We would like to thank the reviewers for their feedback. We are glad that each of the reviewers had positive comments about the submission. R1 stated that the paper is “original" and has "interesting results for neuroscience". R2 "liked the originality of the work" and found the algorithm "clearly exposed". R3 expressed that “the paper is clear, well-written and technically sound".

The reviewers wanted to better understand the utility of the developmental strategy over traditional CNN’s. As one example, CNN’s are primarily used for analyzing 2d planar images. However, there is a growing need to analyze spherical images acquired by omnidirectional cameras on cars, drones (ref: omnidirectional camera- Davide Scaramuzza, UPenn). A 2d projection of a spherical image distorts the image, necessitating that a neural-net tile 3d curved surfaces, a challenge for hand-crafted CNNs. Our developmental algorithm, unlike traditional CNNs, self-organizes based on local rules, and can form pooling layers that tile curved surfaces. Here, we show the self-organization of a pooling layer on a sphere (fig-1a).

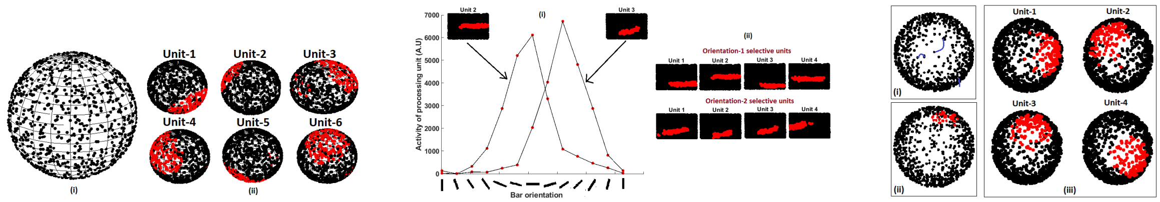

(a) (b) (c)

Figure 1: (a) Self-organizing pooling layers on a sphere. (a-ii) Upstream units connect to spatial patches of nodes on the sphere. (b) (b-i) Tuning curve shows that units in layer-2 have a preferred orientation. (c) Self-organizing networks on Poincare disks. (c-ii) Snapshot of a traveling bump. (c-iii) Receptive fields of units in layer-II.

R3 raises a major concern regarding the novelty of this work as the emergence of complex feature-selective neurons has been described in the 90s. Our major contribution in this paper is towards demonstrating properties such as *flexibility*, *robustness* and *reconfigurability* of the developmental algorithm and showing that these properties allow us to "grow" artificial systems purely by local rules, which to the authors knowledge, hasn’t been explored before. Flexibility is essential for growing networks on curved surfaces (fig-1a), useful for spherical image analysis. R3 recommends that we show **orientation selectivity**. Reconfigurability enables us to grow and self-organize units in layer-2 that are orientation selective, by altering the properties of the emergent traveling bump of activity (fig-1b). We plot a tuning curve to show that units have a **preferred orientation**. R3 mentions that the self-organized networks aren’t "pooling" in the sense of CNN’s. As the sensor-nodes in the input layer are not evenly spaced, the classical definition of pooling breaks down. We refer to **pooling**, when units across layer-2 are connected to similar sized spatial-patches of nodes in layer-I. Post self-organization a max/average operation can be applied by the unit in layer-2, making it max/average pooling. A tight regulation of spatial-patch size has been observed in our networks as shown in histogram fig-5e in the paper.

R1: (Roson and Bauer) show evidence of dimensionality reduction from retina to LGN. R1 also pointed out that the developmental algorithm could be used for **bio-inspired hardware**. We completely agree with this as we’ve begun adapting this algorithm for implementation on Loihi (Intel’s neuromorphic hardware). Currently, to configure spiking neural nets (SNN) on Loihi to perform tasks, one needs to specify a hand-designed SNN topology, manually define neurons that belong to different layers and specify connections between layers (ref: Chit-Kwan & Wild). Instead of hand-programming these networks, our developmental algorithm would enable neuron clustering into different layers via the growing process and self-organize connections between these layers through spontaneous activity in the lower layers. This would provide a flexible and scalable way to program neuromorphic hardware. As suggested by R1, we shall make additions to the abstract and changes to the convolutional terminology.

R2: Spiking neurons have been used instead of neural masses to allow for future implementation of this algorithm on hardware dedicated for spiking networks (neuromorphic chips). R2 was concerned about the biological plausibility of global inhibition scheme. As the magnitude of inhibition decays exponentially with distance, a small network (2000 nodes) has every node inhibiting $\sim 90\%$ of the nodes in the network (as seen in insects). However, as we scale the network to 50000 nodes, every node inhibits $< 4\%$ of the nodes in the network, making it biologically feasible. We have also shown scaling of the developmental algorithm with the network size in supplementary (S5). R2 also suggested that we demonstrate the formation of pooling networks on a **Poincare disk**, with a non-uniform distribution of sensor-nodes. This is shown in Figure-1c. R2 wanted clarification on how random networks performed with high accuracy. Functionality of the network has been measured by connecting 2 layers of the network, either hand-crafted, grown from a single unit (our developmental algorithm), or random to a perceptron to perform the MNIST task. Only the perceptron was trained while keeping the first 2 layers fixed. This serves as a control to show that "grown" networks extract useful features and perform as well as hand-crafted networks (detailed in supplementary-S4). We have attempted to answer major comments and append results suggested as improvements, but are limited on space to answer all comments.