Author response for *An Analysis of SVD for Deep Rotation Estimation*, Paper ID 11801.

We thank all the reviewers for their efforts and constructive comments, and for recognizing the contribution of our comprehensive mathematical and experimental analysis in support of $\texttt{SVD}^+(M)$. We first address the concerns of relevance and novelty mentioned by two reviewers: **[R2]** "*overall not surprising. . . an artifact of the incentives in computer vision research*," and **[R4]** *"whether the contributions. . . warrant publication in this venue."*

Surprising: There is ample evidence that our contribution would be received as surprising by the research community. Deep learning research for vision/robotics applications is judged on the output (e.g. 3D reconstruction, depth estimates, skeleton pose), not by their network's internal rotation representation. *Thus the incentive is to use the best available representation*. That domain experts do not consider SVD (L40, [19,4,30,24]) indicates our results would be surprising. This is supported by other reviewers, e.g. **[R1]** "*quite surprised by the result (understandably, as many others).*"

Novel: In addition to the thorough experimental analysis, the mathematical analysis is an important component of the exposition and is also a novel contribution. The error analysis derivation (Sec 3.3, e.g. Corollary 1: $\texttt{SVD}^+$ error is 3σ, $\texttt{GS}^+$ error is 6σ), the theoretical and empirical gradient analysis (Sec 3.2, Supp. 3.1), and discussion on continuity (Sec 3.4), are all novel contributions. This analysis provides the theoretical grounding supporting $\texttt{SVD}^+$ in neural networks.

Relevant: Rotation estimation in neural networks has **[R3]** "*broad applicability to many NeurIPS-related subject areas,*" and is **[R1]** *"a central question in 3D computer vision."* Given the surprising and comprehensive empirical findings, along with a novel mathematical analysis tailored for deep learning and for comparison to state-of-art methods (SVD vs GS [47]), this work is very relevant to the NeurIPS community.

[R3] "*if the continuity described in section 3.4 is the same type of 'global right-inverse' continuity described in [47].*" We use "continuity" in the conventional sense of continuous functions and differentiability. The global right-inverse condition imposed by [47] automatically applies to our setting since our 9D representation space by definition contains SO(3) as a subspace (and SO(3) itself is fixed by the projection functions $\texttt{SVD}(M)$, $\texttt{SVD}^+(M)$)

[R1] "*does not seem to investigate why SVD-plus is better (albeit for a comparison with [47] in Corrolary 3).*" Prior work [47, 20] has carefully analyzed the limitations of classic SO(3) representations in neural networks, so we focused our comparative analysis on $\texttt{GS}^+(M)$ [47] since GS is closely related to SVD and is the current state of the art. We believe our analysis (SVD as the natural robust projection onto SO(3), stable gradients, etc) explain its success in the experiments. We will include a discussion placing our analysis in the context of classic representations.

[R1] "*[L110] the noise distribution over M is Gaussian. . . Bingham and Langevin distributions are better suited to model errors over SO(3).*" Here the noise model represents errors introduced by networks when predicting unconstrained 9D outputs rather than errors in SO(3). We will add the references and include a clarification discussion.

[R3] "*Peretroukhin et al, RSS2020. . . published contemporaneously.*" Thanks for the suggestion. Although this paper appeared *after* the NeurIPS deadline, we will include a discussion and add it to all experiments in the final version. Preliminary results indicate it ranks 2nd for Pt. Clouds (Table 1): mean/med err of 1.97/1.06° vs 1.63/0.89 for $\texttt{SVD}^+(M)$.

[R3] L207: "*large errors. . . due to representation discontinuities*." The ShapeNet airplanes used by [47] contain spaceships with perfect 180° symmetry. We will add images to the supplemental, and rephrase the text to indicate that in general, errors for an unseen test set can depend on representation, model generalization, and data ambiguities.

[R1] "*For Euler angles . . . it is prudent to know of the parameterization*." We treat the network output as *XYZ* Euler angles. We did not consider alternatives since we were following previously established experimental settings, e.g. [47], but we will include alternatives (Cayley) in the final version. We thank **R1** for the references on state estimation and control theory, and will include a discussion in our related work and analysis.

[R1] "*empirical analysis . . . would hold for special cases of rotations (eg. about a fixed axis. . . .*" The KITTI dataset (Table 7) is mostly planar motion. We will add other special cases in supplemental by simulating data with 3D shapes.

[R1] "*if this approach can be extended to . . . SE(3) . . . Sim(3).*" $\texttt{SVD}^+(M)$ could be deployed in a straightforward way for regression to product spaces involving SO(3) by simply decoupling SO(3) from the other terms (e.g. regressing $\mathbb{R}^3$ and SO(3) separately). We leave it to future work to analyze different approaches in practice.

[R4] how "*rotation estimation impacts other 'downstream' computations.*" Inverse Kinematics and KITTI depth (Tables 6 and 7) are examples of established applications where accurate rotation estimates impact downstream objectives.

[R3] "*[LR, other] hyper-parameters.*" The conclusions remained with/without LR-decay (Tab. 1 and Supp 4.2), different losses (Supp 4.3) and encoding models (Supp 4.4.2). We will include an experiment with granular change in LR.

Other points: We will release the experiment code as well (**R2**). We sample random rotations according to the Haar measure on SO(3) (**R3**). We found no change between chordal and geodesic loss (Supp Sec 4.3) (**R3**). We will restructure the paper according to the helpful suggestion from **R4** to include more details in the main body. We will update the analysis summary (L174–176) to reiterate the least-squares optimality of SVD is well-known. (**R2**).