We thank all the reviewers for their thoughtful comments and positive feedback. We will first address two major points raised by the reviewers and then answer individual questions.

Major Points:

1. **(Reviewers 1 and 3)** - "Providing more empirical results and refined baselines will improve the experiments section and be useful to study the strength of this algorithm in practice." (paraphrase)
   - We agree and will have more experiments in revision. Specifically, we plan to include few-shot learning evaluations on MiniImageNet [47] and detailed numerical results for both the meta-learning and federated learning settings. Furthermore, following the suggestion of Reviewer 1, we will investigate a comparison with Reptile+Meta-SGD [39], which would also learn a per-coordinate learning rate. One caveat here is that Meta-SGD was designed for a MAML-like approach of only taking one gradient step within-task, whereas Reptile takes many iterations; since Meta-SGD uses higher-order differentiation, repeated application may slow it down. Note that we do discuss (MAML+) Meta-SGD briefly in lines 305 -> 307.
2. **(Reviewer 2)** - "The paper is missing a detailed discussion/examples of the behaviour of V_Ψ, which makes it hard to judge the sensibility of the proven bounds."
   - We agree that examples of V_Ψ are needed to understand the results, as it is the main measure of task-similarity and if V_Ψ is small then so is the average regret. For the case of a fixed comparator (Theorem 3.2), we do give a simple example in lines 136 -> 137; here V_Ψ is proportional to the empirical standard deviation of the optimal task-parameters, so if they are close (i.e. tasks are similar) then V_Ψ is small. As we will describe in more detail in revision, the case when the comparator is varying is similar, as V_Ψ is now proportional to the average deviation of the optimal task-parameters from a shifting sequence of vectors. For example, if we see one task every day for year, and Ψ is a sequence that fixes a single comparator for each month, then V_Ψ^2 is roughly the average over the months of the empirical variance of each month's thirty or so optimal task-parameters from that month's fixed comparator. This can be very small if the variation between tasks is well-described by a seasonal trend.

Reviewer #1:

1. "It would be useful to include a reference for the regret of OGD." (paraphrase)
   - We agree and will add a pointer to the Shalev-Shwartz survey [49, Theorem 2.15 for $R(w) = \frac{1}{2\eta}\|w-\phi\|_2^2$].
2. "Is the form of the regret-upper-bound $\hat{\mathbf{R}}_t(x_t)$ nice in cases more general than just that of OGD?" (paraphrase)
   - Yes! This is a main reason we expect this framework to be broadly applicable. In our paper, we show that several results (specifically Theorems 3.1 and 3.2) hold for any algorithm in the OMD/FTRL family, which includes not just OGD but also other classical methods such as exponentiated gradient/multiplicative weights. Even more generally, regret guarantees often include terms that depend on some measure of distance from an initial state, which are often amenable to study (e.g. because norms are convex). We will elaborate on this in the revision.
3. "On the line 108 -> 109, it's not clear how $\hat{\hat{\mathbf{R}}}_T \le o(T) + \min_x \frac{1}{T}\sum_t \hat{\mathbf{R}}_t$."
   - We believe this asking how the right-hand-side goes to zero. This is a typo - the first term should be $o_T(1)$, i.e. sub-constant, not sub-linear. Similarly, the last term in the statement of Theorem 3.1 should be divided by T (the expressions in the proof are correct as-is). We apologize for both errors and will correct them in revision.
4. "For FedAvg case, does the improvement comes from the meta-learning treatment (where we optimize for the initialization) or the ARUBA algorithm itself (e.g. the fact that the learning rate is adaptable)?" (paraphrase)
   - In fact to get FedAvg+ARUBA we do not modify FedAvg except to adapt the learning rate - the global model is still learned in the same way. This is possible because FedAvg is equivalent to Reptile with the outer-loop update coefficient set to 1.0. So the improvement is indeed coming from the adaptivity.

Reviewer #2:

1. "Complicated and overloaded notations: too many versions of regret and bounds with very similar symbols, the sequence of ψ_t for dynamic regret is not defined for notations of Theorem 3.1."
   - We will make sure that the mathematical presentation is as clean as possible in the revision. One way of reducing the many variations on the regret notations is to depend less on accents and represent regret-upper-bounds by a different capital letter (e.g. $\mathbf{U}$) to better distinguish from regret terms ($\mathbf{R}$). As for the sequence of ψ_t, this can be any arbitrary sequence of vectors in the action space Θ.

Reviewer #3:

1. "What factor leads to the big experimental difference between meta-learning, where the improvement over Adam is relatively small, and federated learning, where the improvement over FedAvg is a lot more significant?" (paraphrase)
   - Our explanation, which we will add in revision, is based on the nature of data. Whereas standard evaluation datasets in few-shot learning consist of tasks with identical amounts of i.i.d. data, the Shakespeare benchmark we use for federated learning has tasks with highly variable amounts of data (two different roles can have very different numbers of lines) and data that is not i.i.d. (lines are not shuffled but split in their order of appearance in the play). Our method may be better able to handle such data, for example by tempering noisy directions in low-data tasks by learning which directions are important based on the distance traveled in high-data tasks.