Generalised Implicit Neural Representations

Daniele Grattarola, Pierre Vandergheynst

Neural Information Processing Systems 2022

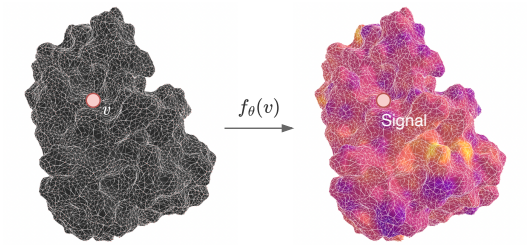

Implicit neural representations

Given a signal or field $f : \mathbb{R}^d \to \mathbb{R}^p$

Sample a regular grid $\mathbf{y}_i = f(\mathbf{x}_i)$

Train neural network $f_\theta : \mathbf{x}_i \mapsto \mathbf{y}_i$

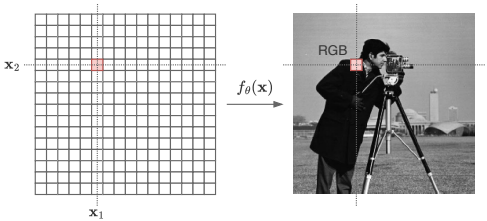

Generalised implicit neural representations [1]

Given a signal or field $f : \mathcal{T} \to \mathbb{R}^p$

Sample a graph signal $\mathbf{y}_i = f(v_i)$

Train neural network $f_\theta : v_i \mapsto \mathbf{y}_i$

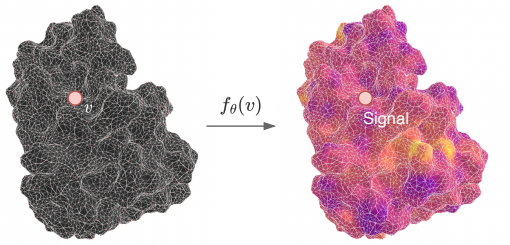

[1] D. Grattarola *et al.*, "Generalised implicit neural representations," *Advances in Neural Information Processing Systems*, 2022.

Generalised implicit neural representations [1]

Use spectral positional encodings as coordinates:

$$\mathbf{e}_i = \sqrt{n}\ \underbrace{\left[\mathbf{u}_{1,i}, \ldots, \mathbf{u}_{k,i}\right]^\top}_{\text{Laplacian eigenvectors}} \in \mathbb{R}^k$$

LB eigenfunctions for $n \to \infty$.

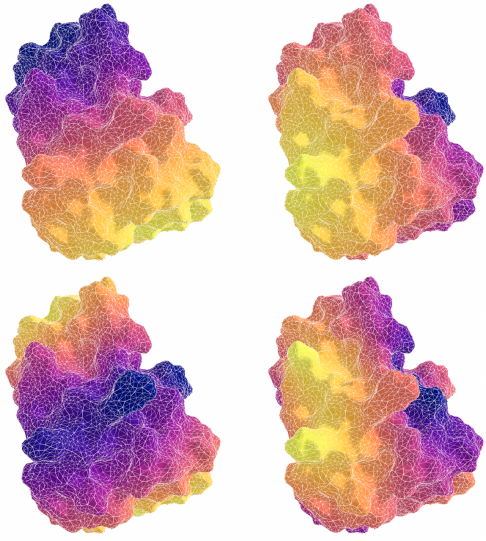

[1] D. Grattarola *et al.*, "Generalised implicit neural representations," *Advances in Neural Information Processing Systems*, 2022.

Experiments

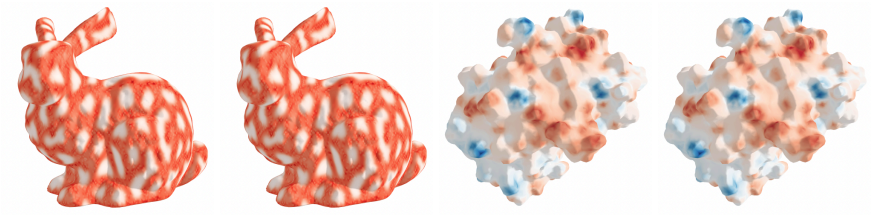

Ground truth signals *vs.* signals predicted by the GINR.

	Bunny	Protein	US Election
R^2	1.000	1.000	0.999
MSE	$9.14 \cdot 10^{-8}$	$1.17 \cdot 10^{-10}$	$1.45 \cdot 10^{-3}$

Approximation error

Experiments - Size of embeddings

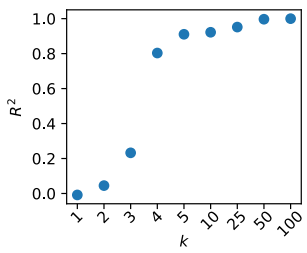

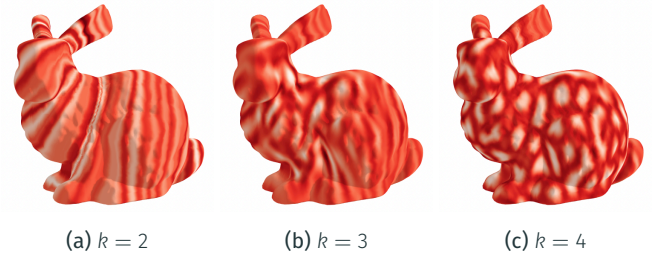

(a) $k = 2$ (b) $k = 3$ (c) $k = 4$

Left: R^2 *vs.* k; **Right:** signals learned by the INR for $k = 2, 3, 4$.

Experiments - Transferability

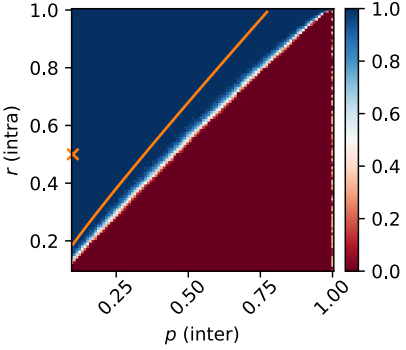

Toy problem with SBM.

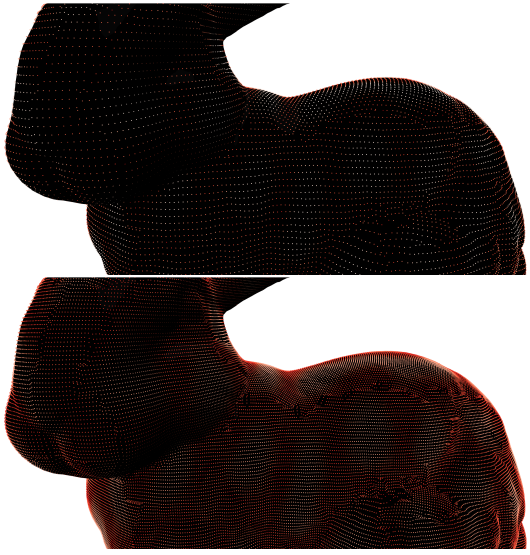

Experiments - Conditional GINRs

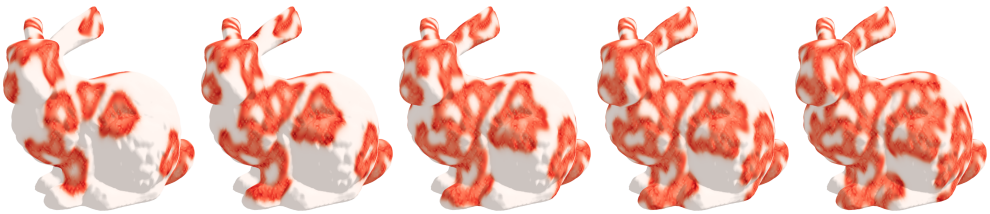

Signals predicted by the conditional GINR $f_\theta(\mathbf{e}_i, t)$

Experiments - Conditional GINRs

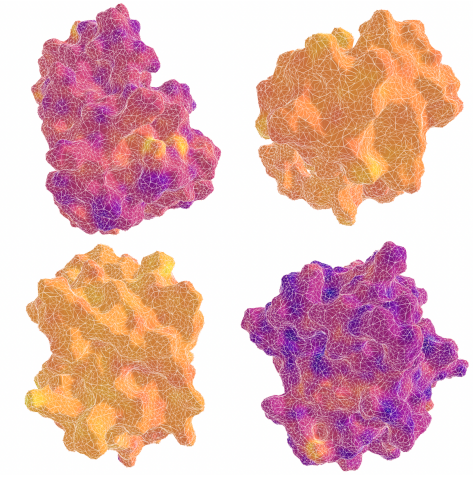

$f_\theta(\mathbf{e}_i, \mathbf{z}_d)$ for node i and protein d.

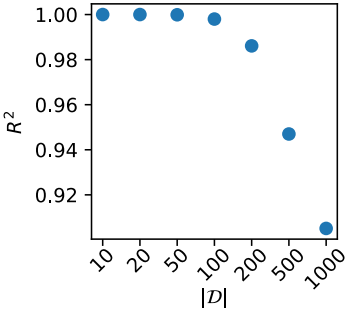

Good performance up to 100s of proteins.

Experiments - Solving differential equations

(a) Laplacian (supervision, Δf_θ)

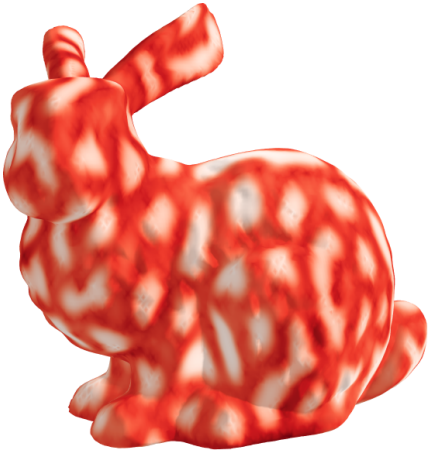

(b) Reconstructed (f_θ)

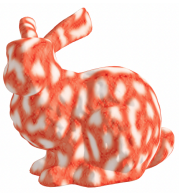

(c) Original (f)

Experiments - Weather modelling

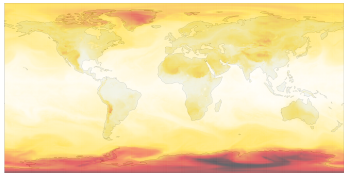

(a) Temperature

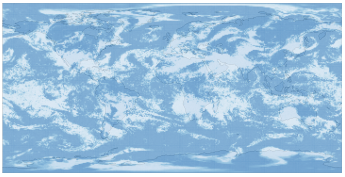

(b) Clouds

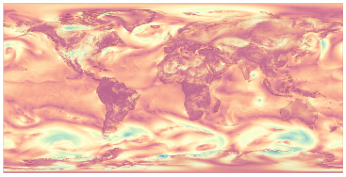

(c) Wind

Data from the National Oceanic and Atmospheric Administration (NOAA).

Conclusion

EPFL

Generalised Implicit Neural Representations

Daniele Grattarola, Pierre Vandergheynst

Neural Information Processing Systems 2022

github.com/danielegrattarola/GINR

arxiv.org/abs/2205.15674

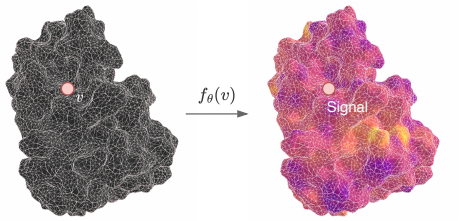

Get in touch: daniele.grattarola@gmail.com or @riceasphait

References i

[1] D. Grattarola and P. Vandergheynst, “Generalised implicit neural representations,” *Advances in Neural Information Processing Systems*, 2022.