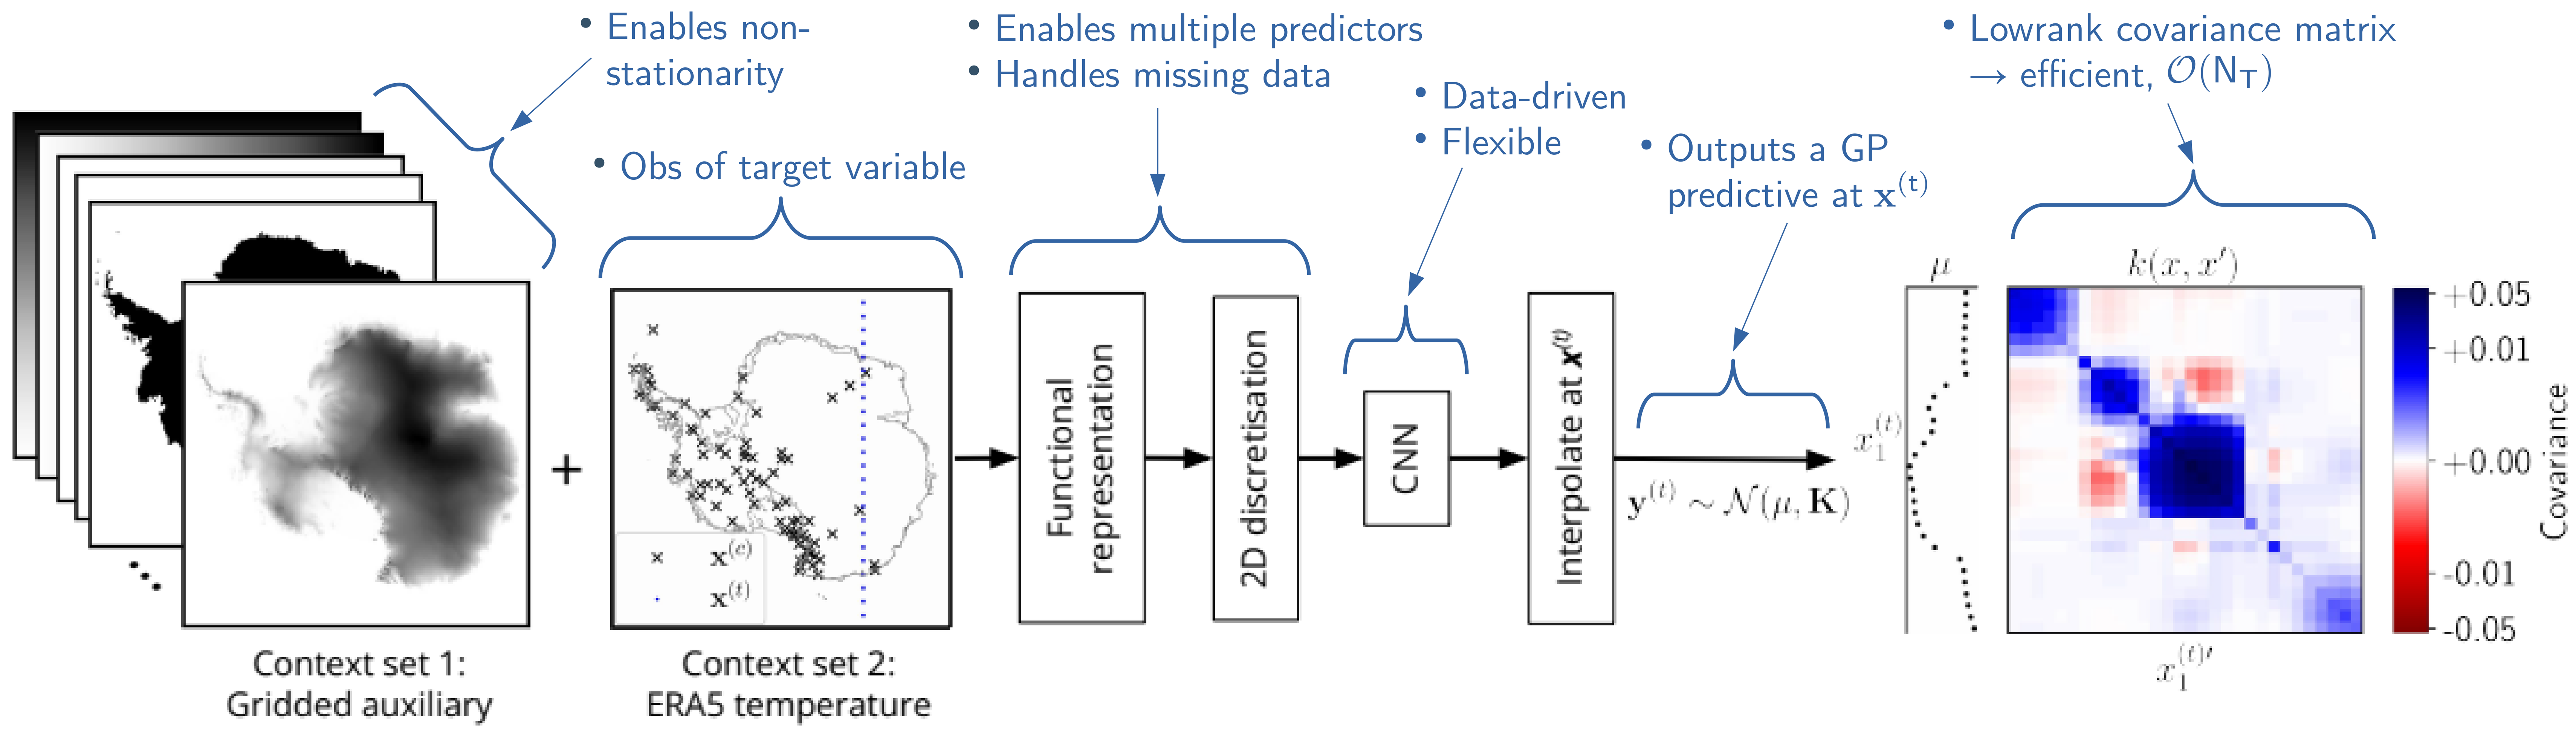

ConvGNP learns non-stationary covariance structure

- ConvGNP outperforms Gaussian processes (GPs) with a non-stationary kernel (Gibbs) and a stationary kernel (EQ) on test data (2018-2019):

Metric	ConvGNP	Gibbs GP	EQ GP
Normalised NLL	**−1.76**	−1.15	−0.72
MAE (°C)	**0.93**	1.34	2.10

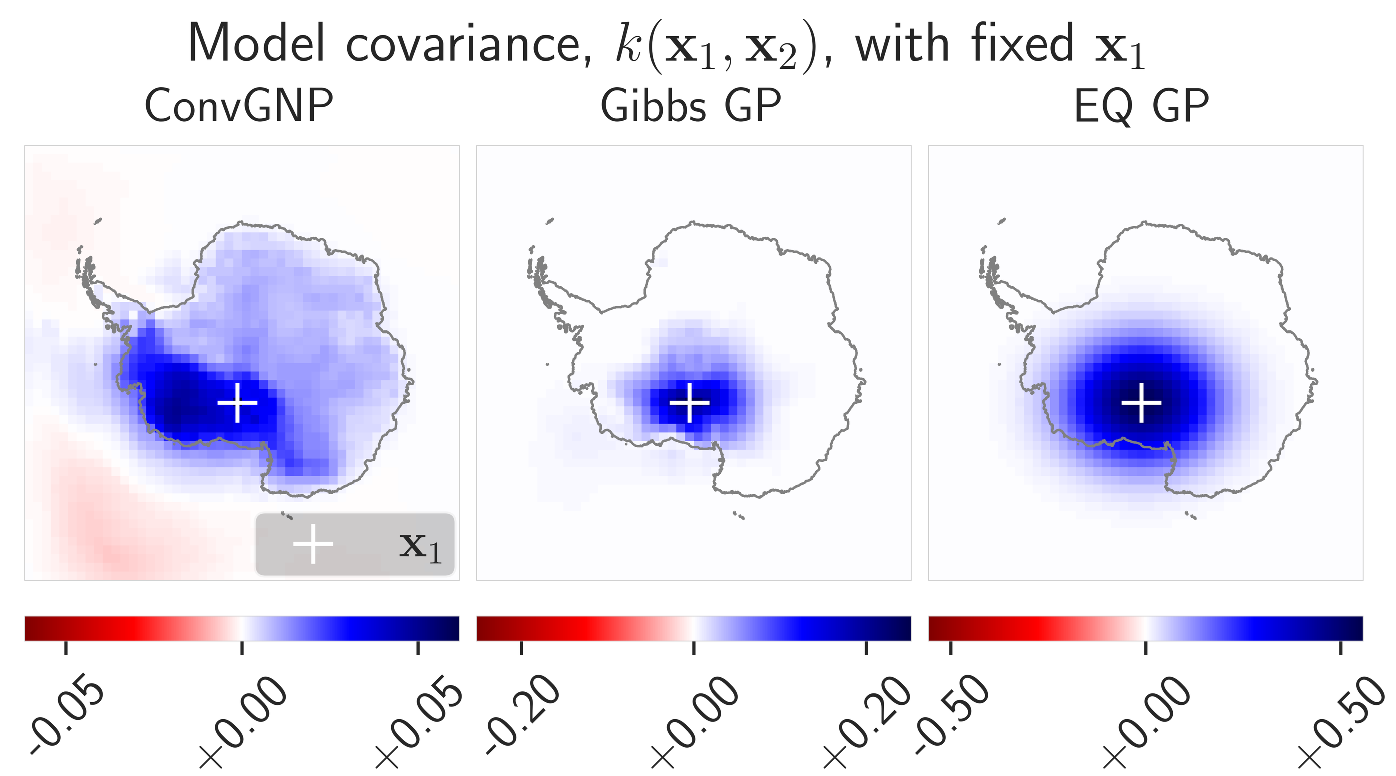

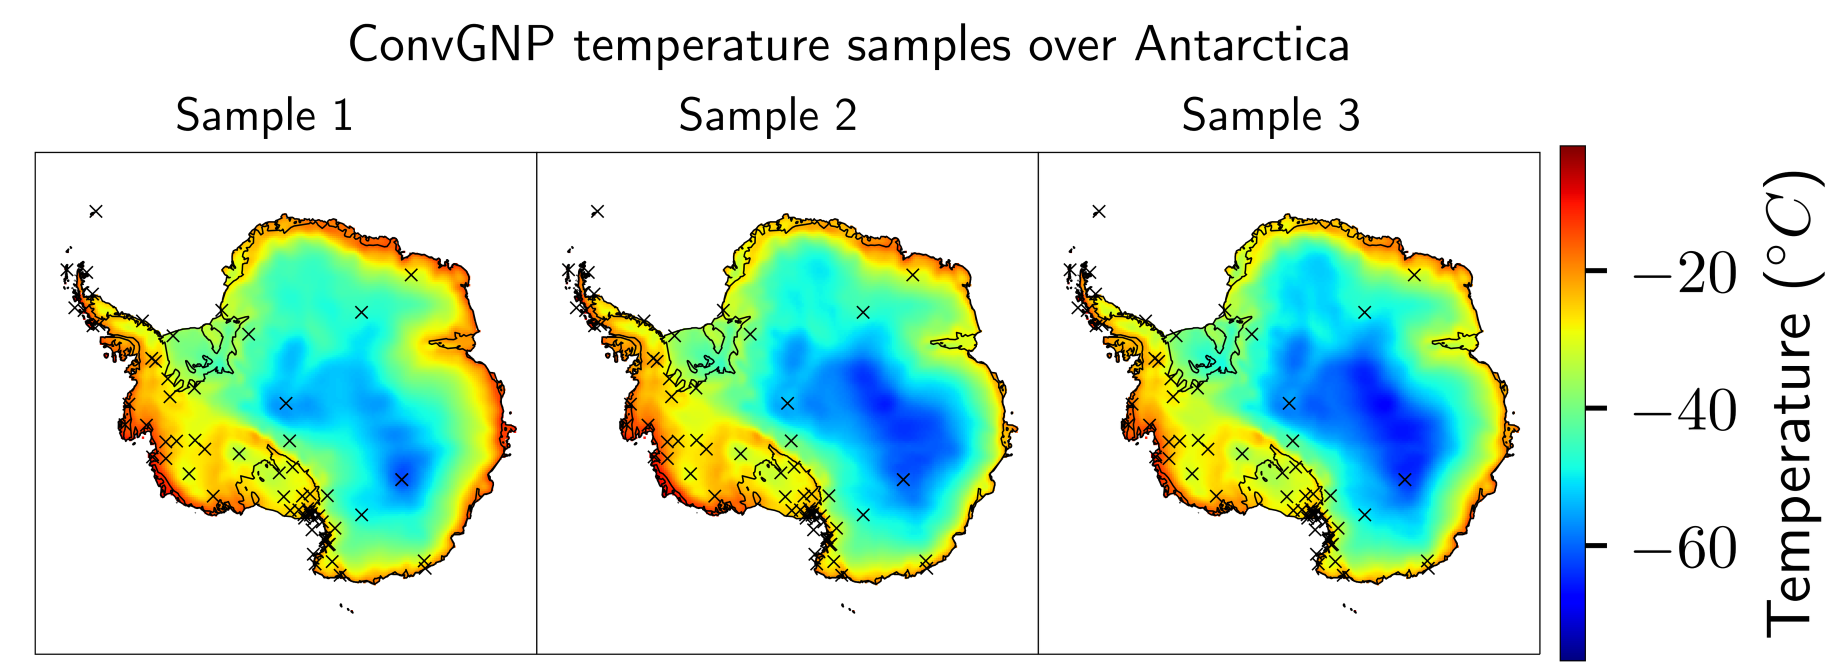

Sensor placement toy experiment

- Greedily place sensors[4] and compute likelihood of true $\mathbf{y}^{(t)}$

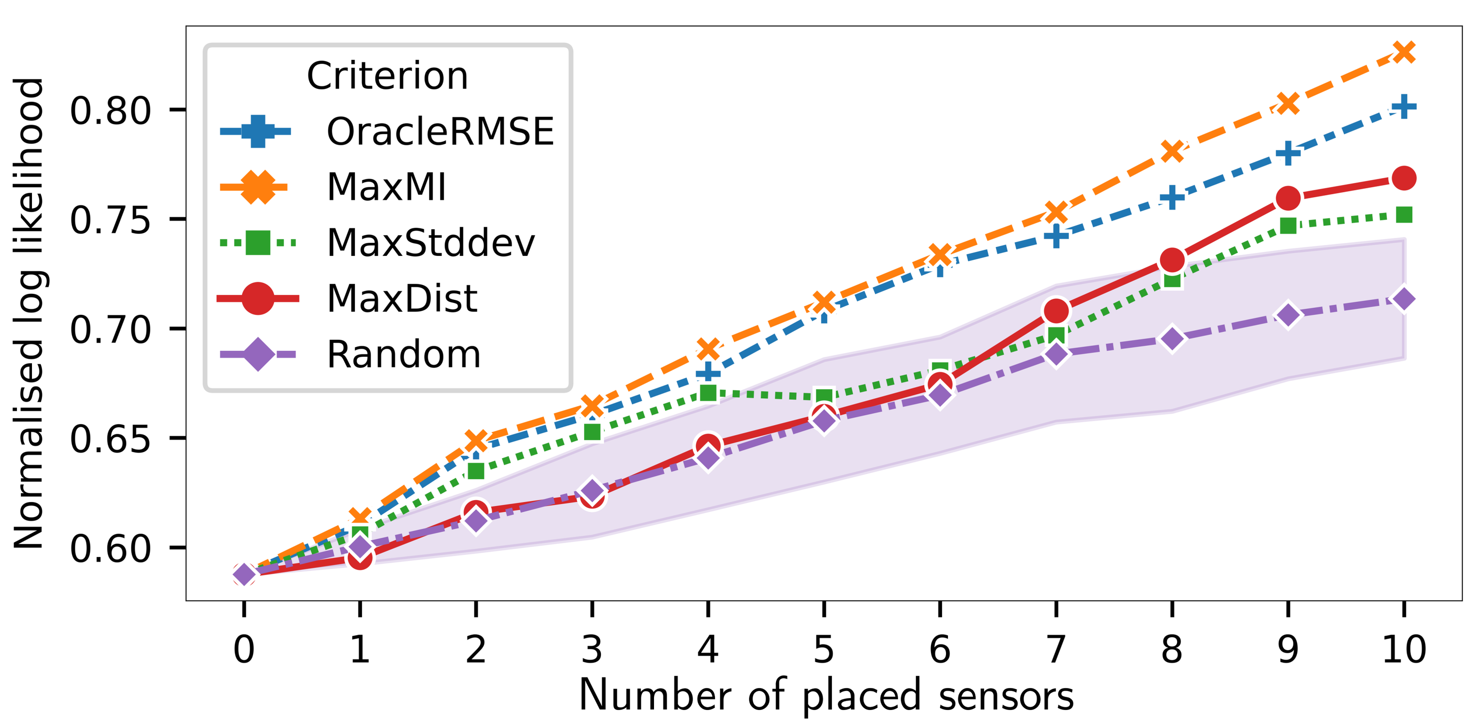

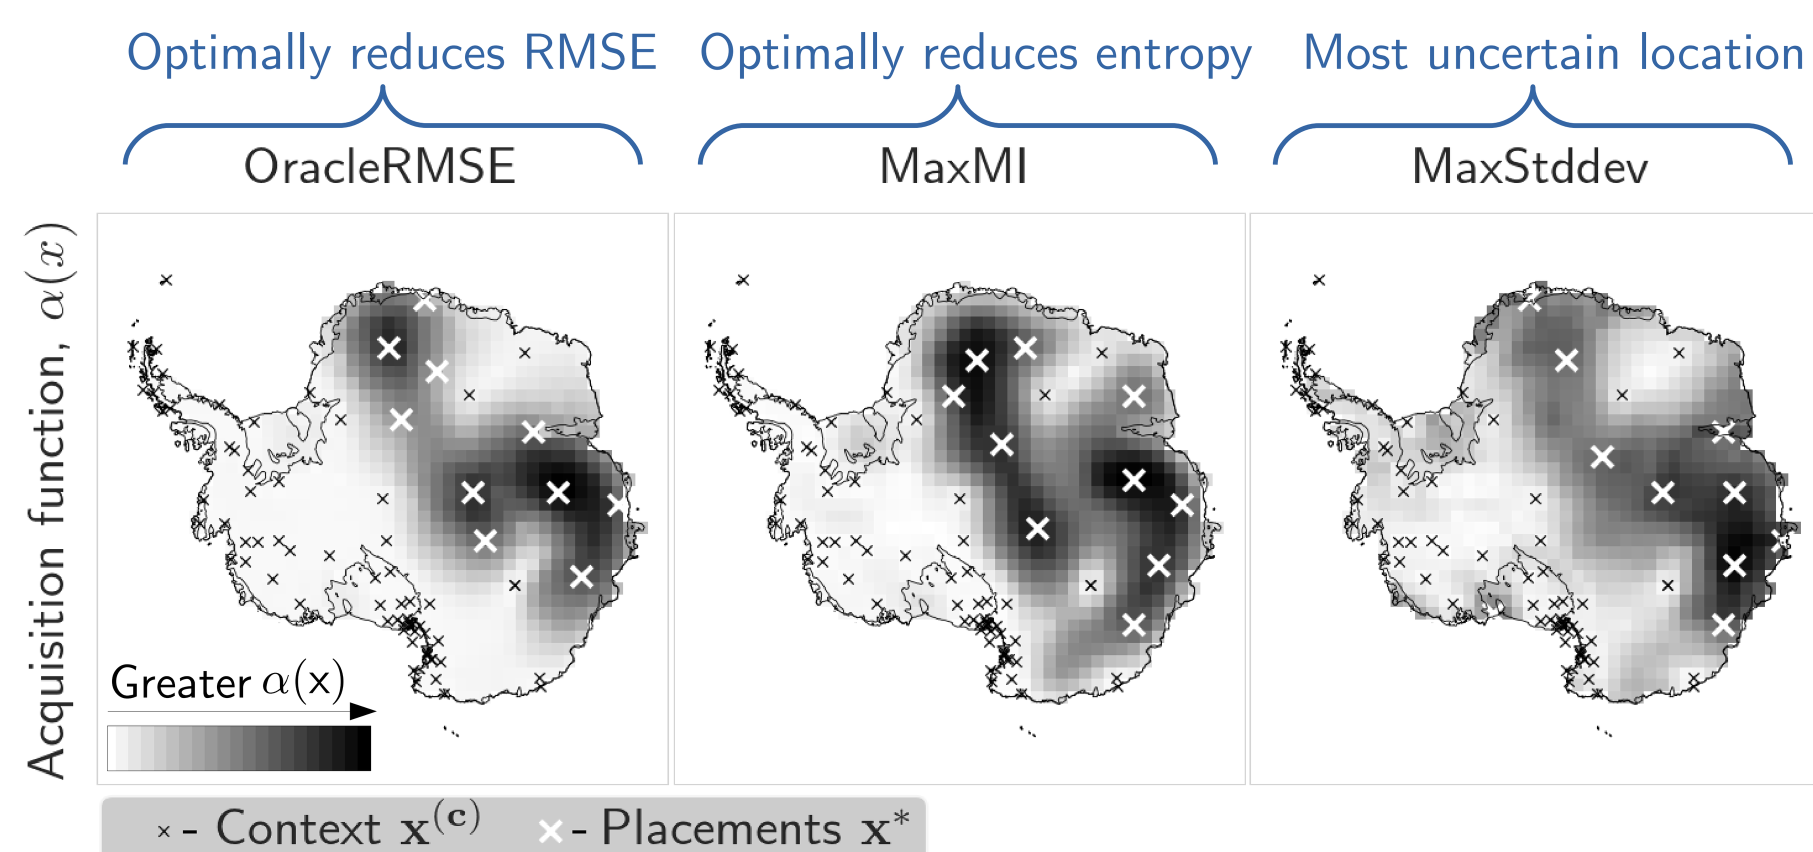

Conclusions

- Vanilla GPs place strong restrictions on the form of the covariance function, $k(\mathbf{x}_1, \mathbf{x}_2)$ → poor performance on non-stationary environmental data
- ConvGNP leverages obs over multiple time steps τ to *meta-learn* arbitrary non-stationary covariance structure
- Active learning using ConvGNP's mutual information can rapidly increase the likelihood of ground truth

Funding: This work was supported by Wave 1 of The UKRI Strategic Priorities Fund under the EPSRC Grant EP/W006022/1, particularly the AI for Science theme within that grant & The Alan Turing Institute.

References
1. S. Markou et al., "Practical Conditional Neural Processes via Tractable Dependent Predictions." ICLR, 2022.
2. W.P. Bruinsma et al., "The Gaussian Neural Process." AABI, 2020.
3. M. Garnelo et al., "Conditional Neural Processes." ICML, 2018.
4. A. Krause et al., "Near-Optimal Sensor Placements in Gaussian Processes." JMLR, 2008.

Paper: arxiv.org/abs/2211.10381