We thank the reviewers for constructive comments and suggestions. Reiterating our contributions: we introduce the spin-weighted spherical CNNs to strike a balance between expressivity and efficiency in the context of equivariant spherical CNNs. Our goal is to be more efficient than the SO(3)-based models introduced by Cohen et al. [7] and more expressive than the purely spherical introduced by Esteves et al. [15]. We achieve it with a computation complexity closer to the more efficient model, while outperforming both models. Moreover, we make a fundamental contribution on equivariant processing of spherical vector fields that enables future scientific applications as suggested by **R3**.

Although reviewers recognized our approach as novel,[R1, R2, R4] interesting,[R1] elegant,[R3] and its promising performance,[R1, R2, R3, R4] they all suggested extra experiments. While we believe our claims are adequately supported, we agree that new experiments will strengthen the paper. We evaluate our model for 3D shape classification[R1] and panoramic image segmentation[R2] and show results in Tables 1 and 2. The classification model takes simpler inputs (4k pixels, two channels instead of the 15k+ pixels, six channels of Jiang et al.[1] and Cohen et al [7]), and a simpler architecture (single branch instead of the two branches of Esteves et al [15]). The performance is competitive nevertheless. One potential improvement that is unique to our model is that we could represent surface normals with spin $s = 0$ and $s = 1$ components, achieving more faithful input representation without sacrificing the equivariance.

	Esteves [15]	Jiang[1]	Ours	Ours + BE
mIoU	0.363	0.383	0.398	0.421

Table 1: Stanford 2D3DS spherical panorama segmentation.

	Esteves [15]	Jiang[1]	Ours	Ours + BE
aligned	88.1	90.5	88.5	89.4
rotated	86.9	-	87.2	86.8

Table 2: ModelNet40 shape classification accuracy [%].

R2 - *"rotation equivariance is not a good property for many data"*. While some datasets are guaranteed to be upright, obviating the need for a globally equivariant model, the benefits of (local) equivariance in terms of filter sharing and efficient use of network capacity still apply. This effect has been demonstrated on CIFAR10/100 by Cohen and Welling [8] and more recently by Weiler and Cesa [36]. Moreover, it is trivial to break equivariance at some layer of an equivariant model, while it is not possible to transform a non-equivariant feature map into an equivariant one. We demonstrate this with the "Breaking-Equivariance" (BE) variants of our model in Tables 1 and 2. We simply add a few standard 2D convolutional layers after the last equivariant feature map. This effectively breaks equivariance so the model can leverage global orientation. As expected, it helps on aligned datasets but not on rotated.

R1, R2 - *"citation and comparison with Jiang et al.[1]"*. This was an oversight of our part; we will include the citations and comparisons elaborating on Tables 1 and 2. Note that the model in Jiang et al[1] is not equivariant to rotations and hence is expected to underperform on rotated datasets, in contrast with ours.

R1 - *"Source code is not provided"*. As mentioned in the introduction, the code and datasets will be released.

R1 - *"benefits of spin basis compared to Wigner D-matrices?"*. The benefits are efficiency and flexibility. The spin basis is related to sets of columns of the Wigner D-matrices. We can choose the number of spins per layer which will define the size of the basis, enabling selection of the desired trade-off between expressivity and efficiency.

R1 - *"further improve the efficiency to match that of [15]"*. Asymptotically, the complexity of Esteves et al [15] is $\mathcal{O}(B^2 \log^2 B)$ while ours is $\mathcal{O}(B^3)$. In practice, this is acceptable in exchange for the expressivity improvements and applicability to a wider range of tasks, while still comparing favorably against the $\mathcal{O}(B^4)$ of Cohen et al [7].

R1 - *"it is not infinite sum"*. In practice, for computation, we need to limit the bandwidth. What equation (2) shows is that, in theory, any square-integrable function on the sphere can be expanded as the infinite sum. This is analogous to the Fourier series of functions on a real interval – it may need infinite terms but in practice we compute a finite number.

R1 - *"when transformation is more complicated ... issues to the model?"*. If inputs are perturbed by transformations that are not rotations, our model needs to learn them from data. This is no different than how standard neural networks operate. If the set of perturbations also include rotations, our model still has the equivariance advantage.

R2, R3 - *"applications to climate, cosmological, geophysical data"*. We truly appreciate the suggestions for possible applications and plan to address them in future work, especially regarding processing of vector fields on the sphere.

R1, R2, R4 - *"overcomplicated, hard to understand the math"*. We believe the prerequisites are at same level of Cohen et al [7] and Esteves et al [15], and we list references to books and papers with the necessary harmonic analysis and spin-weighted functions background. Nevertheless, we will revise the text for clarity and from a didactic standpoint.

All - We thank the reviewers for catching typos, missing data and missing citations. We will update accordingly.

[1] Jiang et al, Spherical CNNs on Unstructured Grids. ICLR'19.