Environmental Sensor Placement with Convolutional Gaussian Neural Processes

Tom Andersson

British Antarctic Survey
NATURAL ENVIRONMENT RESEARCH COUNCIL

UNIVERSITY OF CAMBRIDGE

Microsoft

The Alan Turing Institute

Where to put the next Antarctic weather station?

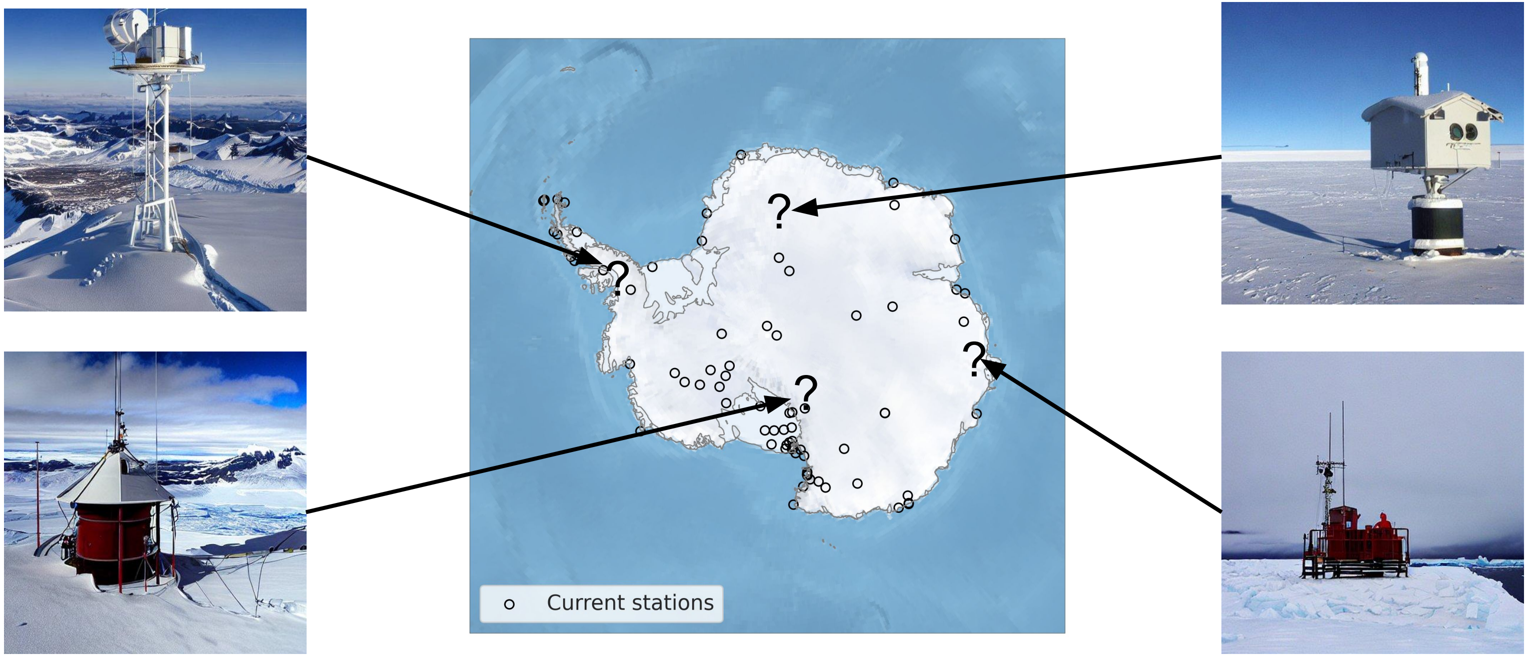

Where to put the next Antarctic weather station?


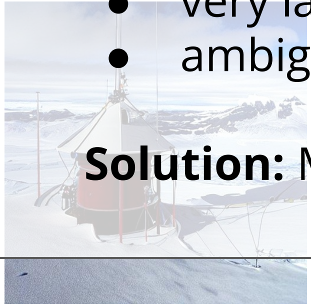
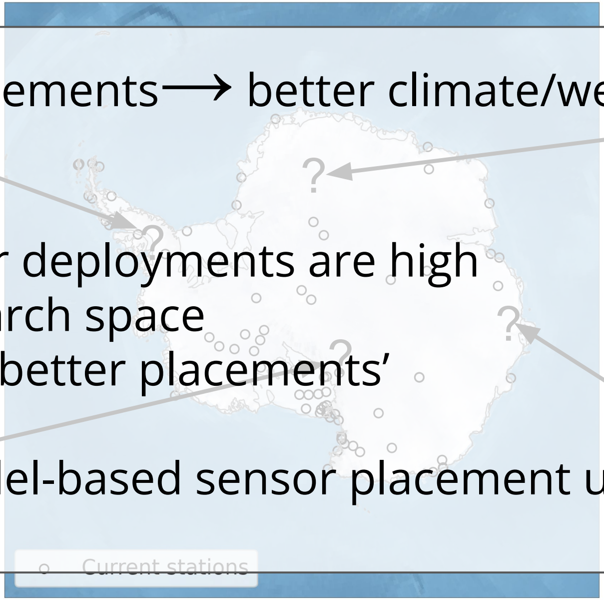


Better sensor placements ⟶ better climate/weather monitoring

But...

- cost of sensor deployments are high
- very large search space
- ambiguity of ‘better placements’

Solution: ML model-based sensor placement using *active learning*

Sensor placement using active learning

- Two ingredients: model $p(\mathbf{y}|\text{data})$ + acquisition function $\alpha(x) \longrightarrow$ placements $\mathbf{x}^*$

Belief about target variable given observed data

Utility of new observations at *x* (e.g. variance of model at *x*)

New *x*-locations that maximise acquisition function

Sensor placement using active learning

- Two ingredients: model $p(\mathbf{y}|\text{data})$ + acquisition function $\alpha(x) \longrightarrow$ placements $\mathbf{x}^*$
  - Gaussian processes (GPs) widely used for $p(\mathbf{y}|\text{data})$

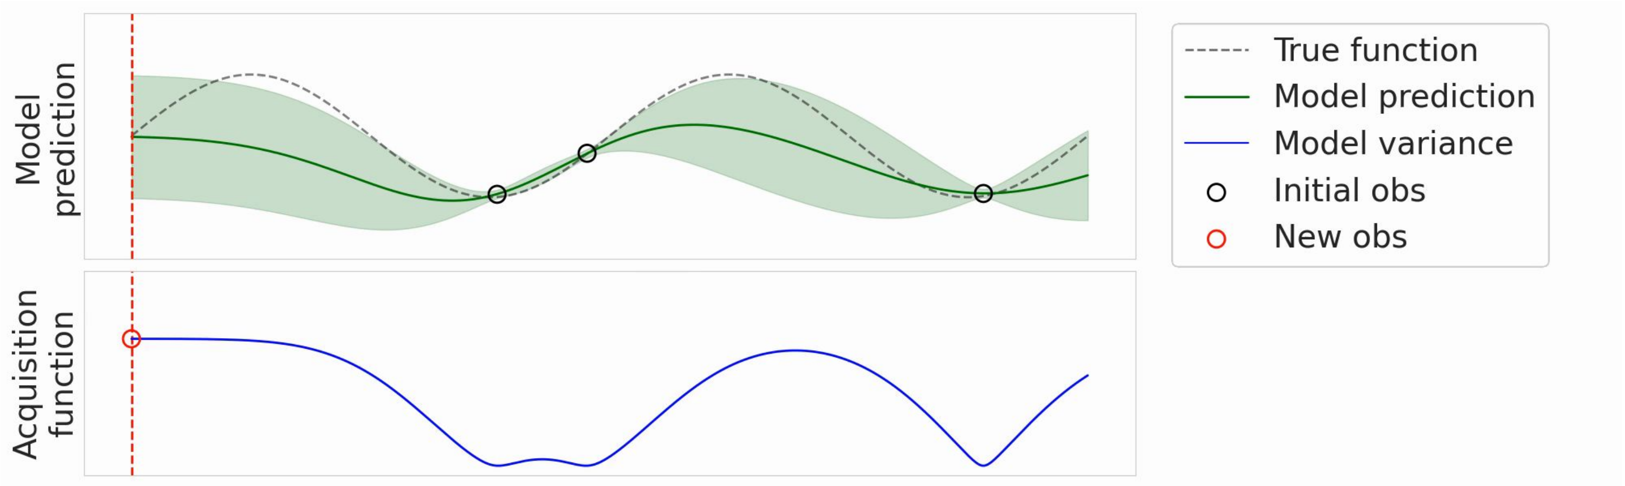

Sensor placement using active learning

- Two ingredients: model $p(\mathbf{y}|\text{data})$ + acquisition function $\alpha(x) \longrightarrow$ placements $\mathbf{x}^*$
  - Gaussian processes (GPs) widely used for $p(\mathbf{y}|\text{data})$

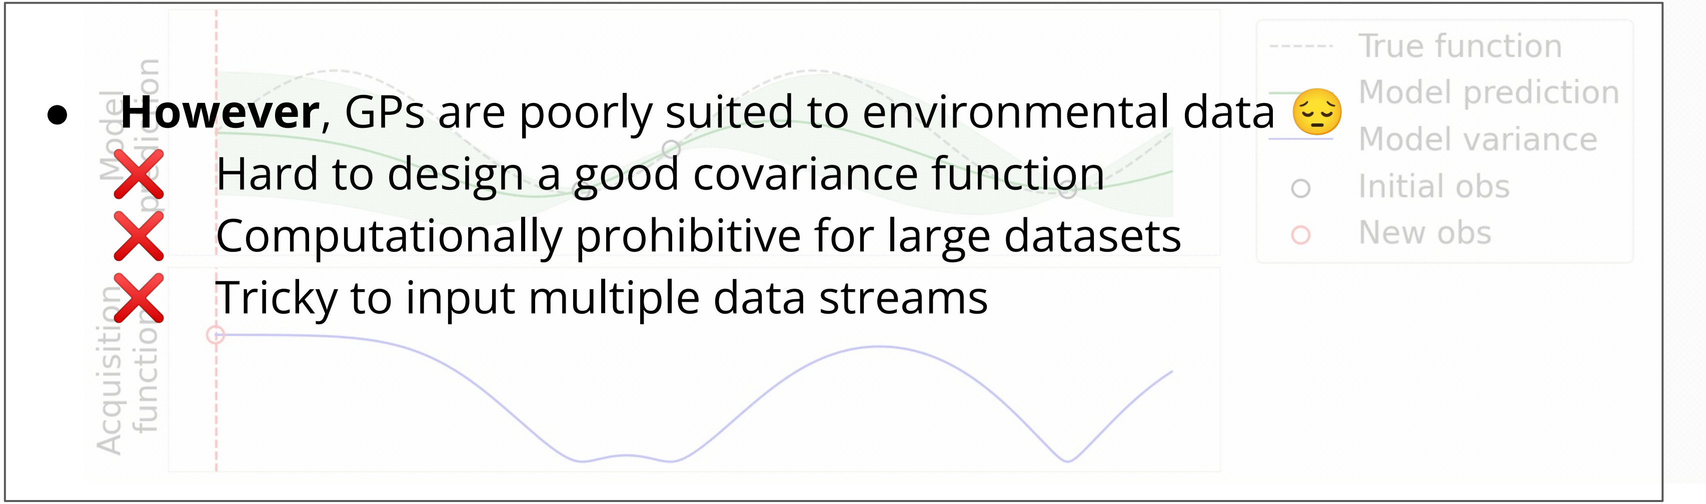

- **However**, GPs are poorly suited to environmental data 😔
  - ❌ Hard to design a good covariance function
  - ❌ Computationally prohibitive for large datasets
  - ❌ Tricky to input multiple data streams

Problem set-up and the ConvGNP

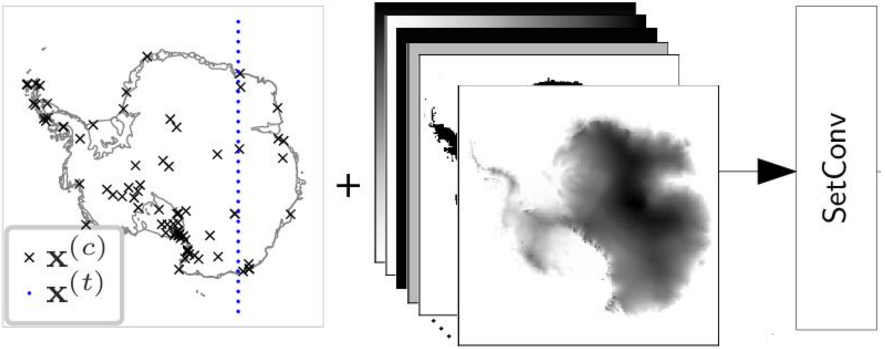

Context set 1:
2m temperature anomaly

Context set 2:
Gridded auxiliary

Problem set-up and the ConvGNP

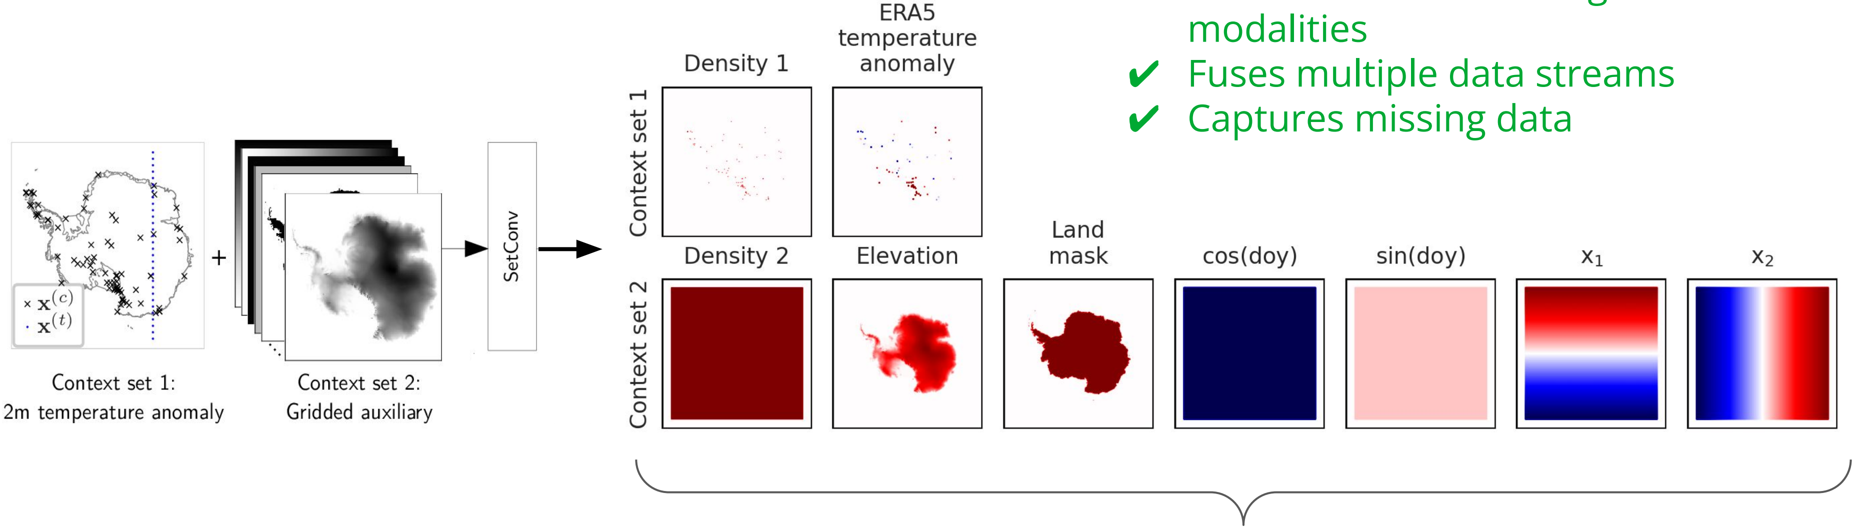

✔ Handles scattered and gridded modalities
✔ Fuses multiple data streams
✔ Captures missing data

3D tensor output by SetConv encoder
- Density channels: observation *locations*
- Data channels: observation *values*

Problem set-up and the ConvGNP

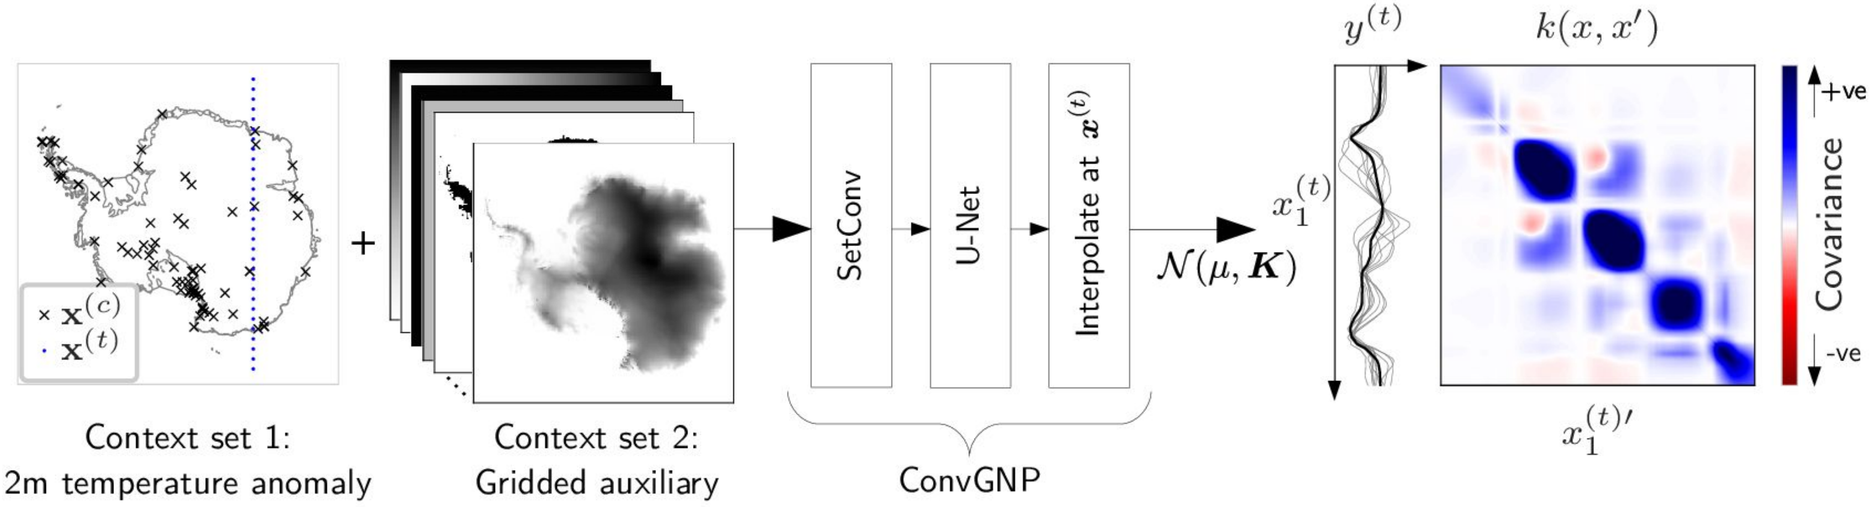

- ✔ Learns to output arbitrary mean & covariance functions given context data
- ✔ Inference scale linearly with dataset size

Training the ConvGNP

We train a ConvGNP to spatially interpolate ERA5 daily-average 2-metre temperature anomaly

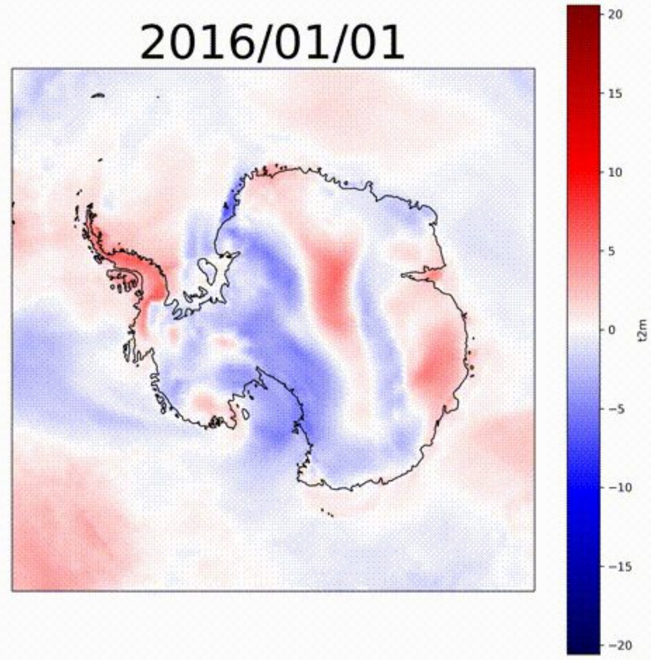

Training the ConvGNP

ERA5 2m temperature anomaly **context** points

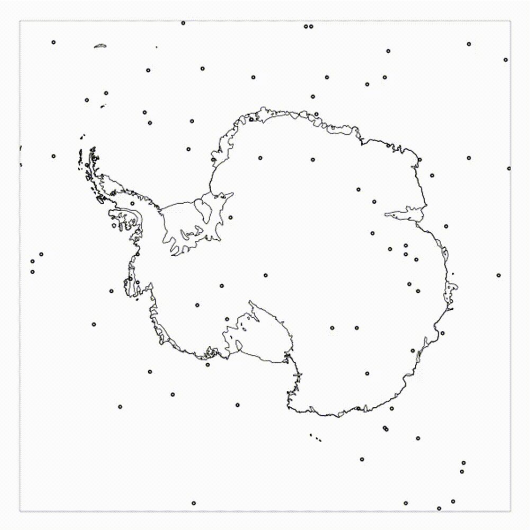

Train ~4 million parameters to maximise target probability

ConvGNP

ERA5 2m temperature anomaly **target** points

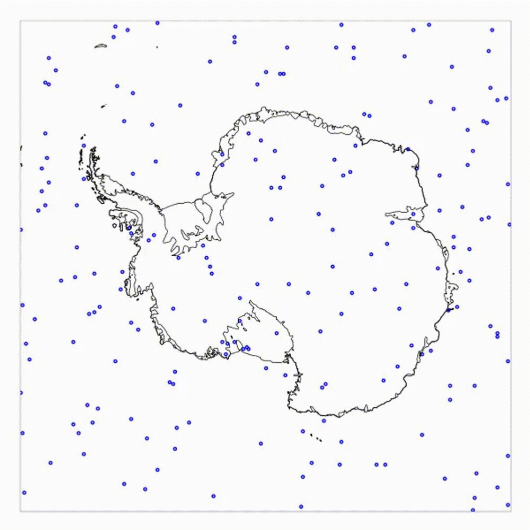

The ConvGNP learns non-stationary covariance, $k(x_1, x_2)$

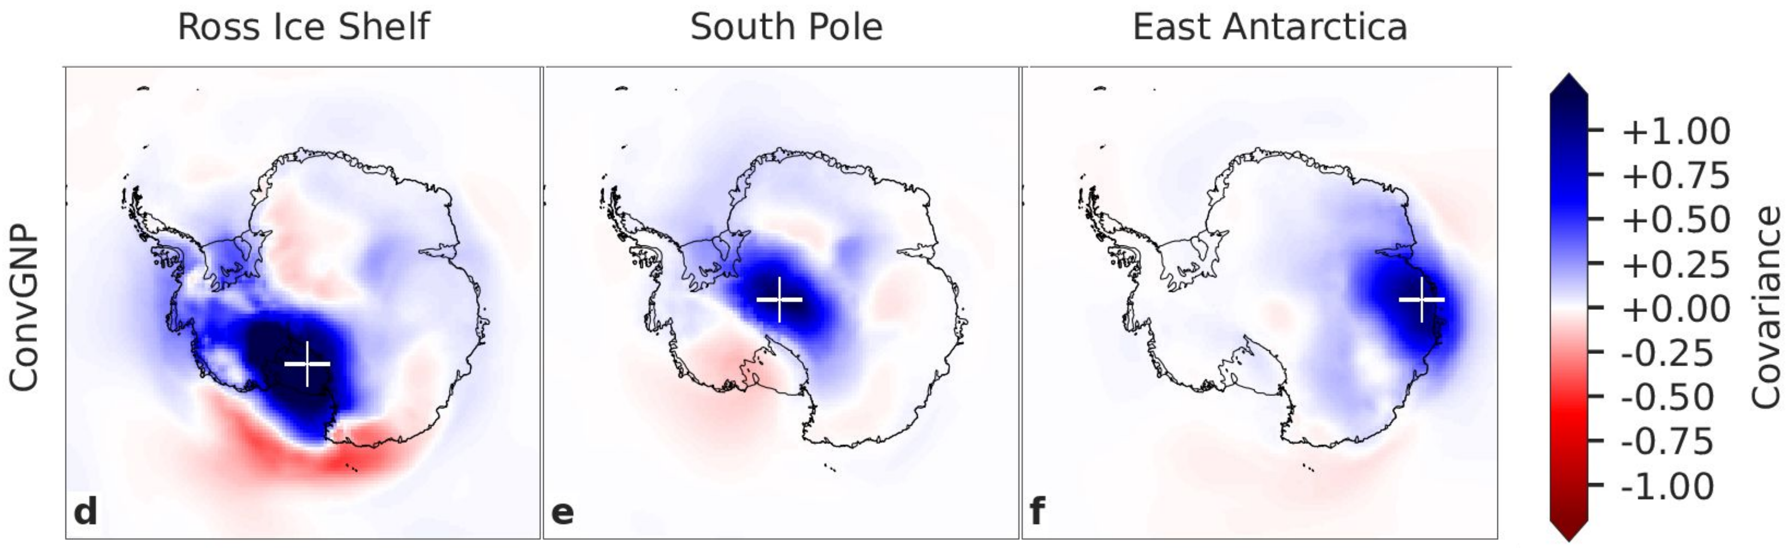

Sensor placement experiment

1) Initialise with ERA5 context points (○) fixed at real Antarctic station locations

2) $\alpha(x)$ = predicted change in variance across continent with query observation at x appended to context set

3) Append best query observation to context set and repeat

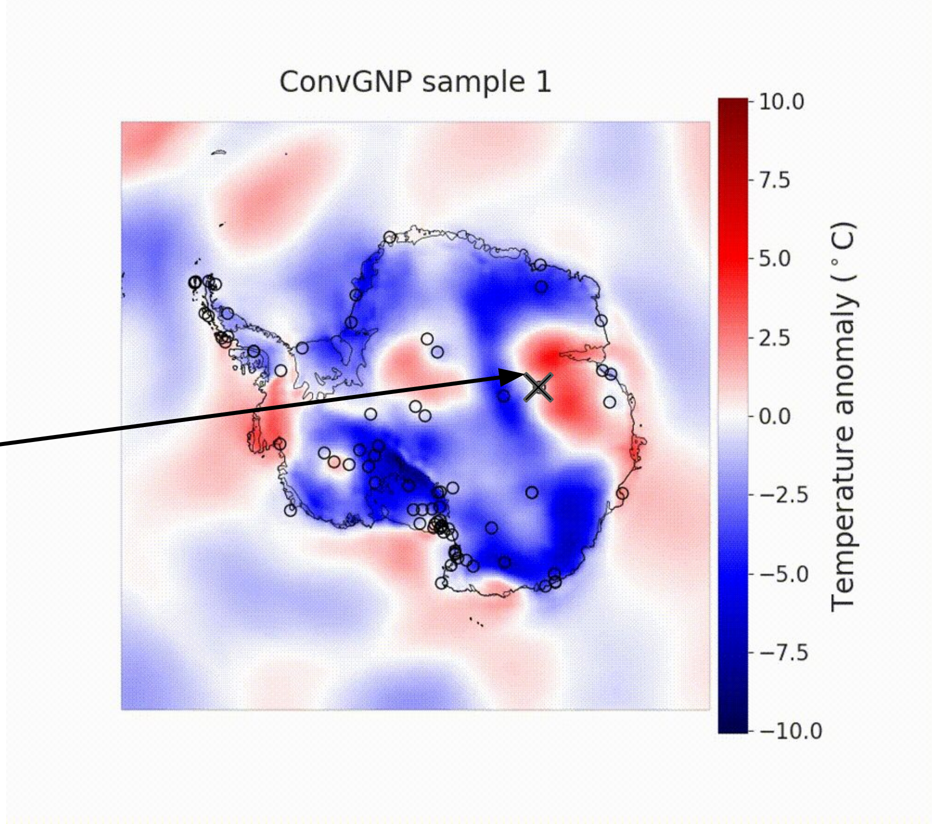

The ConvGNP finds highly informative sensor placements

<u>ConvGNP</u>

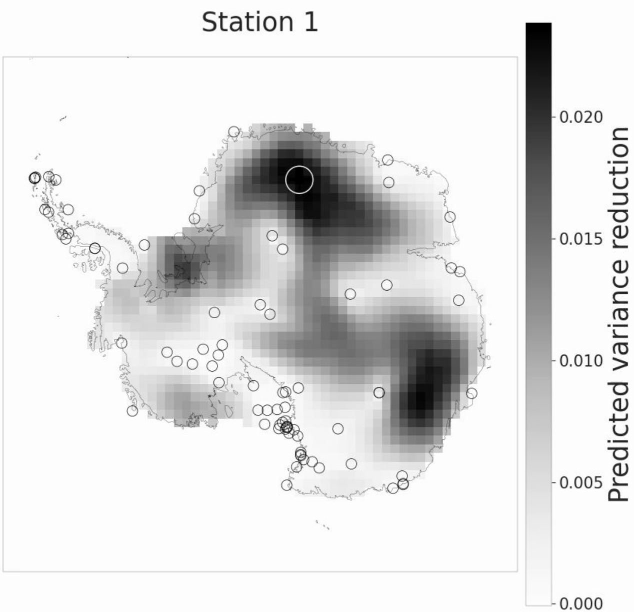

<u>GP baseline</u>

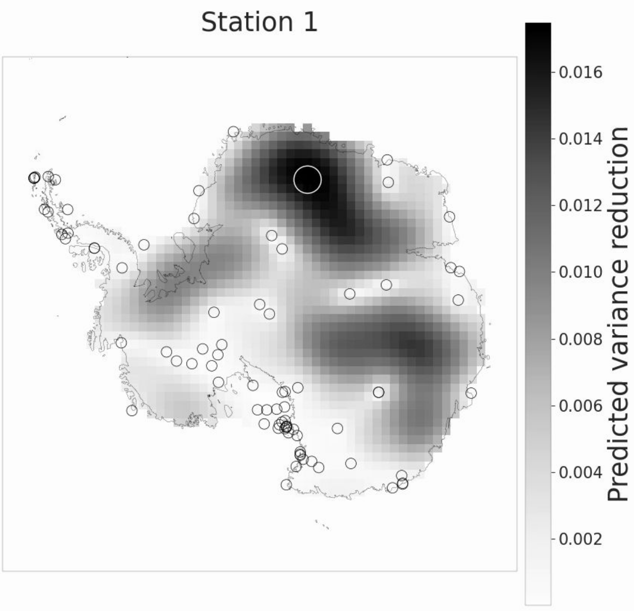

The ConvGNP finds highly informative sensor placements

ConvGNP

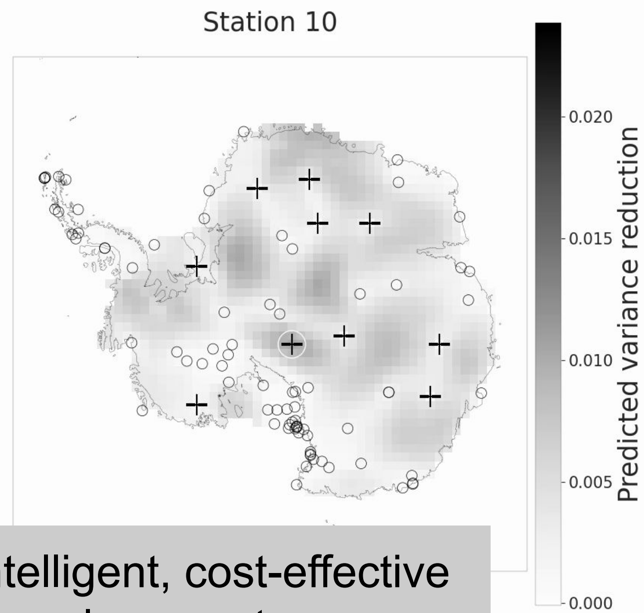

✅ Intelligent, cost-effective placements

GP baseline

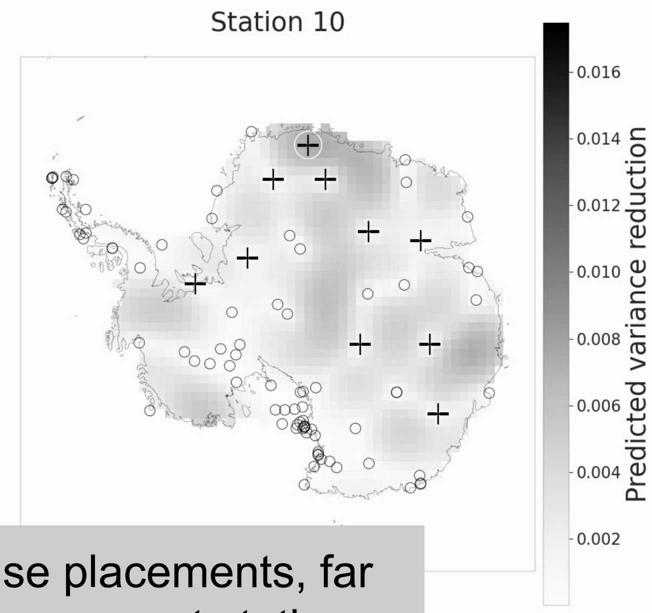

❌ Diffuse placements, far away from current stations

The ConvGNP finds highly informative sensor placements

Reveal ground truth to models in order of proposals.

ConvGNP:

✅ starts off with better RMSE

✅ reduces its error faster

(see paper for probabilistic metrics)

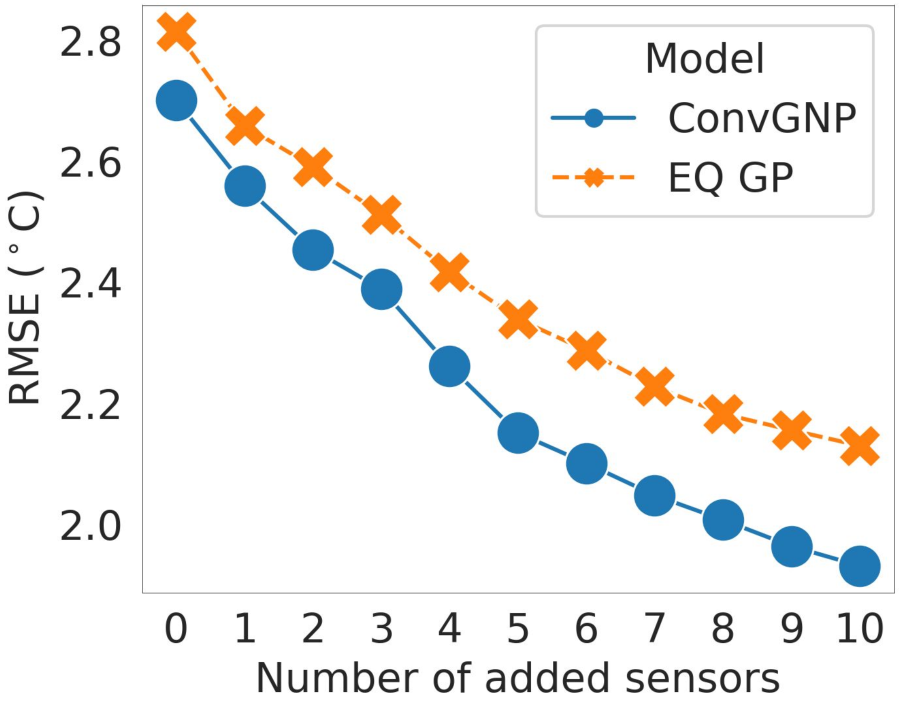

Limitations

1. Our training procedure: Model learns from reanalysis, not real observations
2. ConvGNP: Data hungry (needs to *learn how to condition on data*)

Some future work

- Sim2Real by fine-tuning on real observations (Jonas Scholz's Cambridge MPhil project)
- Python package (`deepsensor`)
- Propose sensor *trajectories* for AUVs

The future = ConvGNP as a Digital Twin component?

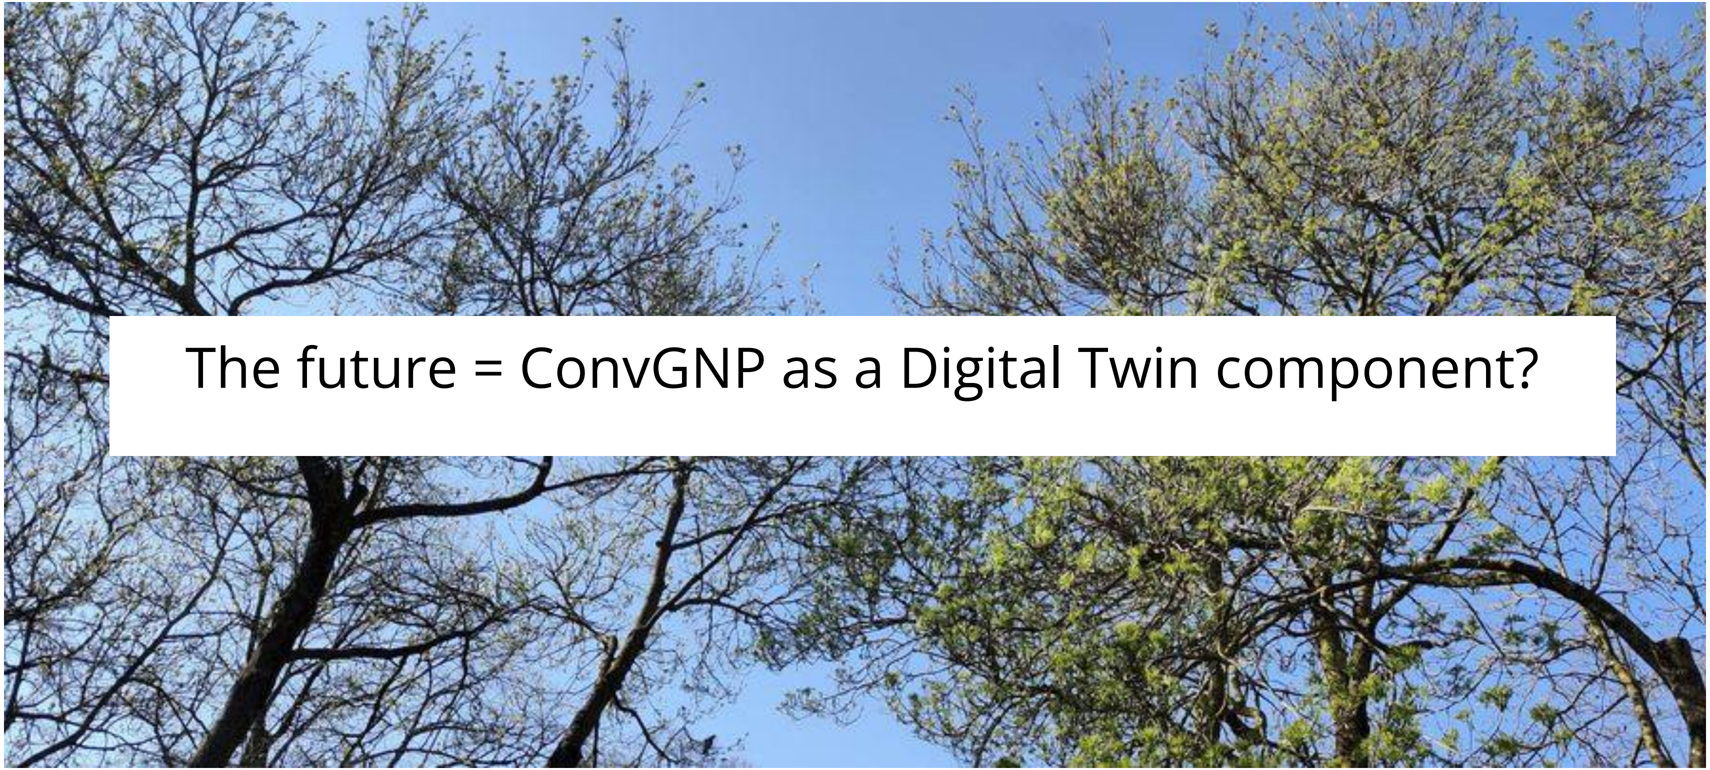

Thanks for listening!

Pre-print

Environmental Sensor Placement with Convolutional Gaussian Neural Processes

Contact

tomand@bas.ac.uk

tom_r_andersson

People


Wessel Bruinsma

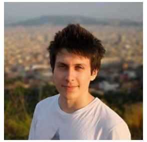
Stratis Markou


James Requiema


Alejandro Coca-Castro


Anna Vaughan


Anna-Louise Ellis


Matthew Lazzara


Dani Jones


Scott Hosking


Rich Turner

Places

British Antarctic Survey
NATURAL ENVIRONMENT RESEARCH COUNCIL

UNIVERSITY OF CAMBRIDGE

Microsoft

The Alan Turing Institute

Appendix

The ConvGNP accurately predicts true performance gain

Spatial pattern (joint) probability | Point-wise (marginal) probability | RMSE

Uncertainty-based ConvGNP **prediction**

True ConvGNP performance gain

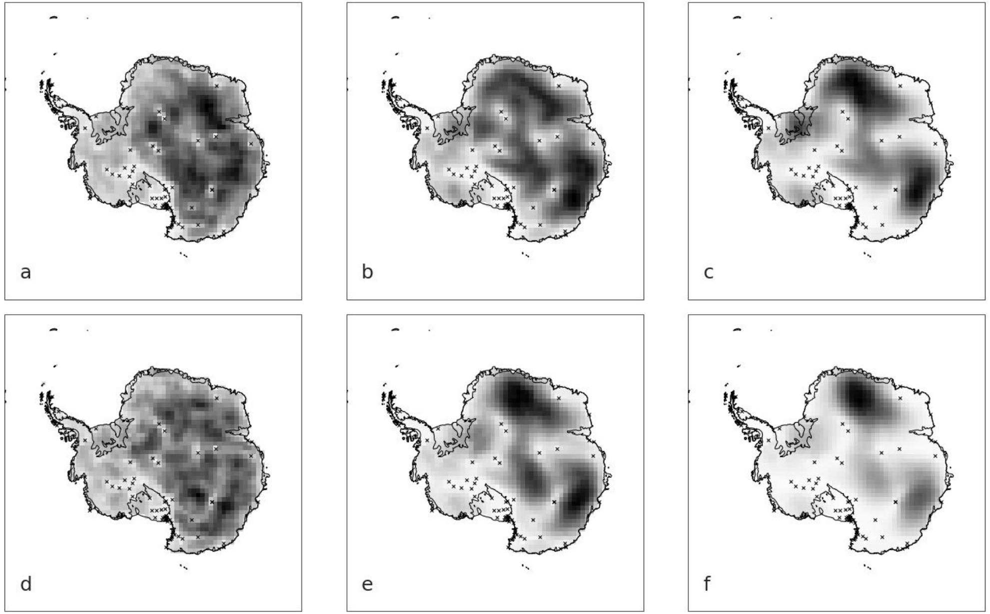

Trading off informativeness with cost

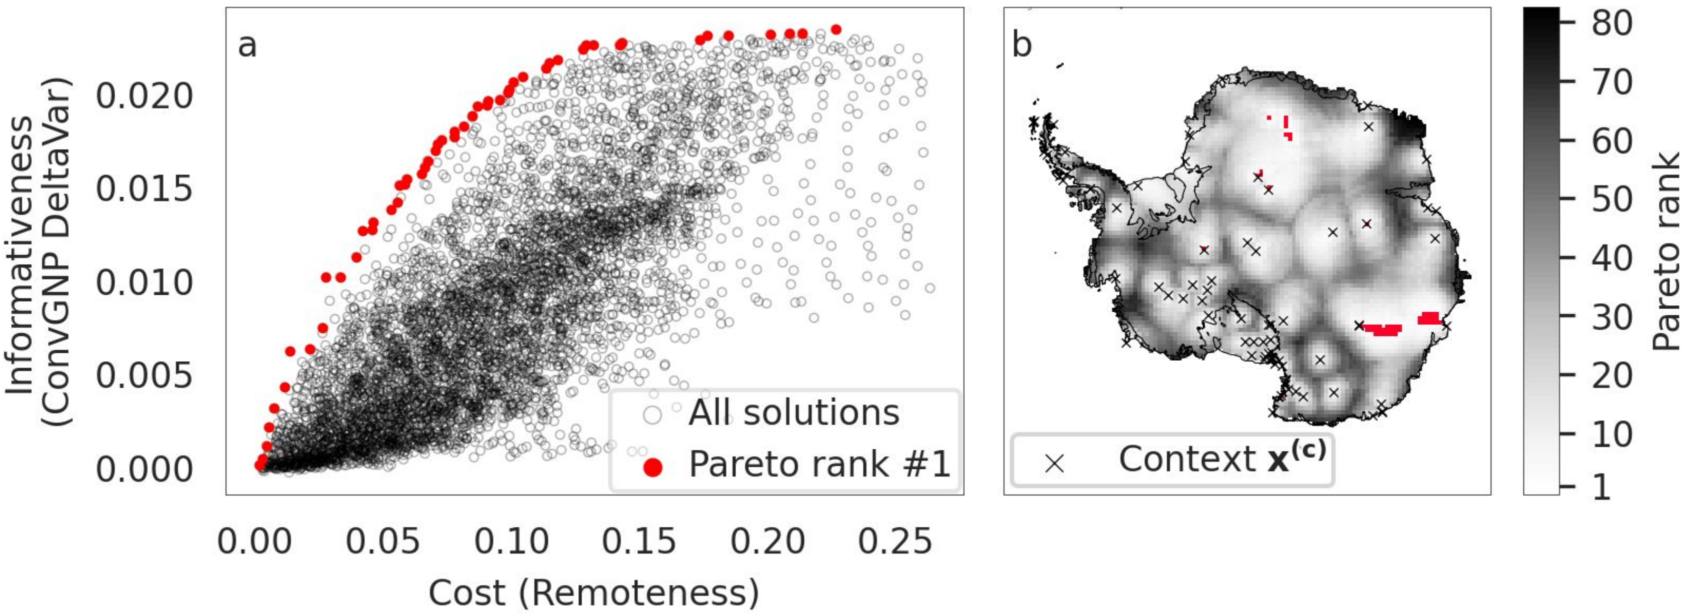

Regression results on unseen data (2018-2019)

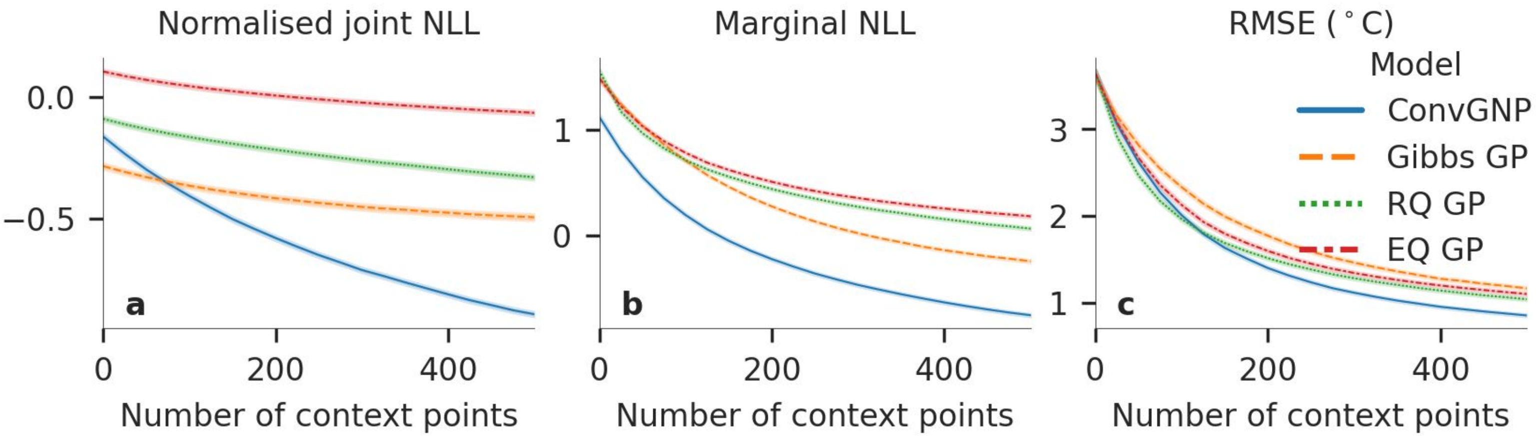

Oracle acquisition function results

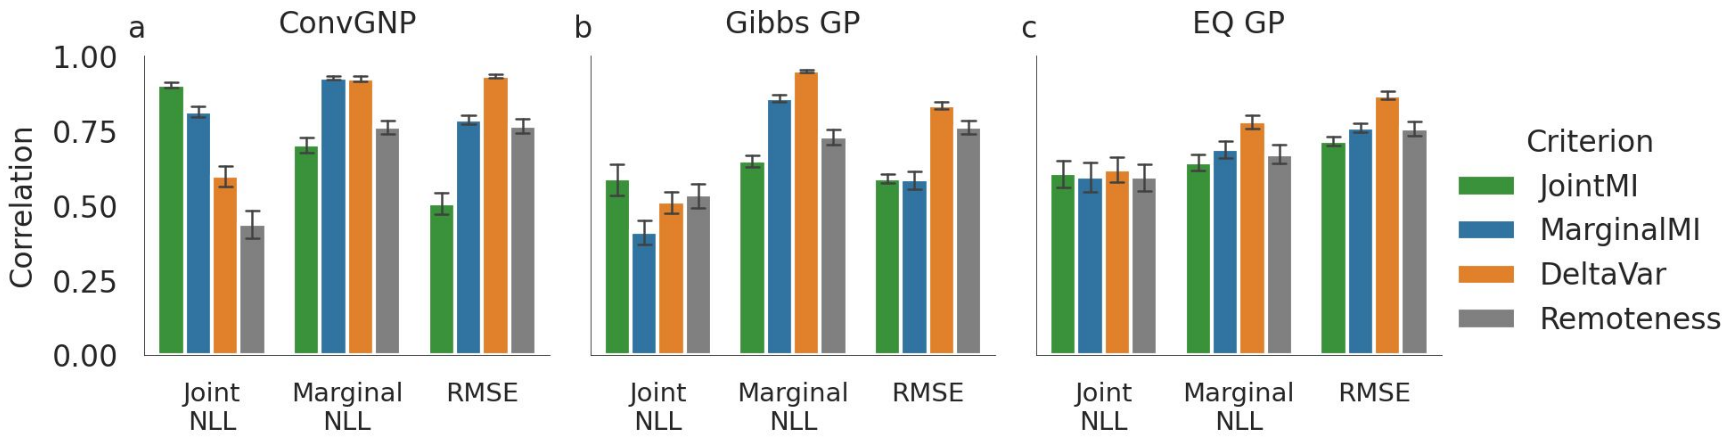

Sensor placement results

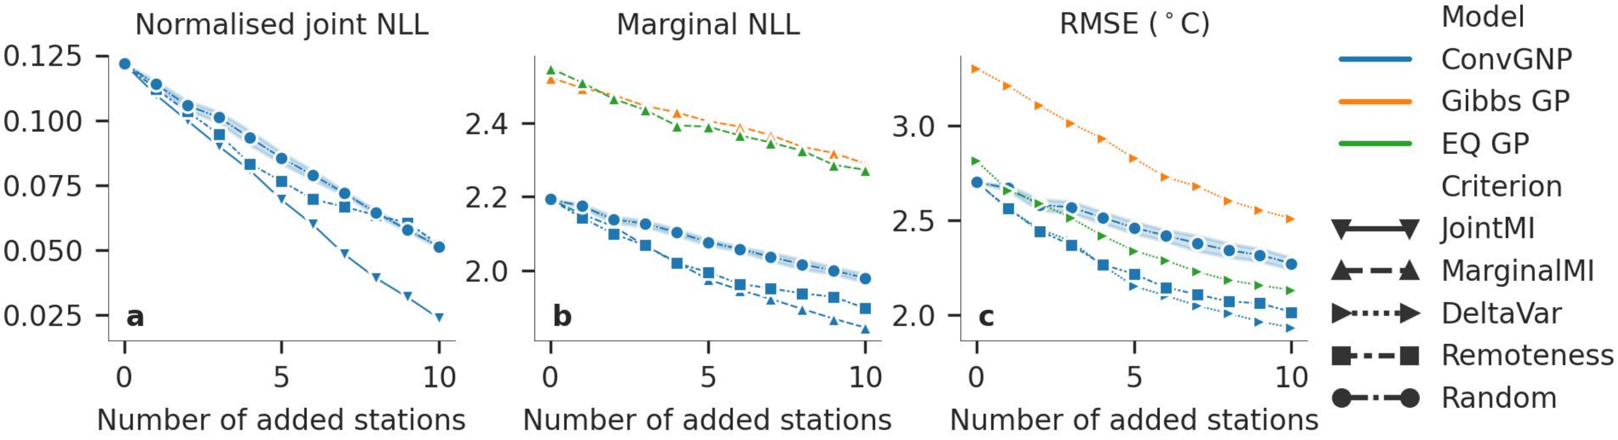

The ConvGNP learns seasonally-varying non-stationary covariance

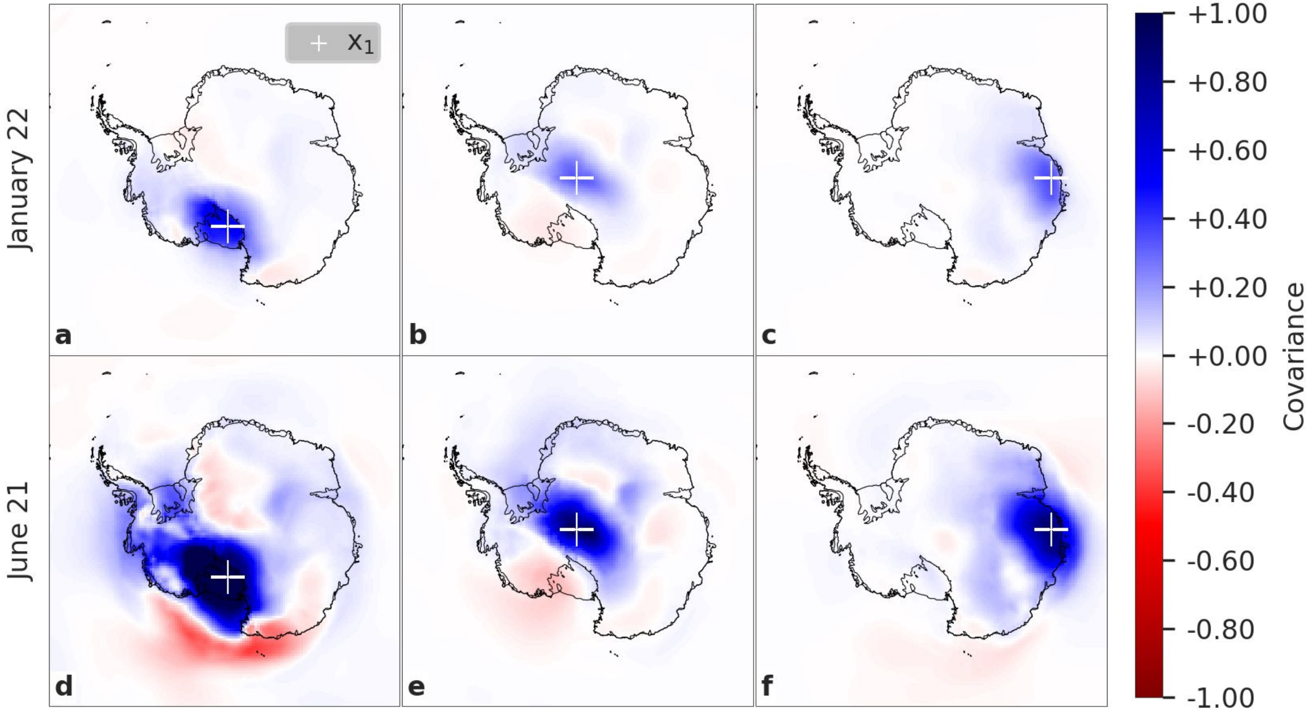

The ConvGNP learns seasonally-varying correlation

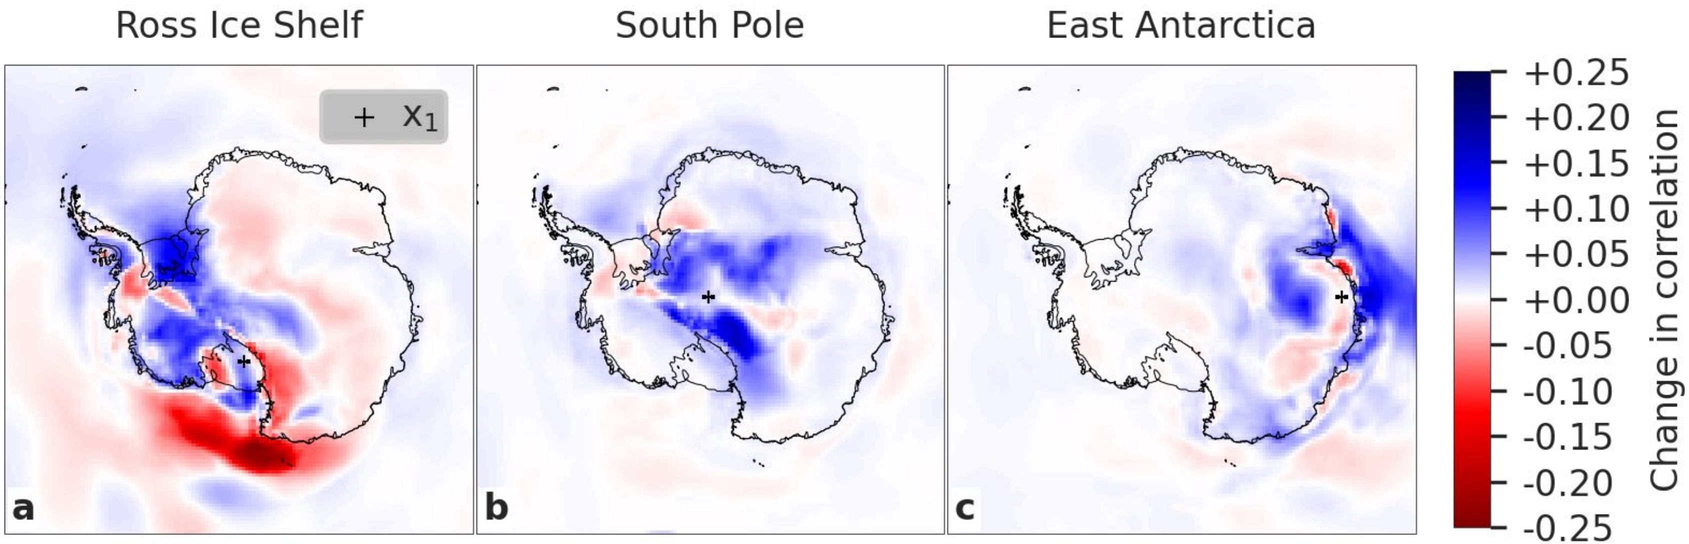

After training, the ConvGNP extrapolates plausible scenarios away from data

ConvGNP samples with ERA5 context points (○) fixed at real Antarctic station locations

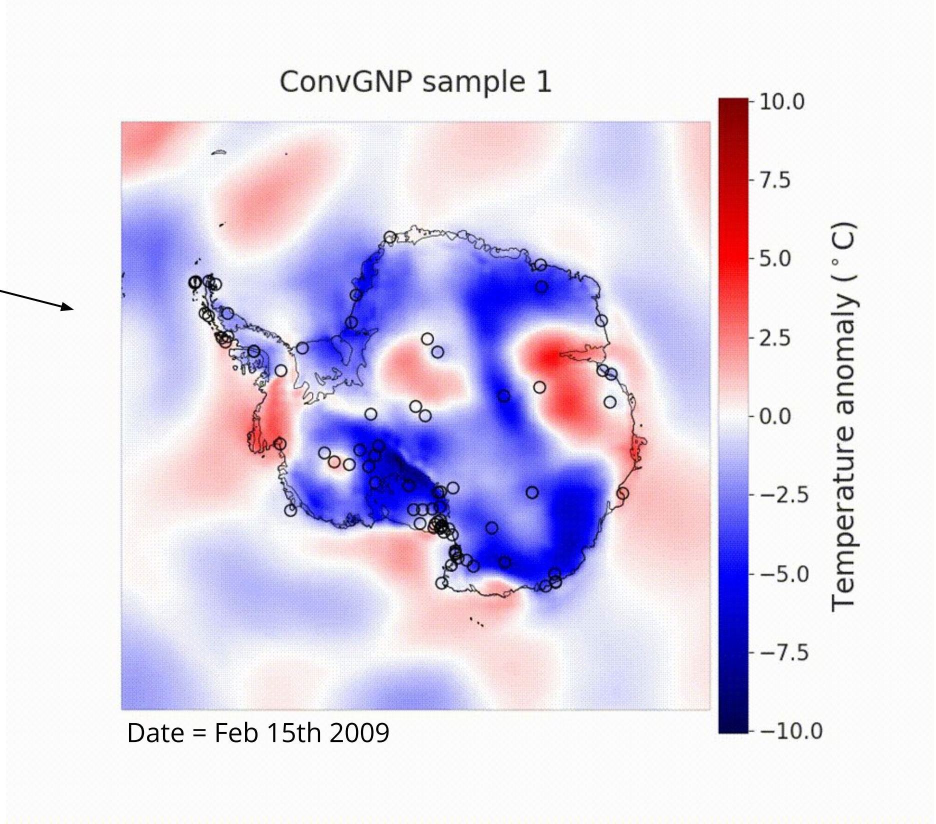

Comparison of models and acquisition function

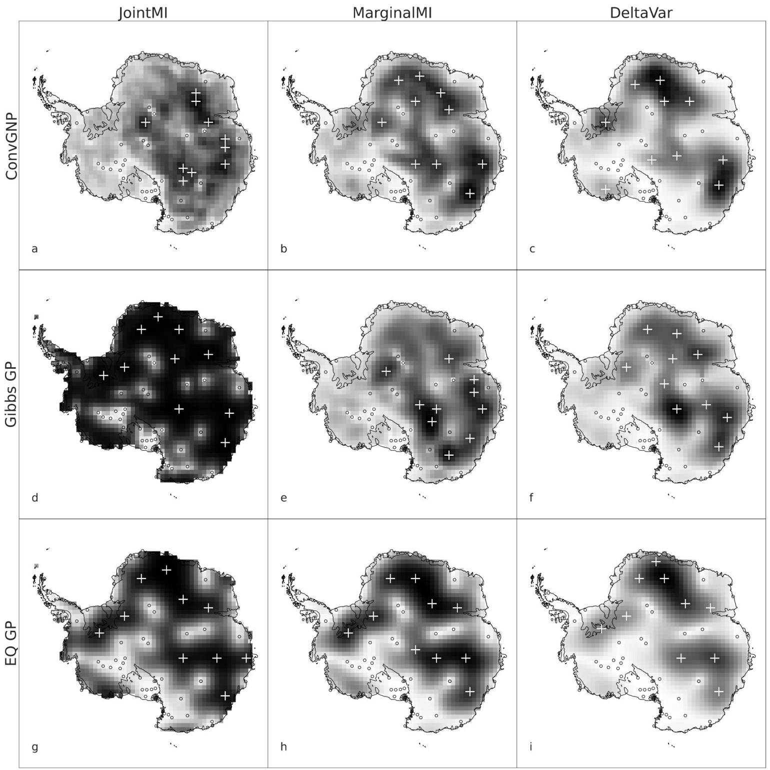

Comparison/integration of ML and physics-based observing system design

Numerical modelling approach	**Machine learning analogue(?)**
Observing system experiments (OSEs)	Variable ablation interpretability techniques
Adjoint modelling	Saliency analysis using backpropagation
Ensemble sensitivity analysis (ESA)	Uncertainty-based active learning (our work)

Neural process timeline

- 2018: M. Garnelo et al., “Conditional Neural Processes.” *ICML*
- 2020: J. Gordon et al., “Convolutional Conditional Neural Processes.” *ICLR*
- 2021: W.P. Bruinsma et al., “The Gaussian Neural Process.” *AABI*
- 2022: S. Markou et al., “Practical Conditional Neural Processes via Tractable Dependent Predictions.” *ICLR*
- 2022: A. Vaughan et al., “Convolutional Conditional Neural Processes for Local Climate Downscaling.” *GMD*
- 2023: T.R. Andersson et al., “Environmental Sensor Placement with Convolutional Gaussian Neural Processes.” *EDS, in review* (our work)