We thank our reviewers for their positive comments, and the time they spent carefully reading our paper. Their feedback to clarify some aspects of our work and suggestions to perform additional experiments will improve our manuscript.

Relation to prior work We are glad to inaclude and discuss the papers **(R1,2)** suggested in our related work section. These manuscripts appeared at the time of our submission and could unfortunately not be discussed (Fourier Features arXiv, June'20 and MeshFreeFlowNet arXiv, May'20) but are indeed related to our work.

Additional experiments **(R2) Qualitative & Quantitative comparisons with NERF's approach [5]** do appear in the paper and suppl. under the name Positional Encoding (PE). Specifically, we compared our work to PE for image fitting (Fig.1.) and SDF (suppl.). In particular, we show in Fig.1 that PE does not yield sensible gradients and Laplacians despite fitting the image well. Those derivatives are crucial in our other experiments involving PDEs, and SIREN was the only architecture resulting in sensible higher order derivatives. **Quantitative experiments for SDF fitting (R2)** were performed using the Chamfer distance, eg. for the statue in a $[-1, 1]$ box, we obtain: ReLU=3.78e-4, PE=6.19e-5, SIREN=5.50e-5. We will add those new results to the paper. **Initialization Sensitivity (R3)** We trained SIREN and a Softplus-MLP on Poisson reconstruction for deviations of the initialization's std. dev. from 5% to 50%. SIREN outperforms the Softplus-MLP at all noise levels and is robust to changes in the std. dev. up to 20%, degrading only significantly at deviations of 50%. Further, we show that independent runs of SIRENs have low error std. dev. (suppl. Tab.3). **Deeper networks, normalization layers and other architectures (R4)** We experimented with a variety of normalization layers. We note that the sine activation has normalizing properties, and the proposed initialization scheme guarantees standard normal distributed activations at initialization (see suppl.). Likely as a result of this, normalization layers did not lead to performance gains. Similarly, we trained SIRENs with 5, 10 and 15 layers for solving Poisson's equation (learning rate=1e-6). The attached figure shows number of iterations vs. image L2 loss. Deeper networks lead to faster convergence & better final performance, but also introduce noise in the training process, which, however, does not impact stability or convergence (see difficulty of training below). Lastly, we tested sine activations in a set encoder architecture (suppl. Tab. 4). All those new experiments warrant exciting future work and we will add them in the suppl.

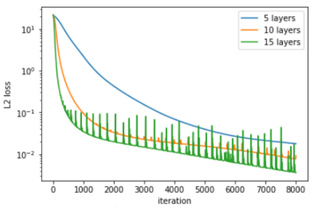

Comment on generalization experiments **(R1)** Generalizing across the space of SIRENs is an exciting avenue for future work and we will comment on this in the discussion section. We chose the CelebA image completion task because those experiments in Conditional Neural Processes (Garnelo et al., 2018) seem to be an appropriate benchmark of encoder-based generalization across a space of functions.

Intuition for sine activation & relation to spectral domain processing **(R4)** Our motivation to use sine activations lies in their infinite VC dimension, their non-locality and shift-invariance arising from their periodicity, as well as the promise that they might address the low-frequency bias of ReLU MLPs. Those comments will be added in the paper. **(R1,2,3)** suggest to comment on the connections with spectral domain processing. We first note that a SIREN with a *single hidden layer* can be seen as performing a frequency decomposition as in Fourier Series. However, the function parameterized by a multi-layer SIREN cannot be trivially identified with such a series. SIRENs, like other deep nets, build on *non-linearities*, while spectral methods for PDEs rely on using *linear* decomposition on Fourier Bases. Hence, relationships to spectral methods are not obvious. **(R3)** suggests transforming the inputs into the spectral domain: this is related to the approach taken by PEs [5], which we benchmark.

Comment on Poisson reconstruction **(R1)** pointed out that Poisson equation can be used for surface reconstruction from normals. This is correct and our SDF loss contains $|\nabla\Phi_x - \vec{n}|$; If this was the only term, it would exactly solve a Poisson equation but our loss has additional constraints to solve for functions also satisfying an Eikonal equation. We chose to demonstrate the Poisson reconstruction and editing on images since they are minimal yet useful examples.

Clarify boundary conditions **(R1)** We use different types of boundary values (Neumann, Dirichlet, and mixed BVs) throughout the paper. In the Helmholtz equation, the BVs appear via the PML formulation (energy at infinity is 0), in the SDF they appear as the points on the surface, which must have a value of 0, the gradient/Laplacian experiments do not have BVs expressed on the function itself (Dirichlet) but on its derivatives (Neumann), hence their solution is up to a constant (resp. an additional ramp). This will be clarified in the paper as well as in the suppl.

Comment on difficulty of training **(R4)** SIREN is easy to train and converges quickly and reliably for the applications we demonstrated in the paper. When using a high learning rate with ADAM, some spikes in the error curve typically appear during training (see above figure). However, these do not impact learning, the loss recovers immediately, and the error systematically decreases. We will comment on this aspect in the supplemental.

Clarify the training with derivatives **(R4)** Training using the derivatives (w.r.t inputs) of the model is performed via automatic differentiation (autodiff). Since the inputs to our networks are the coordinates x and the output is the function value $\Phi(x)$, autodiff can be used to evaluate the derivative $\partial\Phi/\partial x$ by computing the gradient of the network's output w.r.t to its inputs. This gradient calculation is automatically added to the computational graph with autodiff, enabling the optimization of the weights during training. This will be stated in the paper and expanded in the suppl.

Clarify input/outputs **(R4)** We will clearly state the input/output pairs for all the experiments in the paper and will include a summary table in the suppl. Reviewers can also refer to the suppl. video describing the setting (input/outputs and the loss formulation) for each experiment.