Scale-Agnostic Super-Resolution in MRI using Feature-Based Coordinate Networks

Dave Van Veen, Rogier van der Sluijs, Batu Ozturkler, Arjun Desai,
Christian Bluethgen, Robert D. Boutin, Marc H. Willis, Gordon Wetzstein,
David Lindell, Shreyas Vasanawala, John Pauly, Akshay S. Chaudhari

Abstract

Deep learning has shown great promise to encode high-frequency information for the task of super-resolution in MRI. However, typical methods require upsampling at fixed, discrete scales due to their convolutional structure. This is undesirable for clinical interpretation and places limits on acquiring homogeneous training data [1].

We propose a *scale-agnostic* framework for MRI super-resolution, using a coordinate network as a decoder. The continuous nature of this decoder enables:

- Querying at arbitrary resolutions
- Decoupling between training and querying scales

We compare this (*coord*) to a similar framework with a convolutional decoder (*conv*) and evaluate performance in the context of various denoising strategies.

Background

Coordinate networks (**left**) are recent powerful tools for representing signals with fully-connected multi-layer perceptrons by mapping signal coordinates to their corresponding pixel values. This is continuous w.r.t. network weights—in contrast to standard pixel-based representations, which operate on discrete grids (**right**).

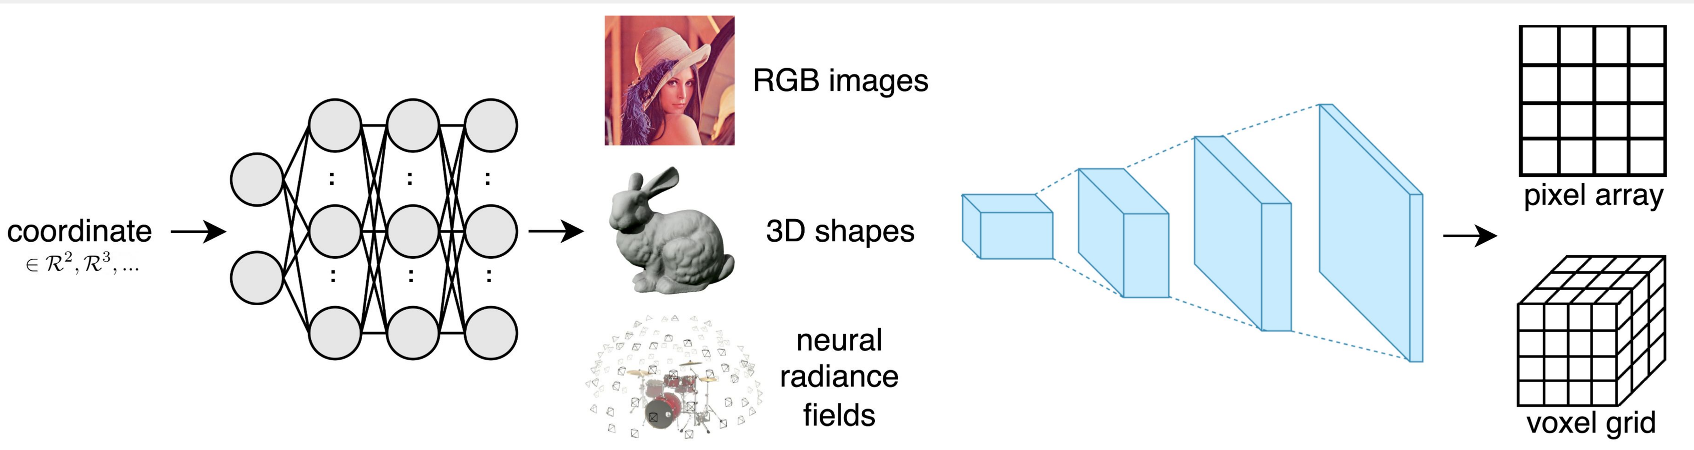

Methods

Our training pipeline (**bottom**) maps a low-resolution input to a latent representation, which is subsequently queried at arbitrary resolution using a coordinate network. The continuous nature of the decoder enables this framework to be scale-agnostic, e.g. we could train on a range of scales 1-2× and query at 1.8×.

We notice *coord* outperforms *conv* at representing higher frequencies; consequently it is more apt to represent noise. As such we incorporate both early stopping and denoising regularization.

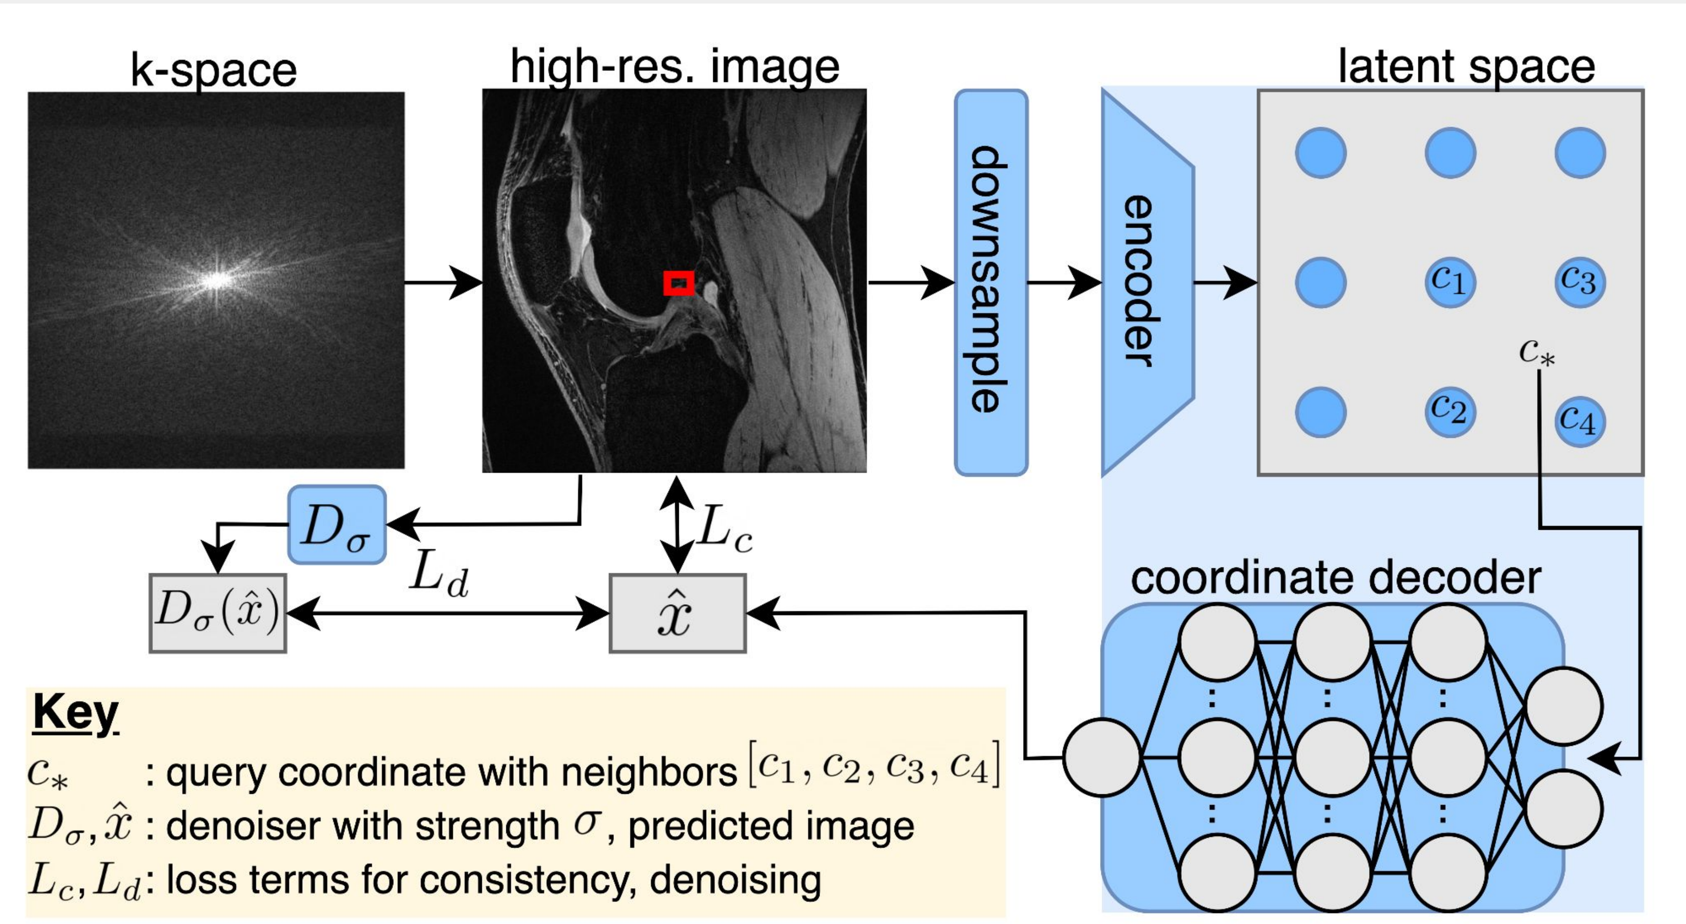

Results

Top: Evaluations of *coord* and *conv*, each trained at 2×. *Coord* can also be queried at 4× without re-training because it is scale-agnostic (F). *Coord* benefits from denoising ($\lambda \neq 0$, D vs. E), while *conv* does not.

Bottom: In addition to being scale-agnostic, *coord* obtains similar performance (higher VIF and lower PSNR), the former of which is more indicative of super-resolution and diagnostic quality [2].

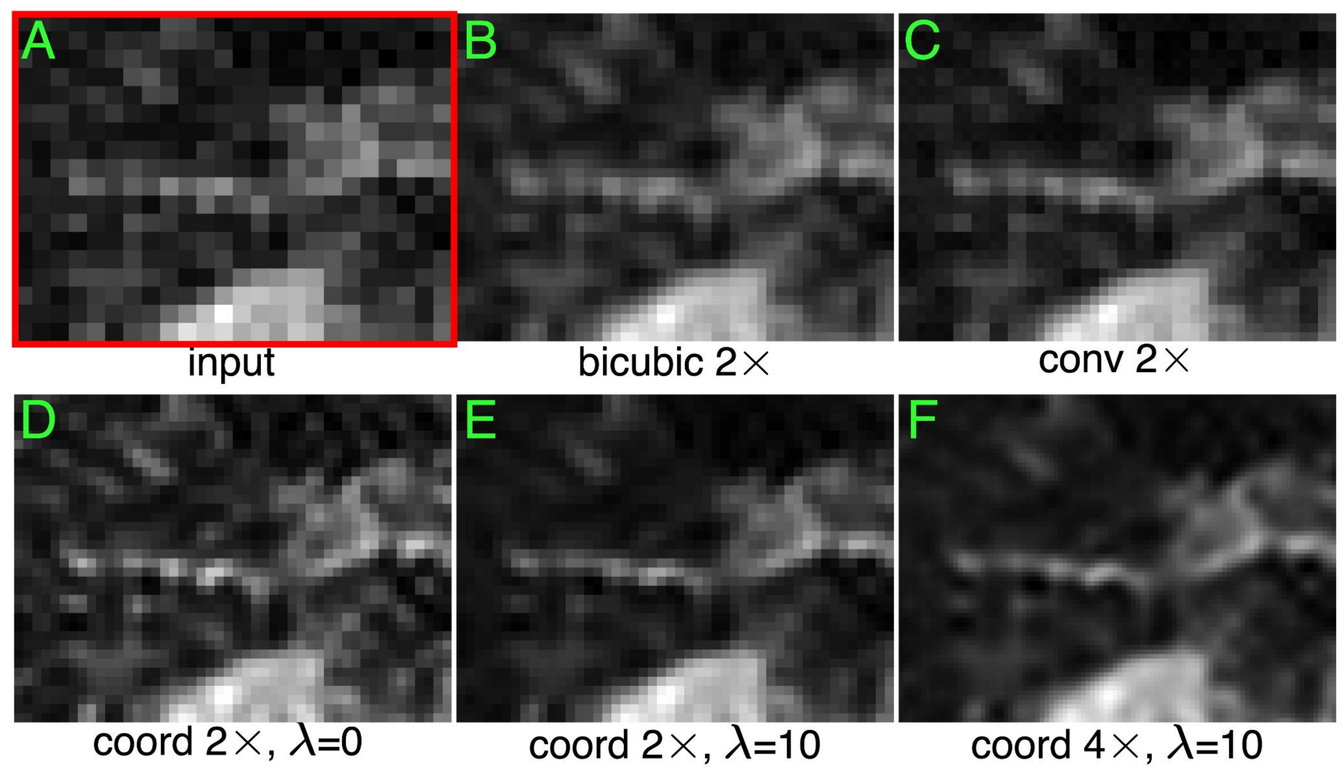

	Method	λ	Scale-agnostic	Train scale	Query scale (VIF/PSNR) 1.5×	2×	3×
Default	Bicubic	—	✓	—	.87/34.2	.79/31.4	.64/27.6
Default	Coord	10	✓	1-2×	.92/33.8	.87/31.0	.76/26.9
Default	Conv	0		2×	—	.82/32.3	—
Ablation	Coord	0	✓	1-2×	.95/33.0	.88/30.2	.74/26.3
Ablation	Coord	10	✓	2×	.95/33.2	.89/30.6	.77/26.2
Ablation	Conv	10		2×	—	.81/32.2	—

Reader Study

To enable clinically-relevant evaluation at scales beyond ground-truth, radiologists scored images using a five-point Likert scale (**right**). This demonstrates a slight preference for *coord* in terms of noise and sharpness. The study was built with Voxel: an emerging, web-based viewer for medical imaging tensors (**bottom**).

Compared to *conv* in [noise/sharpness], *coord* is...

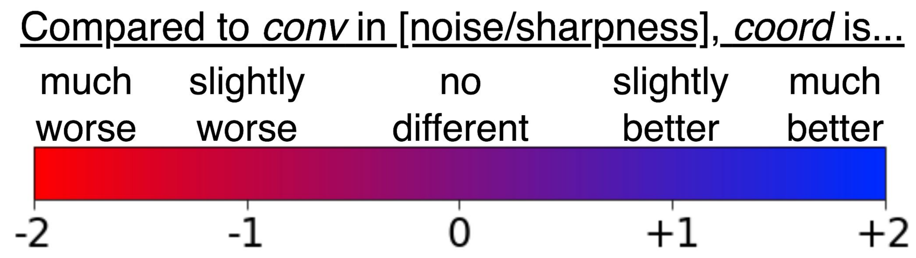

Reader	Noise	Sharpness
1	.70 ± .64	1.4 ± .69
2	.78 ± .46	.02 ± .14
Pooled	.74 ± .56	.70 ± .84

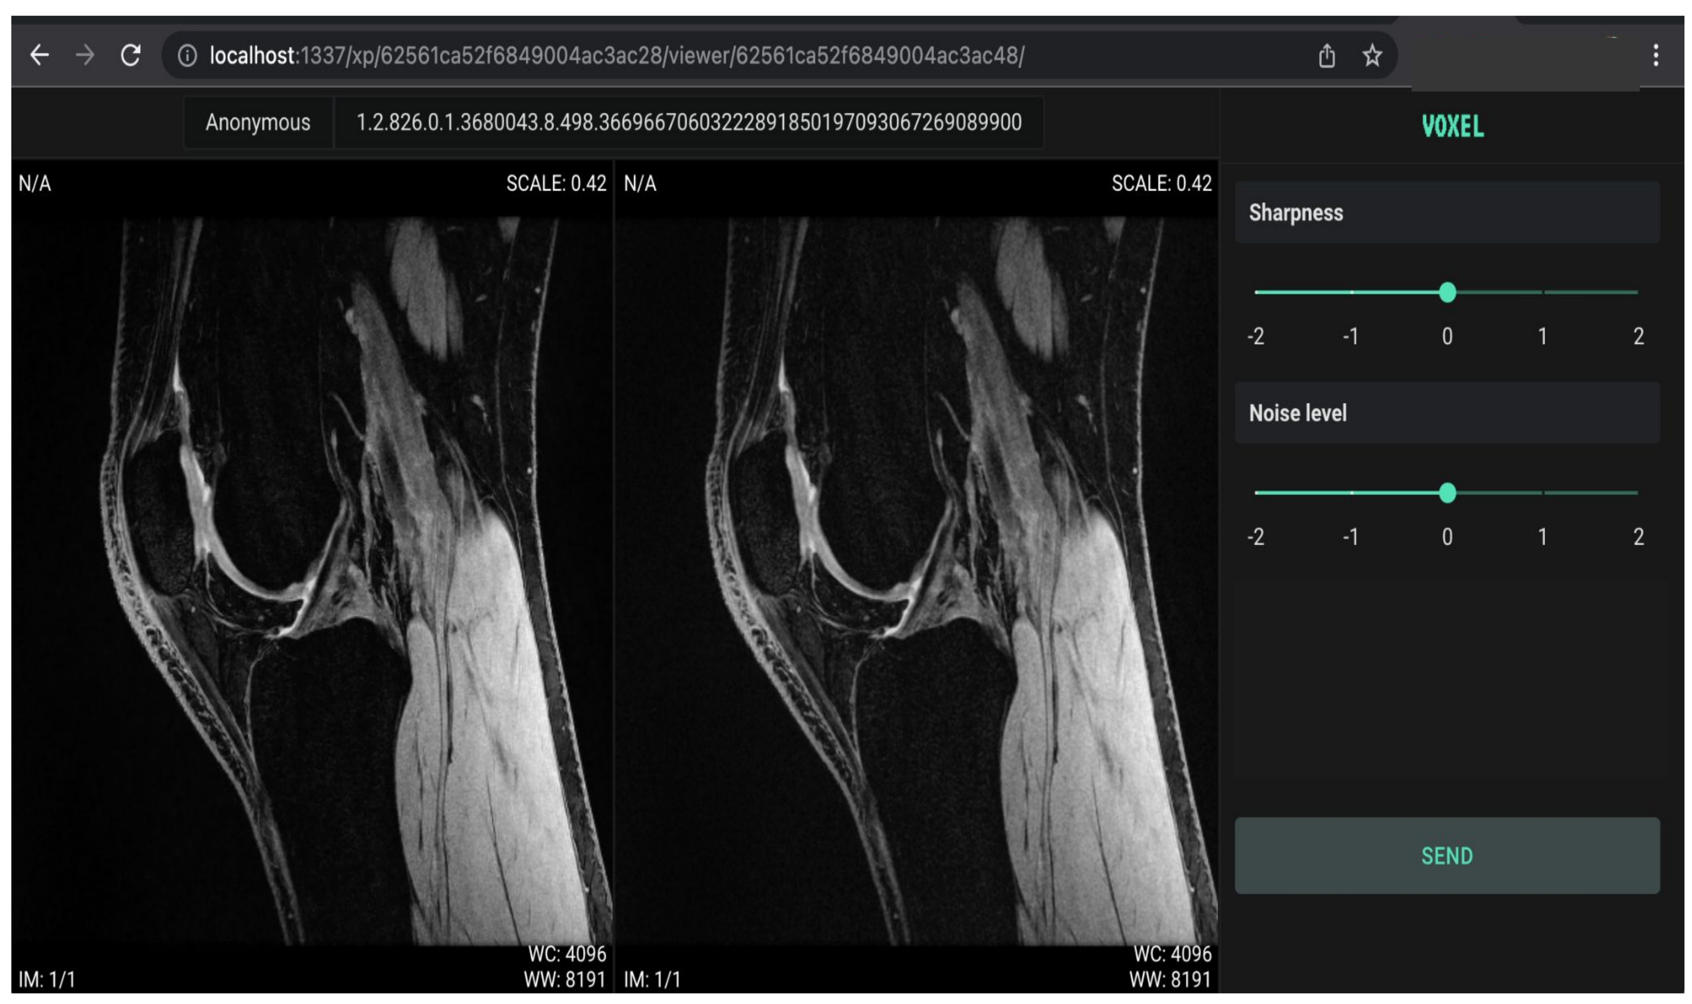

References

[1] Chaudhari et al., Prospective deployment of deep learning in MRI: best practices. *JMRI*, 2021.
[2] Mason et al., Comparison of image metrics to radiologists' scoring of MR diagnostic quality. *IEEE - TMI*, 2019.

Contact: Dave Van Veen | vanveen@stanford.edu | davevanveen.com