We would like to thank the reviewers for their careful study of our paper and helpful suggestions. There is some spread in the review scores, with reasonable arguments on all sides, but we think that we can address all of the issues raised.

All reviewers seem to agree that the paper makes a significant contribution in unifying the rapidly expanding literature on equivariant CNNs into a single coherent mathematical framework / language (this language being foundational to modern theoretical physics as well), showing how existing work fits in this framework, showing how convolutions emerge naturally from the requirement of equivariance, and characterizing the admissible (equivariant) convolution kernels. We are also pleased to see that the reviewers generally appreciate the clarity and elegance enabled by the mathematical machinery used in the paper. Finally, it is agreed that the material is inherently technical, that unification requires a certain amount of abstraction, and that at the same time the prerequisite mathematics is not part of the standard ML background knowledge. The main concern raised by the reviewers centers on this difficulty of exposition.

It is our aim to make the paper accessible to the mathematically oriented ML researcher, as well as to mathematicians and physicists who want to enter ML. The former should be able to grasp the general ideas and intuition (and know where to learn more), while the latter should be able to relate the ideas to things they already know. Thus the lengthy mathematical build up serves the dual purpose of introducing the concepts (for those not familiar) and establishing notation (for those who are). That said, we do agree with R1&R4 that the build-up is too long and starts from too elementary a point. A complete novice will not be able to follow anyway. We have thus decided to move all of Section 2 (groups, homogeneous spaces) to the appendix. This gives us ~ 1 page of space to address other concerns (see below).

We will keep sections 3-5 (bundles, fields, induced representations) in the main paper, but change the emphasis from formal definitions to intuition and visualization, with some definitions moved to the appendix. In this way, the ML audience can more easily get the general idea while math and physics folks will know enough from reading the names (bundle, field, etc.). We believe that in this way we can address the concerns of R1 and R4 without giving up on the expository style appreciated by R5 and R7.

Given the space freed up by this refactoring, we have been able to substantially improve the discussion based on reviewer suggestions in the following ways:

1. We added an overview after related work where we provide intuition for the main concepts of the theory (bundles, fields, induced rep, equivariant kernel), and discuss their role in the theory. This provides a roadmap for the paper, and makes it clear to the reader where we are going with the definitions that follow (R1).
2. Explain the concept of a principal and associated bundle, as well as the induced representation via the example of a vector field on the sphere (with new figure).
3. When introducing a concept, we explain more explicitly how it relates to familiar ML concepts such as convolutional feature maps and filter banks (R1).
4. Added a short section to explain how the main theorems inform the implementation of G-CNNs (e.g. by parameterizing the kernel as a linear combination of basis kernels that solve the kernel constraint) (R1).
5. We have expanded the examples and made them more readable. Moreover, we explicitly mention where the theory identifies gaps in the literature (R4), such as vector fields (and other fields) on the sphere or diffusion tensor images. These examples can only be described satisfactorily using bundles (R1, R4), a point we clarify in this section.
6. We expanded the discussion of the theorems and their significance. As mentioned by R5, it is important for practitioners to know that the kernel they use is the most general, and one is not imposing additional constraints (unnecessarily limiting expressivity) beyond what is necessary for equivariance to hold.

In response to the question regarding novelty by R4, we would like to point out that the novelty vis a vis Kondor & Trivedi is that we allow for arbitrary fields (not just scalar). In other words, we cover general steerable CNNs and not just regular G-CNNs. Given our framework, the difference is simple to state and may sound small, but without our framework it is actually not at all obvious how to extend K&T to steerable CNNs, or even that fiber bundles / fields and steerable CNNs are in any way related. Concerning novelty wrt application papers, we note that although it is not the main point of our work, our systematic approach does reveal some gaps in the literature (e.g. a network that can process vector fields on the sphere, which is useful in climate science). It also covers methods yet to be discovered.

Although as mentioned in the paper, we do not introduce fundamentally new mathematics, we believe that the observation that all G-CNNs can be described so cleanly in the language of fiber bundles, which is so fundamental to modern mathematics and physics, is highly significant. We believe that together with the work of Kondor & Trivedi, our paper provides a solid foundation for the theory of G-CNNs on homogeneous spaces, and is likely to have a significant influence on future work in this area. It is true that our paper deviates from the standard ML-paper mold, but believe this should in itself not be a reason for rejection. Finally, given the above-mentioned improvements to structure and clarity inspired by the reviewer comments, we think our paper provides a very readable account of the theory.