We thank the reviewers for their thoughtful feedback. We are encouraged they found that: our proposed evaluation framework is interesting, well motivated and captures the important desiderata of continual learning (CL) [**R2**, **R3**, **R4**]; our proposed algorithm C-MAML is suitable for the setting, promising and performs very well [**R2**, **R3**, **R4**]; our empirical evaluation is strong [**R3**]; our proposed unifying framework is well described and clears up confusion caused by prior work or that our work was well placed within the literature [**R2**, **R3**, **R4**]. We are further encouraged by **R4**'s approval on all of OSAKA's features, realistic focus and high relevance for NeurIPS.

@**R2** **OSAKA seems too specific while failing to make a broader argument for other approaches to CL** We wholeheartedly disagree. OSAKA's purpose is to align CL research with the *deployment* of autonomous CL systems, e.g., a virtual assistant or a general-purpose robot that would keep on learning about new users and new environments (more examples in Sec. 1 & 3). This scenario encompasses most of the reasons why we study CL. As hypothesised and then shown in the empirical section, other approaches to CL are not well suited to handle the broadness of real-life's requirements or of OSAKA's challenges and thus can not be applied at deployment time.

@**R2** **Pre-training limits the generality of OSAKA and adds computational needs.** We disagree. OSAKA aligns with the deployment of CL systems in real life (Sec. 1 & 3) and it would be more realistic to deploy an agent with some knowledge of the world. Nevertheless, pre-training is not mandatory, although prescribed, and we have a baseline that does use it (C-MAML). Furthermore, it is currently more computationally efficient to learn on i.i.d. data at pre-training than on non-stationary data at CL time and pre-training is a one-time cost compared to CL which is a recurring one.

@**R3** **Why putting features of different frameworks, e.g., from MOCA, together is useful for continual learning evaluation?** We unified *and extended* these features to create a more realistic setting than the ones studied in the previous literature. Other frameworks study some of the features in silos but when methods are tested in less realistic settings some methods perform better than they should [12]. See Sec. 6.3 (under dynamic representations) for such an example.

@**R3** **several continual learning papers already measure the cumulative performance in task-agnostic setups, e.g. MOCA.** We have made a thorough literature review and we haven't found papers that measure cumulative performance apart from concurrent work MOCA. However, we are happy to add these papers if the reviewer includes them in their final review. Furthermore, Appendix B.1 is devoted to contrasting MOCA and OSAKA and describes their differences: context-dependant targets, OoD tasks and the expansion of the set of labels. **R4** (and to some extent **R2**) agrees with the novelty of OSAKA as well as its potential to greatly enhance the field of CL.

@**R4** **I think it is strange that MAML performs better in the 0.90 setting.** The reviewer's intuition is right. However, C-MAML needs to predict correctly the context switches otherwise it will get mixed gradients from different tasks. Thus, $\alpha = 0.90$ can be more challenging for methods with dynamic representations when the OoD tasks are not too far from the pre-training ones, as in the Tiered-Imagenet experiment. We added this insight in the paper.

@**R2** **How would you adapt OSAKA to RL?** We can replace the current vision tasks by ones from a multi-task RL benchmark such as different tasks in robotics, or change the *contexts* within a standard benchmark, such as different gravity levels in a Mujoco environment. One would control the non-stationarity via the time allocation (or number of episodes) in each task/context.

@**R2** **Is task-revisiting an implicit form of replay that reduces catastrophic forgetting?** Yes, because real-life agents encounter the same or similar tasks through their lifetime [**R4**]. Thus, we can build systems that benefit from this natural implicit replay rather then spend computations to constrain models to *always remember all past distributions*.

@**R2** **Section 2 on me was tiring and confusing.** As **R2** acknowledged, this is "partly inherited from prior work" and we understand the confusion. It seems however that **R3** and **R4** thought our unifying framework was well described. We would greatly appreciate if the reviewer could expand on this point.

@**R2** **Comments in the pseudo-code would be helpful.** Complete algorithms including comments are in Appendix A.

@**R2** **Related work is only immediately related to the submission.** Appendix B provides an exhaustive related work.

@**R2** **Lack of a technical introduction to MAML.** You can find the MAML algorithm in all the algorithms in Appendix A under the pre-training phase.

@**R2** **About the change detection mechanism.** See Figure 4 in Appendix for some results on its accuracy.

@**R2** **Without task revisits, does ϕ stop being suitable for few-shot learning?** It stays suitable because it is still trained with the MAML loss, which optimizes for few-shot learning.

@**R2**,**R3** Thanks for pointing out typos, an error in the bolding of results, and a redundant part in the introduction. We have corrected them. We will also use the extra allowed page to increase readability in Sec. 2, answer **R2**'s questions in more details and further justify the unification of OSAKA's features [**R3**].