We thank the reviewers for their thorough reading and thoughtful comments. Below we address their major concerns. All minor points are corrected in the revised draft.

Reviewers R1 & R2 agreed with some of the concerns we raised in the original submission regarding **scalability and generalization** of our results (old Fig. 2). Since submission, we have made substantial progress in scaling our approach and networks. We now use an evolutionary search strategy (i.e., CMA-ES; Hansen, arXiv 2016) in the outer loop of our algorithm, to replace SGD and finite differences. CMA-ES has significantly improved scaling properties and allowed us reproduce our results with larger (up to N=100) rate networks (see Figure below) and substantially shortened computing times for both rate and spiking neuron simulations (not shown here). We will include these new methods and results in our manuscript. The rapidly expanding optimization literature will allow us to meta-learn plasticity rules for problems of larger scale and complexity (e.g., via reinforcement learning strategies, such as Abdolmaleki et al. arXiv 2018). This strategy has already greatly improved the performance of our code. The next challenge will be developing a tool to go beyond single GPU computers to simulate many of such networks in parallel. Even if questions such as how to make a spiking network learn and recall stimuli are difficult problems to date, we are confident that they must be tackled. Our work presents a first avenue towards this end. It should be noted that scalability is idiosyncratic, i.e., specific to the problem rather than framework-wide, and challenges will change for network-level problems. For example, the loss function to discover Oja's rule aims to match every synaptic weight to a desired value. This is not required for other problems, and the loss-landscape for a network task may thus be easier, or at least *different* but not necessarily more difficult to navigate.

Reviewers R2 & R4 also raised questions about the **relevance of our work for neuroscience and additional value of a meta-learning approach.** Uncovering the plasticity rules that underlie brain functions such as learning or memory is a key and open question in neuroscience. Plasticity rules directly deduced from experimental data have thus far failed to exhibit rules that accomplish goals useful for the brain on their own. Meta-learning approaches similar to those develop here could help bridge the gap between empirical observations and function. Our study is a proof of principle to show that known/analytically predictable rules can be successfully recovered by a meta-learning approach on a biologically interpretable search space of plasticity rules.

Analytical approaches deducing a small number of plasticity rules that optimally implement a given function have also shown to be successful (Pehlevan and Chklovskii IEEE 2019). However, one cannot be sure that the learning rules used by the brain are analytically tractable, or which functions are being optimized by the brain -if any-. The dizzying number of different synapse types and thus learning rules at play in the brain make alternative, numeric strategies imperative, because we can also deduce rules that don't necessarily behave *well*. The ultimate goal of our approach is to find learning rules from large datasets, by comparing, e.g., patterns of activation before and after learning, and asking by what rules a simulated network has to be altered to achieve comparable dynamics. Such goals seem within reach. From the ML perspective, the search for unsupervised learning rules able to produce useful representations that could provide alternatives to backpropagation is an active area of research (Jaderberg et al. ICML 2017). Our approach provides real-world challenges that sit at the interface of neuroscience and could contribute to such alternatives.

In summary, our approach is complementary to both analytical and experimental ones, as we can apply any loss function *in silico*, e.g., functions that incorporate experimental data, desired behavioral outputs or functional hypotheses. Our meta-learning approach has many advantages with regard to its generality, flexibility and interpretability that we believe will provide a great backdrop for stimulating discussions at NeurIPS 2020.

Minor points. We thank the reviewers for their suggestions of **prior work**, which we now cite and discuss. R2 asked **how a non-differentiable system (old Fig. 4) could be trained**. - We previously used finite differences to compute gradients of the loss w.r.t. plasticity parameters, using perturbations that change the number of output spikes in at least one data set. We now also use CMA-ES where a selection mechanism governs plasticity parameters updates instead of gradient descent. R4 commented on **the plausibility of the spiking set-up** (old Fig. 4). - While it is not realistic to assume inhibitory plasticity acts in isolation, we chose a simple rule as a first goal in a spiking network. It is a rule ascribing a given function to a network, stable and analytically tractable. R4 will hopefully be relieved to know that including excitatory plasticity is currently in progress and we will report our progress for combined plastic inhibitory and excitatory synapses in the camera ready version.

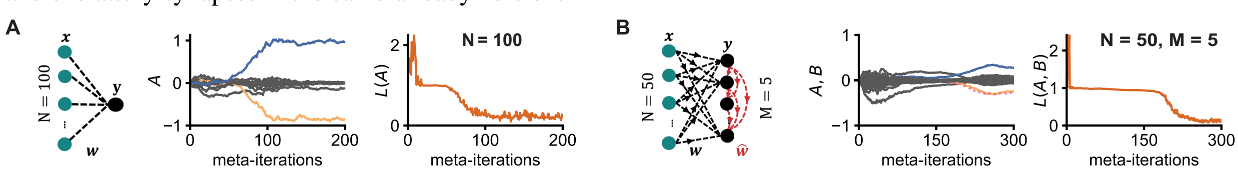

Figure 1: **Extended results** for old Fig. 1, with $N = 100$ inputs (A) and Fig. 3, with $N = 50$ inputs and $M = 5$ outputs (B), using CMA-ES instead of SGD + finite differences. The same color scheme and notations as in the original figures were used. In (B) the parameters of the two co-optimized rules have been collapsed in a single plot.