We thank the reviewers for their constructive comments. Here we will focus on the main concerns. We will address smaller points and clarity of the equations, use of minus signs and symbols into account when we revise the paper.

1) Weight symmetry: R1,2,4 raised questions about the neuroscientific plausibility of weight symmetry and if small deviations from symmetry in the brain would be detrimental. BrainProp does not suffer from the weight transport problem but solves it because synaptic updates of feedforward and feedback connections are proportional so that their strength becomes proportional during learning. While exact reciprocity/proportionality of feedforward and feedback connections is not present at the level of single neurons in the brain, such reciprocity at the level of cortical columns (with many connections in both directions) is plausible. Indeed, the columns have mechanisms to switch on plasticity, such as specific types of interneurons in a column that gate plasticity, as outlined in ref. 27. **R1** asked us to expand on evidence for reciprocity of feedforward and feedback connections. The selective attention literature provides strong evidence for this (e.g. ref. 27). We will here illustrate it for eye movements (our revision will mention feature-based attention too). If an object is selected in the frontal eye fields (FEF) for an eye movement, neuronal response in early visual cortex elicited by the same object are enhanced a short time later (as anticipated by **R1**) (ref. 29). Moore and colleagues used microstimulation in the FEF [1] and observed that feedback is channeled to a tiny hotspot in the visual field where activity is enhanced, implying a exquisite specificity of feedback connections. In the neuroscience field, the consensus is that attentional feedback selectively reaches those neurons in sensory cortex that gave input to the motor response (just as in BrainProp). Importantly, these neuroscientific findings imply reciprocity of feedforward and feedback connections. **R1** was interested in how feedback connections, which carry a distinct signal, influence plasticity. This is a neuroscientific finding, which has been reviewed in e.g. ref. 27, which also discusses cellular mechanisms. **R1,4** asked for results on learning of symmetric weights. In the few days before the deadline of this rebuttal, we ran an experiment on CIFAR10 on the smaller convolutional architecture with randomly initialized feedback weights and weight decay (untying the weights): we reached equivalent accuracy, and conclude that BrainProp learns approximate weight symmetry. Feedback alignment fails on simple problems and is known not work at all in deeper networks.

2) Comparison to AGREL: R1,3,4 asked about the relation between BrainProp and AGREL. BrainProp follows from new insights of how plasticity can be switched on/off in the appropriate columns in a deep network (Fig. 3), whereas AGREL dealt with a single hidden layer. In spite of weight updates based on information that is locally available at the synapse, BrainProp solves many tasks that have not been yet been learned by competing schemes.

3) Learning speed: R1 would like to see a more thorough analysis of learning speed compared to EBP. The apparently larger gap at the start of learning is caused by the initial slow learning phase of BrainProp, when it has to find classes by trial and error because there is no teacher. Thereafter, it catches up by only making errors at category boundaries, i.e. images that matter to get better. Our revision will include a thorough analysis of the initial and later learning speed and how it depends on the number of classes (note that we have already tested problems with 10, 100 and 200 classes).

4) Homogeneity of neuromodulatory effects: R1 asks about the homogeneity of e.g. dopamine influences. The release of neuromodulators is indeed homogeneous, see e.g. [2] and ref. 20.

5) Formal comparison to BP and notation of y^N: R2 asked us to clarify the formal comparison between RPE x FB and EBP. We will improve our description. If we sample all actions uniformly and synaptic updates occur afterwards, the result is mathematical equivalence to EBP. We agree with the point of **R3** about the notation of y^N and will improve it by introducing an extra symbol for the one-hot post selection activity vector.

6) Action selection: R3,4 ask us to clarify the action-selection process. We included 2% of explorative choices, for which the network used a softmax function (mostly sampling nearby categories) and the other 98% of choices were greedy. The results do not depend critically on this percentage of explorative choices. We will include a succinct algorithmic description as suggested by **R3**.

7) Equilibrium propagation: We agree with **R3** and will improve our description.

8) Classical RL tasks: R4 asks about the application of BrainProp to classical RL tasks. This was the topic of AuGMEnT, a learning scheme that works for a network with one hidden layer. When comparing BrainProp to AuGMEnT it follows that BrainProp should be able to solve challenging RL tasks, although we have not yet explored this. Future work might consider e.g. Atari games because BrainProp can train networks with sufficient depth to build useful internal representations of task events.

9) Future challenges: R2 suggests to include a discussion of future challenges, such as generalization to spiking neurons. We have initial results on this, where excitatory feedback neurons carry a linear signal of limited dynamic range by using a fixed activity offset, similar to baseline activity of dopamine neurons (ref. 19). Moreover, Brainprop generalizes to continuous actions spaces (point raised by **R3**) when applied to the actor-critic RL paradigm Deep Deterministic Policy Gradient (DDPG). BrainProp could be used for a biologically realistic implementation of DDPG in deeper networks.

[1] Armstrong, Fitzgerald and Moore, *Changes in visual receptive fields with microstimulation of frontal cortex.*

[2] Gaspar, Stepniewska and Kaas, *Topography and collateralization of the dopaminergic projections to motor and lateral prefrontal cortex in owl monkeys.*