

Beyond Uniformity: Regularizing Implicit Neural Representations through a Lipschitz Lens

Julian McGinnis^{1,2,*}, Suprosanna Shit^{3,4,*}, Florian A. Hölzl^{1,5}, Paul Friedrich⁶, Paul Büschl³,
Vasiliki Sideri-Lampretsa¹, Mark Mühlau¹, Philippe C. Cattin⁶, Björn Menze³,
Daniel Rueckert^{1,2,7,†}, Benedikt Wiestler^{1,2,†}


* Equal contribution † Equal senior contribution

¹Technical University of Munich ²MCML ³University of Zurich

⁴ETH AI Centre ⁵Hasso-Plattner-Institute ⁶University of Basel ⁷Imperial College London

Motivation


Implicit Neural Representations (INRs) have emerged as a promising, modality-agnostic modeling framework for inverse problems.



Notably, recent studies [1,2] have revealed that generalization and interpolation properties are often inferior to grid representations, which may be intuitively explained by ***the absence of implicit regularization*** [3].


\rightarrow To mitigate this, we propose ***Lipschitz regularization in the generalized K-Lipschitz setting***.

How is the total Lipschitz bound for a Neural Net derived?



- (i) Lipschitz bounds induced by embeddings (e.g. Fourier Features)
- (ii) Lipschitz bounds induced by activation (e.g. SIREN)
- (iii) Lipschitz bounds induced by maximum singular value of each layer's weights


1-Lipschitz INRs: Expressiveness and Lipschitz Capacity Utilization



→ The utilization of Lipschitz capacity (= 1 is ideal) depends on the effectiveness of gradient-perserving activation functions (e.g. MaxMin, Householder) and spectral regularization methods (e.g. Björck and SLL).



How to move from 1-Lipschitz to K-Lipschitz settings? Budget Estimation

Find strategies to estimate a reasonable, meaningful budget K for different tasks, such as deformable lung registration, where we use **(i) signal-theoretic** or **(ii) domain-driven estimates**:

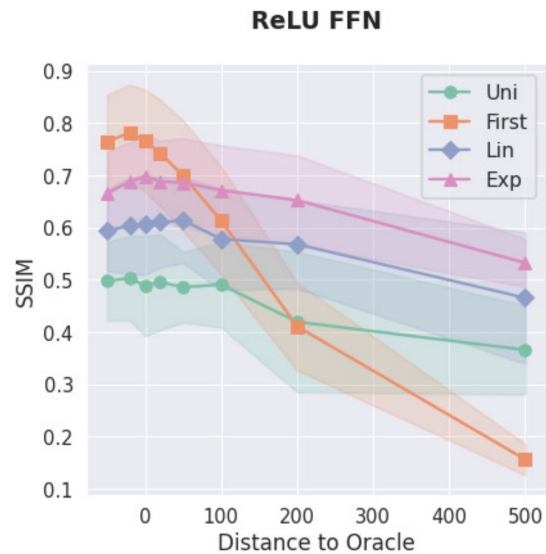
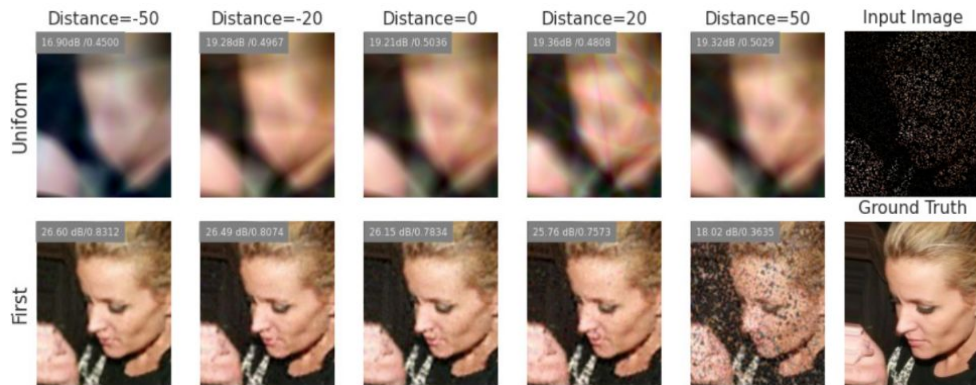


Clinical evidence suggests that a strain approaching two marks a threshold for tissue failure [5], motivating our choice of $K=2$.

Effect of Budget Estimation vs Budget Allocation



Effect of Budget Estimation vs Budget Allocation



Key Takeaways

(1) **1-Lipschitz Setting:**

- Choice of architecture is the main driver of reconstruction quality.
- In particular, the combination of **normalization** and **activation function** matters more than the budget allocation strategy.

(2) **K-Lipschitz Setting:**

- Data driven estimation shows the importance of **accurate budget estimation** for inverse problems.
- For the same architecture and budget K , **non-uniform allocation** consistently outperforms uniform allocation, making allocation strategy a strategic and important factor.

Future work will examine **how we can make both, budget estimation and allocation, learnable.**

For more details, please checkout [our page](#) and [code](#).



Citations

- [1] Fridovich-Keil et al. *Plenoxels: Radiance fields without neural networks*, CVPR 2022
- [2] Kim et al. *Grids Often Outperform Implicit Neural Representation at Compressing Dense Signals*, NeurIPS 2025
- [3] Ramasinghe et al. *On the frequency-bias of coordinate-mlps*, NeurIPS 2022
- [4] Dupont et al. *From data to functa: Your data point is a function and you can treat it like one*, ICML 2022
- [5] Chiumello et al. *Lung stress and strain during mechanical ventilation for acute respiratory distress syndrome*. American journal of respiratory critical care medicine, 2008.