ICML
International Conference
On Machine Learning

Localized, High-Resolution Geographic Representations with Slepian Functions

Arjun Rao Ruth Crasto Tessa Ooms David Rolnick Konstantin Klemmer Marc Rußwurm

University of Colorado Boulder
Microsoft
Mila
UCL
WAGENINGEN UNIVERSITY & RESEARCH
UNI BONN
LEGEND

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML
International Conference
On Machine Learning

Geographic prediction requires functions defined on the Earth's surface.

Geographic location encoders

$$f(x) = \mathrm{NN}(\Phi(x)), \qquad x = (\lambda, \phi) \in S^2$$

Positional encoding/basis function

Geographic location encoders

$$f(x) = \mathrm{NN}(\Phi(x)), \qquad x = (\lambda, \phi) \in S^2$$

Spherical Harmonics are the canonical basis on S^2

$$\mathcal{H}_L = \mathrm{span}\{Y_\ell^m : 0 \le \ell \le L,\ -\ell \le m \le \ell\}$$

$$\Phi_{\mathrm{SH}}(x) = \left[Y_\ell^m(x) : 0 \le \ell \le L,\ -\ell \le m \le \ell\right] \in \mathbb{R}^{D_L}$$

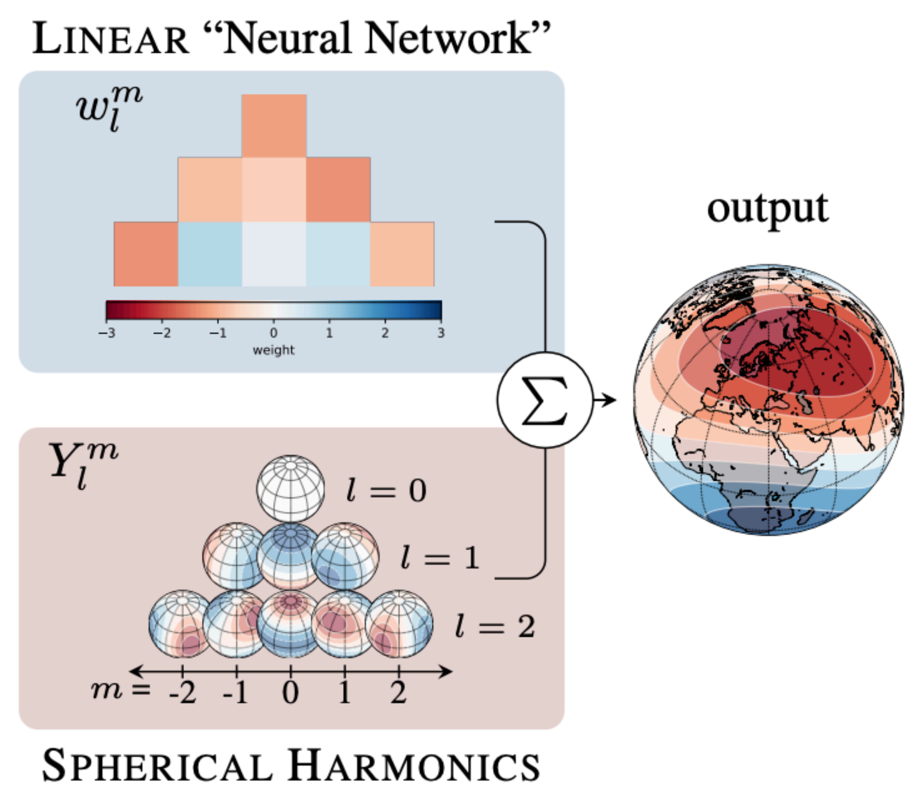

Rußwurm, Marc, et al. "Geographic Location Encoding with Spherical Harmonics and Sinusoidal Representation Networks," ICLR 2024

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML
International Conference
On Machine Learning

Resolving local detail requires increasing global resolution

$$D_L = (L+1)^2$$

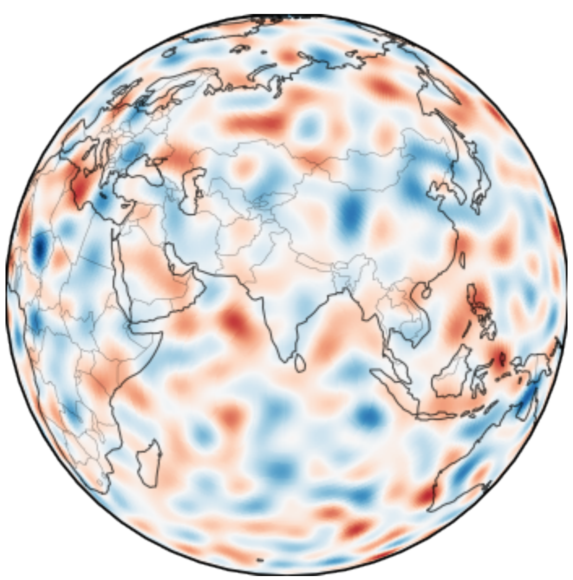

Spherical Harmonics

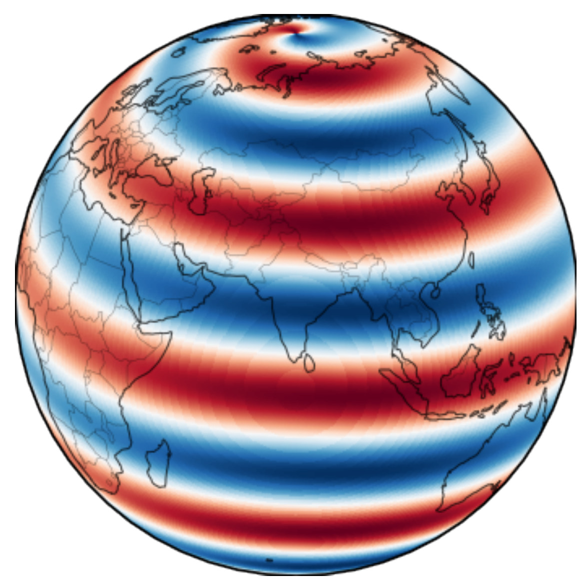

SphereM

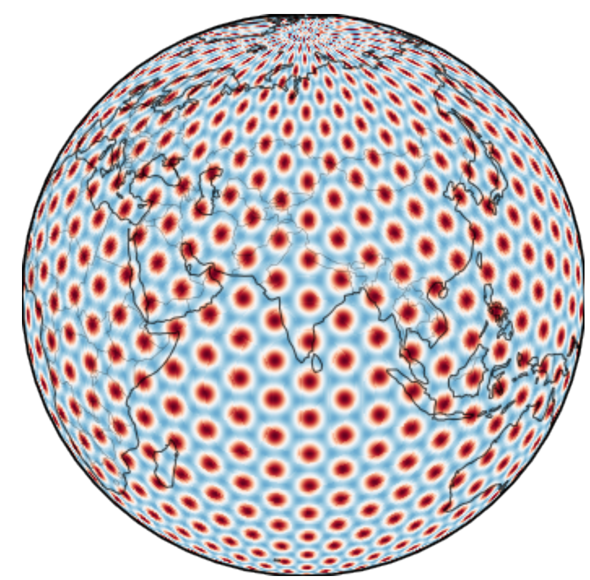

Theory

Mai, Gengchen, et al. "Sphere2Vec: A General-Purpose Location Representation Learning over a Spherical Surface for Large-Scale Geospatial Predictions ," ISPRS Photogrammetry and Remote Sensing

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML
International Conference
On Machine Learning

Problem: Geographic prediction within highly localized regions

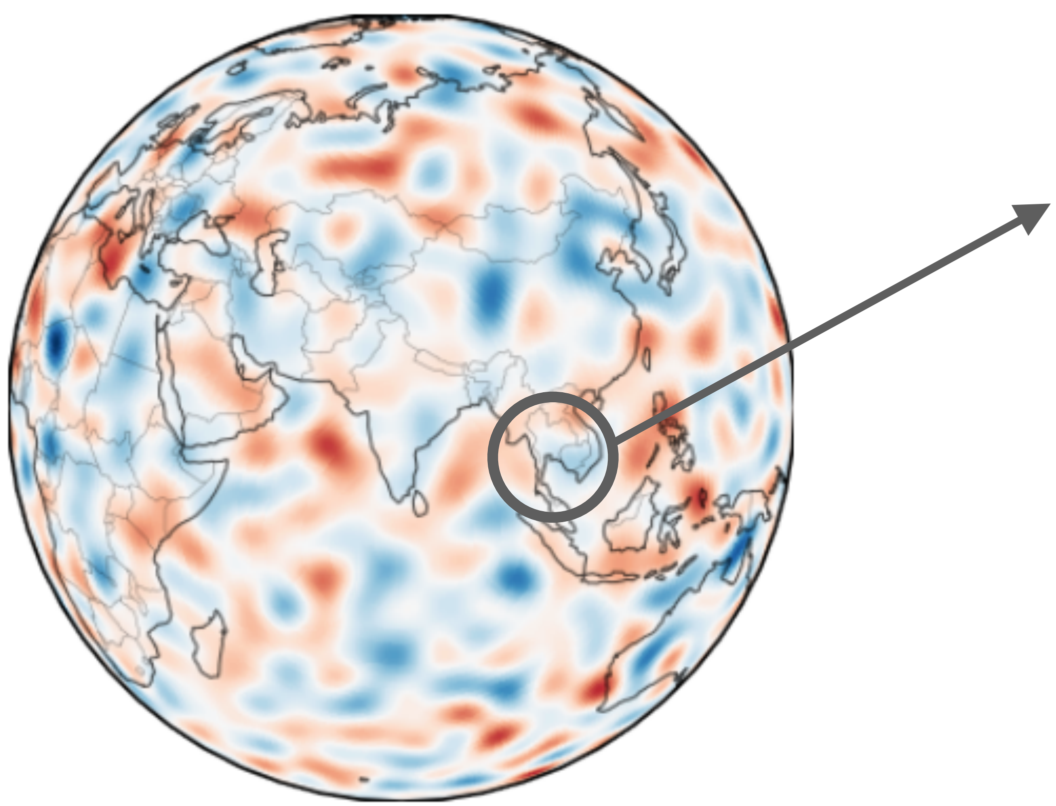

What if we only care about doing well here?

Spherical Harmonics allocates “representational capacity” uniformly and poorly scales to high resolutions demanded by localized spatial prediction tasks.

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

Given a region R on the sphere, find band-limited functions that concentrate their energy in R.

$$\mu = \frac{\int_R |h(\mathbf{x})|^2 \, ds(\mathbf{x})}{\int_{S^2} |h(\mathbf{x})|^2 \, ds(\mathbf{x})} \quad \in [0,1]$$

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

Given a region R on the sphere, find band-limited functions that concentrate their energy in R.

$$\mu = \frac{\int_R |h(\mathbf{x})|^2 \, ds(\mathbf{x})}{\int_{S^2} |h(\mathbf{x})|^2 \, ds(\mathbf{x})} \in [0,1] \qquad \frac{\text{Fraction of energy inside } R}{\text{Total energy on } S^2}$$

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

Given a region R on the sphere, find band-limited functions that concentrate their energy in R.

Energy concentration within R

$$\mu = \frac{\int_R |h(\mathbf{x})|^2 \, ds(\mathbf{x})}{\int_{S^2} |h(\mathbf{x})|^2 \, ds(\mathbf{x})} \in [0,1]$$

Candidate band-limited function

Maximizing this ratio yields a concentration eigenvalue problem.

$$\mathbf{Kh} = \mu\mathbf{h}, \qquad K_{\ell m,\ell' m'} = \int_R Y_\ell^m(\mathbf{x}) \, Y_{\ell'}^{m'}(\mathbf{x}) \, ds(\mathbf{x})$$

Localized, High-Resolution Geographic Representations with Slepian Functions

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML International Conference On Machine Learning

Shannon Number: $N = (L+1)^2 \times (1-\cos\theta)/2$

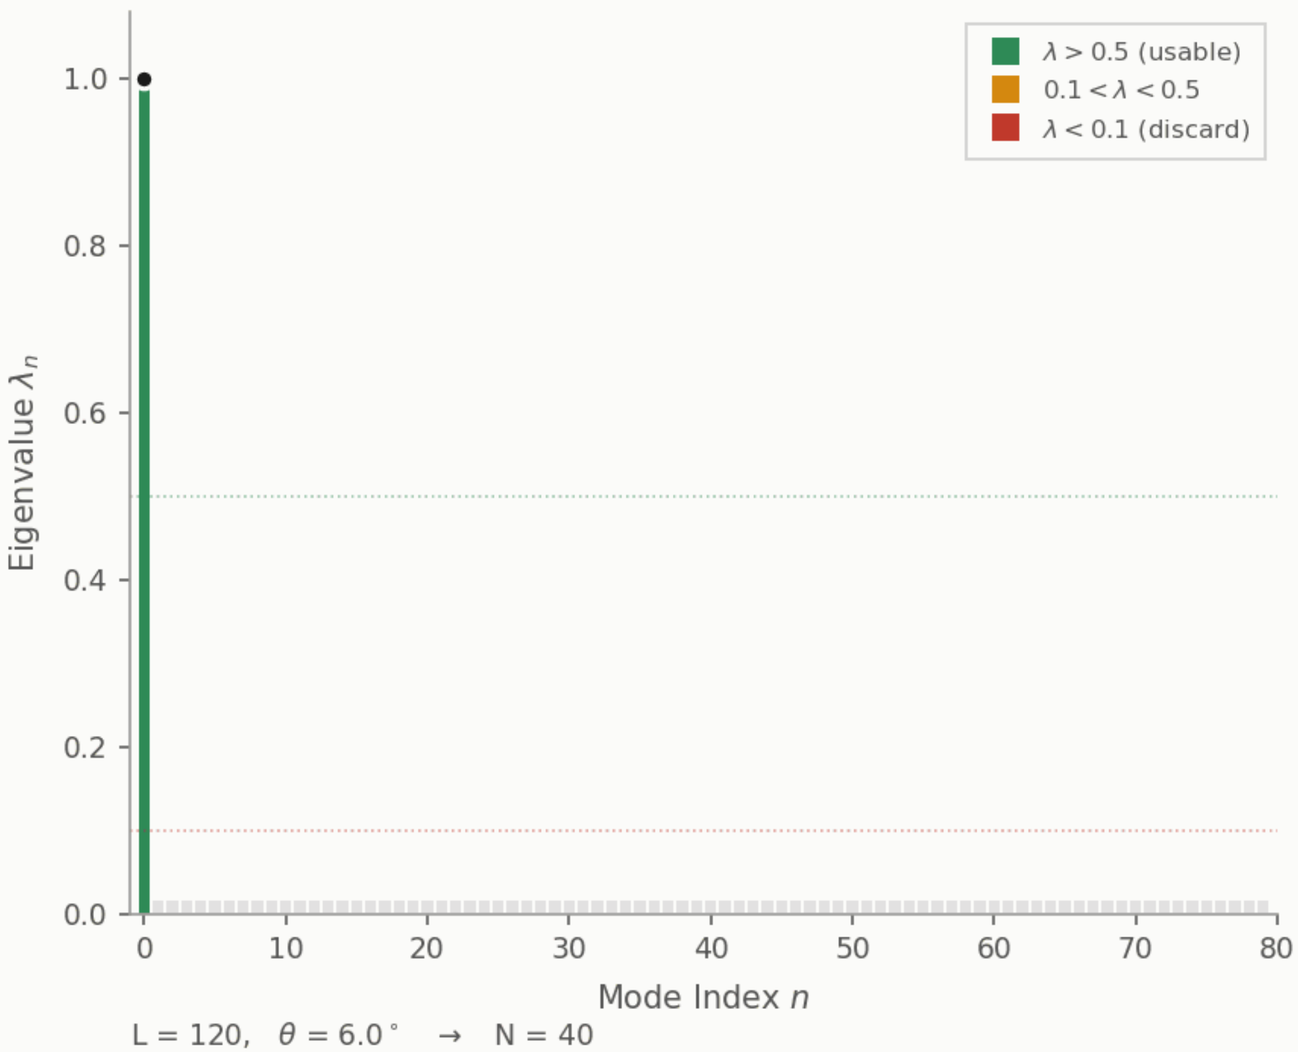

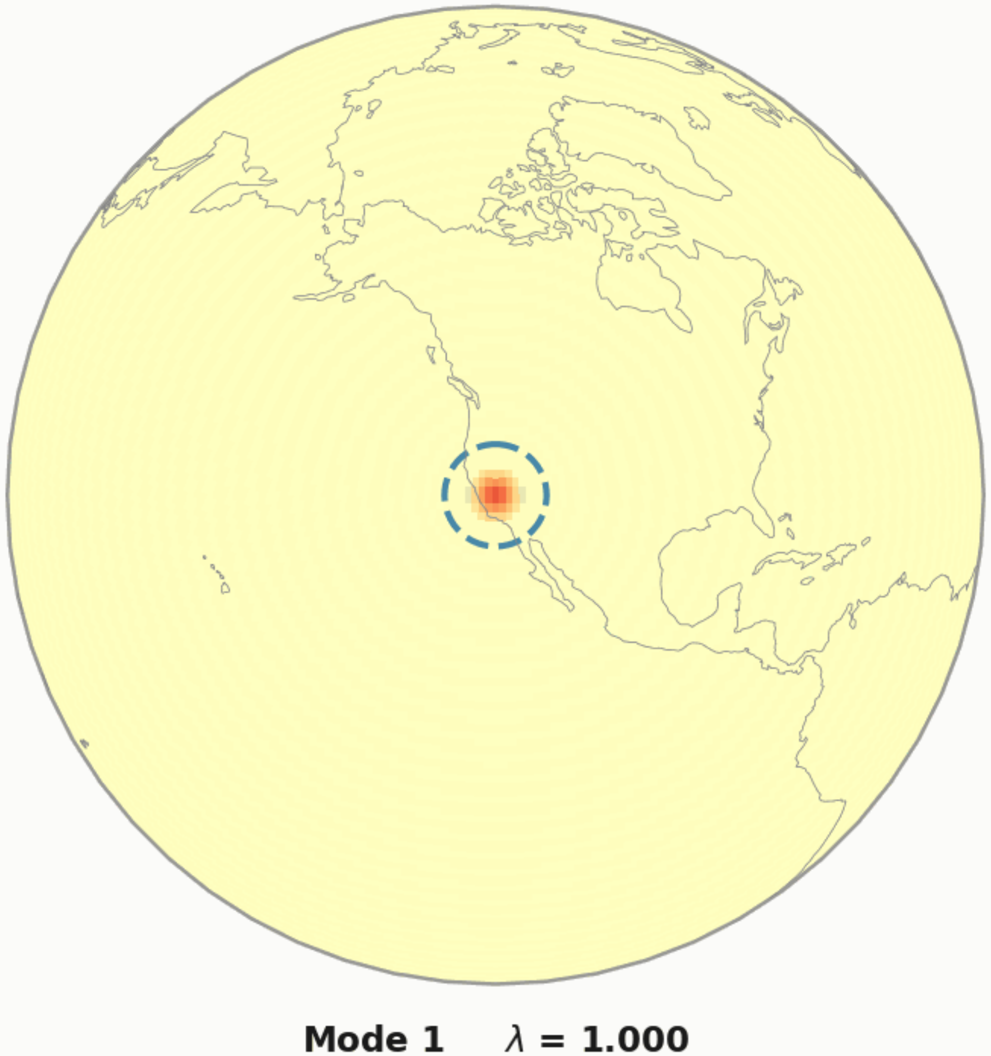

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML
International Conference
On Machine Learning

Hybrid Slepian-SH location encoding captures both local and global geographic context.

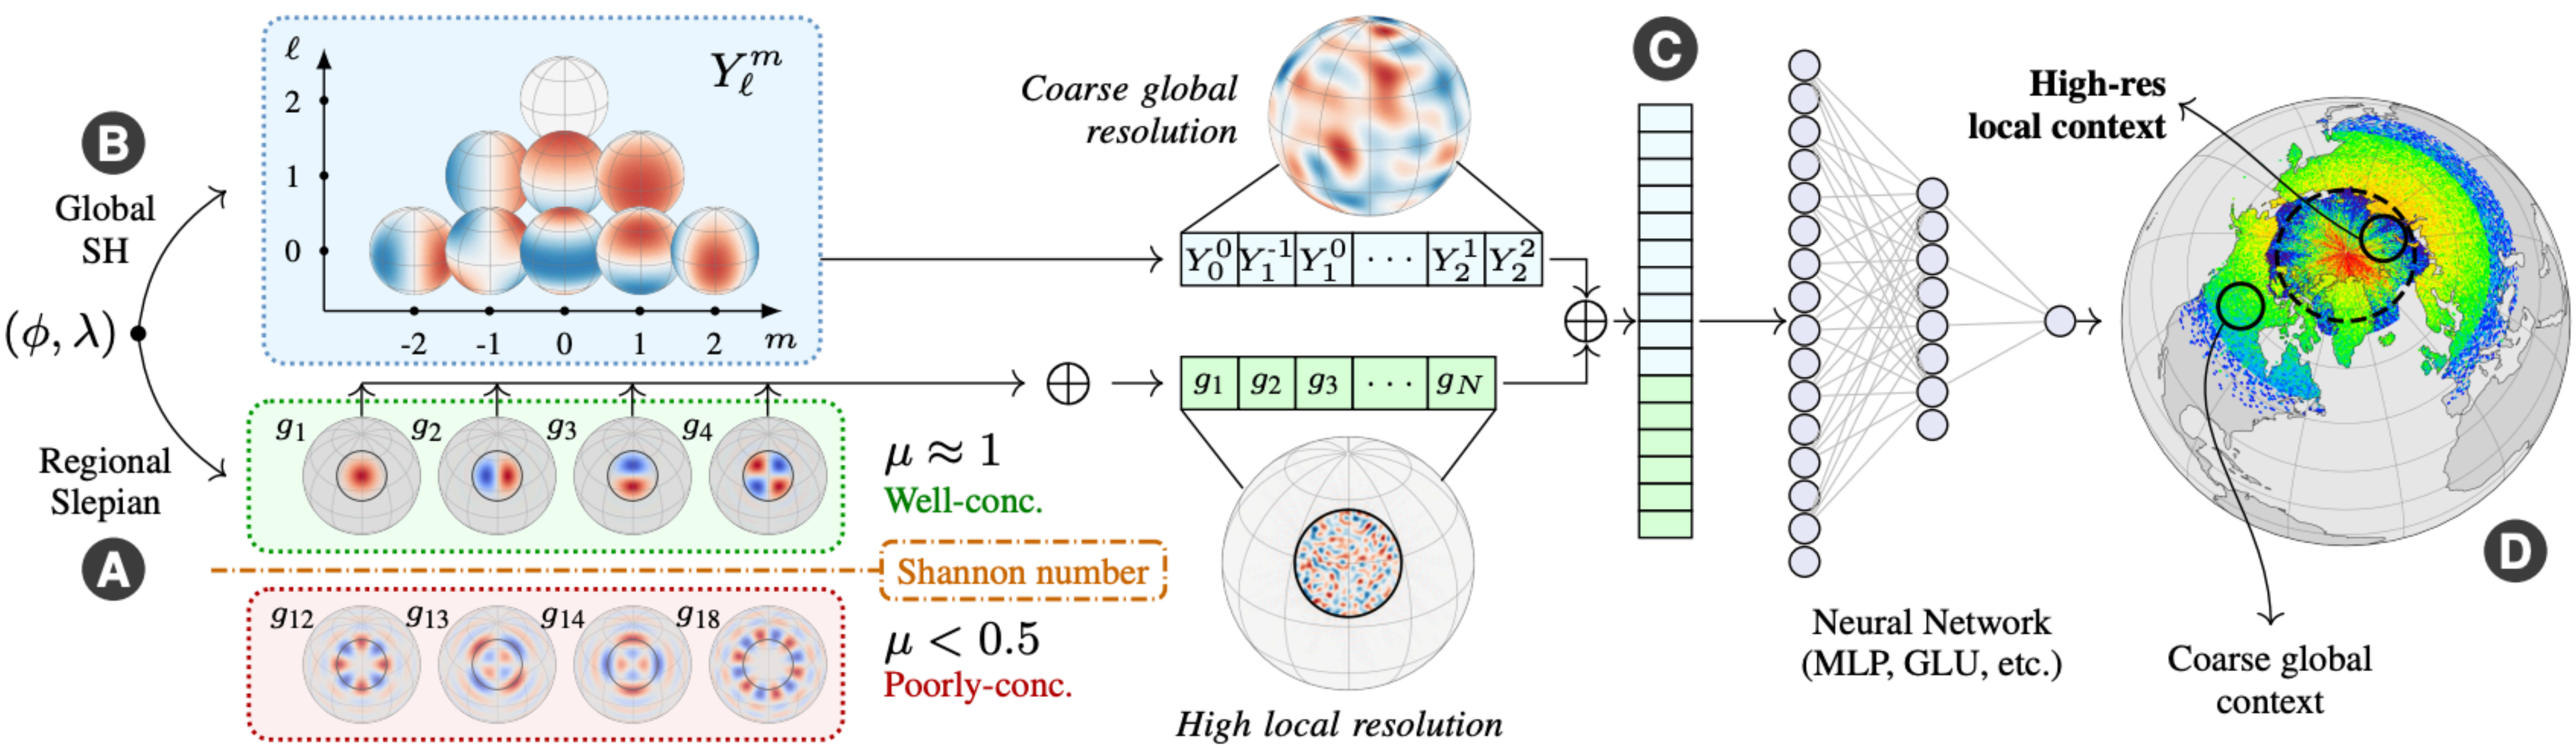

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML
International Conference On Machine Learning

Finding 1: Slepian functions preserve high-resolution detail for spatial prediction tasks.

Finding 2: A Hybrid Slepian-SH encoder resolves the local-global tradeoff.

Finding 3: Slepian encodings are computationally efficient at high-resolutions.

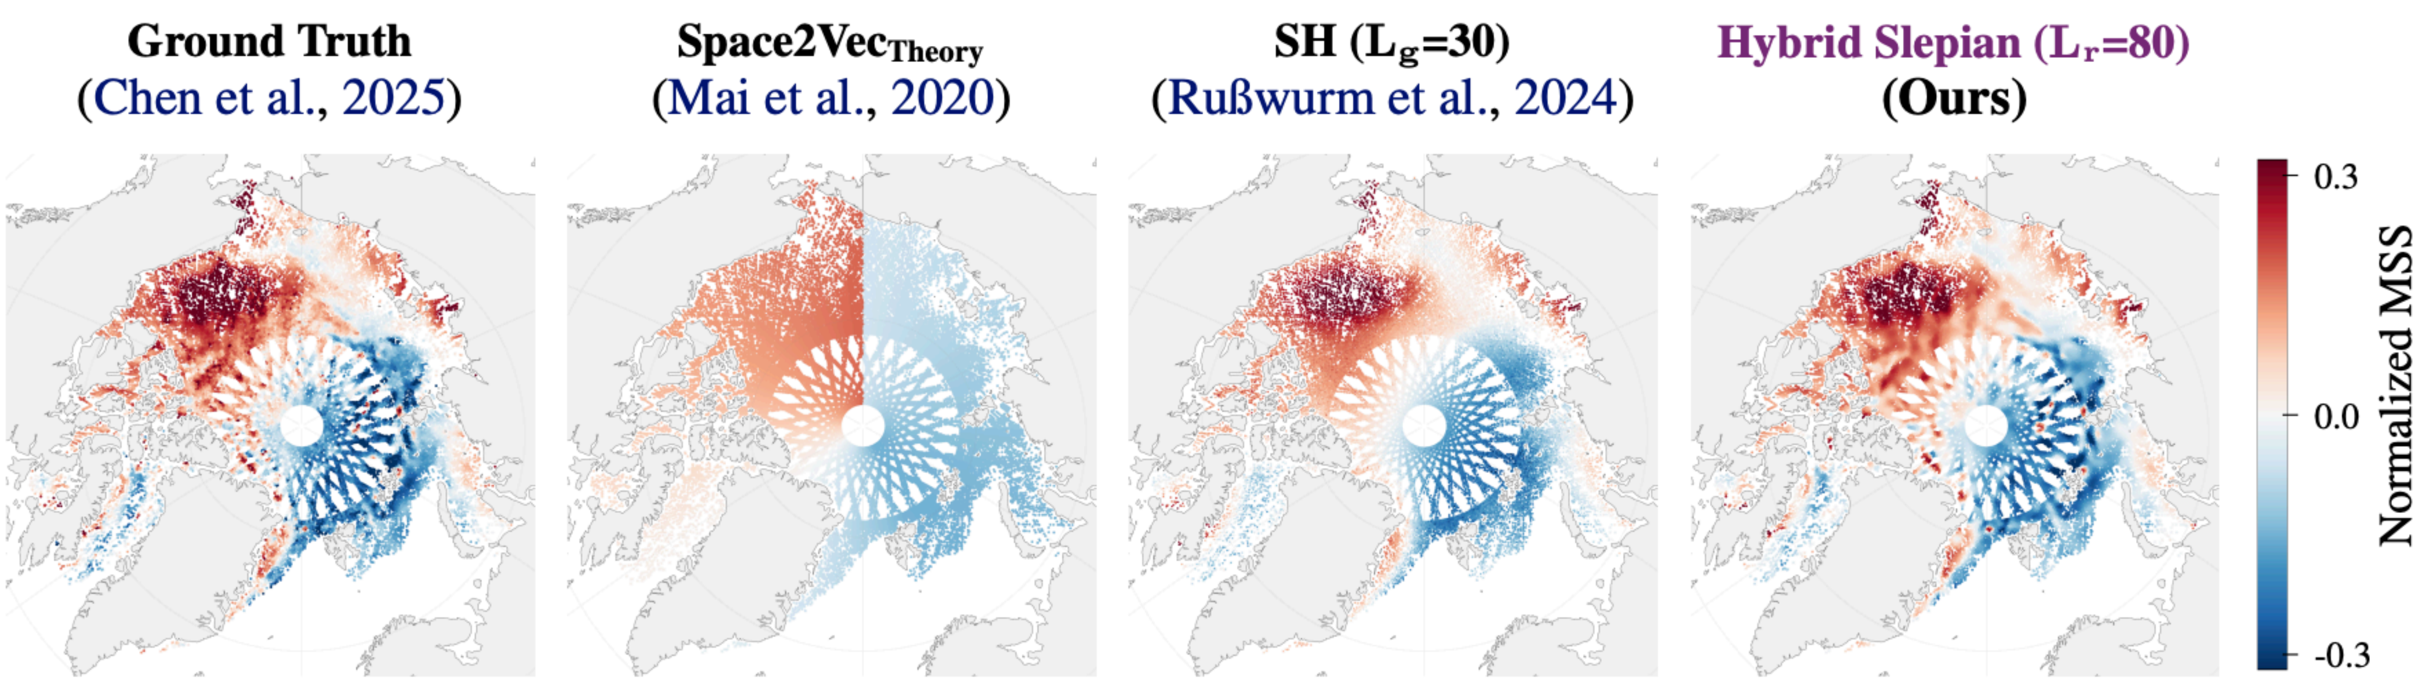

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML
International Conference On Machine Learning

Finding 1: Slepian functions preserve high-resolution detail for spatial prediction tasks.

Finding 2: A Hybrid Slepian-SH encoder resolves the local-global tradeoff.

Finding 3: Slepian encodings are computationally efficient at high-resolutions.

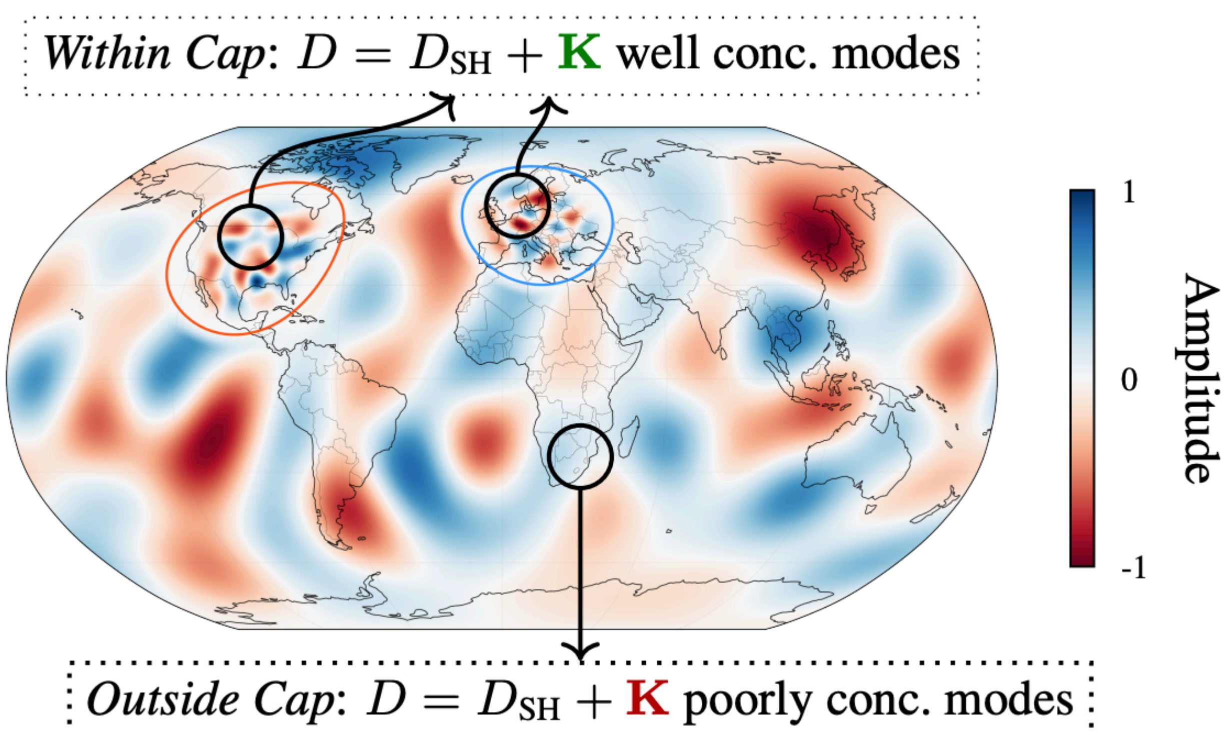

Encoding	Config	S&T Birds ↑	IUCN ↑
Wrap	–	0.250	0.009
SH	$L_g = 10$	0.637	0.342
Slepian only	$L_g = 0,\ L_r = 40$	0.298	0.013
Hybrid Slepian	$L_g = 10$, best L_r	**0.704**	**0.479**

LinNet setting. Best L_r selected per benchmark.

Cole, Elijah, et al. “Spatial Implicit Neural Representations for Global-Scale Species Mapping,” ICML 2023

Localized, High-Resolution Geographic Representations with Slepian Functions

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

ICML International Conference On Machine Learning

Finding 1: Slepian functions preserve high-resolution detail for spatial prediction tasks.

Finding 2: A Hybrid Slepian-SH encoder resolves the local-global tradeoff.

Finding 3: Slepian encodings are computationally efficient at high-resolutions.

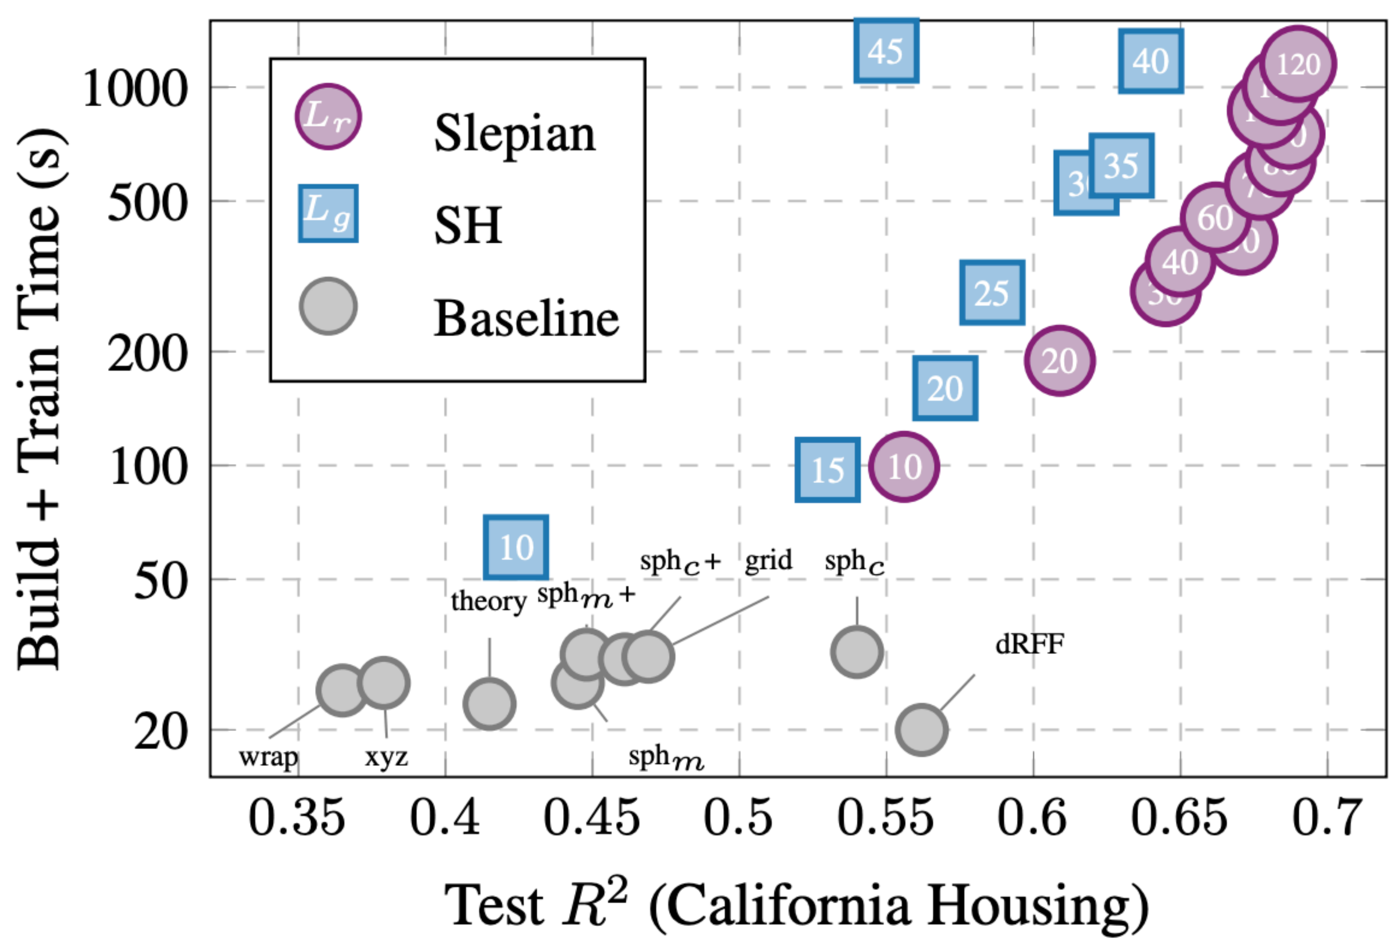

Localized, High-Resolution Geographic Representations with Slepian Functions

Arjun Rao, Ruth Crasto, Tessa Ooms, David Rolnick, Konstantin Klemmer, Marc Rußwurm

TL;DR Takeaways

1 This paper proposes a fairly simple way to learn high-resolution, local geographic context along with a coarse global context.

2 It does it in a way that's computationally feasible.

arXiv: 2602.00392

Paper

Code

University of Colorado Boulder Microsoft Mila UCL WAGENINGEN UNIVERSITY & RESEARCH UNI BONN LEGEND