We thank the reviewers for their feedback. We are encouraged that they found that the empirical study is very important (R3), the experiments are systematic (R2), reasonable / sufficient (R5), and outperform prior work (R5), and that the method is simple (R4), easily reproducible with code (R2,R3), and useful for all RL researchers (R3). We address the reviewers' points of feedback and comments below, and will incorporate all of their feedback.

─────────────────────────────────── **General Feedback** ───────────────────────────────────

Novelty (R5). "R5: *The novelty is not clarified. In summary, the contributions are not enough, which includes some data augmentations.*" Thank you for your feedback! Our contributions and novelty are stated in the introduction of the paper. They are (i) first extensive study of data augmentations in the context of RL on various standard benchmarks (as R2, R3 and R4 mentioned) and (ii) introduction of new data augmentations for both pixel and state-based RL (as R2 and R4 mentioned). It was not obvious before this work that data augmentations were useful in the RL setting, since they are usually not utilized in pixel-based RL literature. We will be sure to emphasize our contributions in the text to make these points clearer.

Simplicity (R5). "R5: *The method is simple that just like combining data augmentation with RL.*". We agree that the method is simple, though we view this as a positive. As R4 pointed out, *"A major advantage of the proposed data augmentations is that they do not require any modifications to the baseline RL algorithms, and their benefits can be realized by simply appending the original observations with the augmented data."*

Generating images (R5). "R5: *We expect to see the challenges as generating images.*" Thank you for the suggestion! Generating images / augmentations will be an exciting avenue for future research.

─────────────────────────────────── **Text & Logic** ───────────────────────────────────

Explaining why data aug works (R4). "R4: *Not entirely sure why the data augmentation ... works, and this paper does not provide a very good explanation*". Thank you for pointing out that our explanation could be clearer. Though we do provide empirical investigations into why data augmentations perform well in Figures 2(b) and 4 of the main draft, we can certainly improve our explanations. Concretely, we ablate which parts of random crop contribute most in Figure 6 of the appendix, where we find that translation invariance is the most important aspect of cropping. We can move this result to the main body and discuss it there to make the explanation clearer.

Justification on random amplitude scaling (R4). "R4: *The argument that randomly scaling the true state variable would make the algorithm robust to noise ... should not apply in this verification.*" We remark that there is some randomness from initial state distribution even though tested simulation environments have deterministic transition distributions. Because of that, our argument can be valid in our experimental setting. However, we will clarify this in the final draft. Thank you very much for your pointer.

Introducing new data augs (R3). "R3: *Random translation is claimed to be one of the two newly introduced augmentations ... but it's never specifically introduced.*" Thank you for pointing this out to us. We will correct the text in final draft to avoid confusion.

Figure 12 legibility. (R2) "R2: *It's hard to read Figure 12. Maybe smoothing the curves can help. Figure 12 only plots the curves for at most 200k steps.*" We will replace figures with more training timesteps for clarity in the final draft.

Random amplitude scaling with multiple variables. (R2) "R2: *L261-263: If multiple variables are used, the relative differences can also be changed.*" That's a good point. As you mentioned, the relative differences can be changed in the case of random amplitude scaling with multiple variables. Because of that, random amplitude with a single scalar achieves the better performance on most environments. We will clarify this part in the final draft.

─────────────────────────────────── **Experiments** ───────────────────────────────────

Results on state-based RL (R3). "R3: *Results on state-based RL ... are not that significant.*" We believe that our experimental results on OpenAI Gym are extensive and demonstrate the strength of RAD in that (a) we consider strong baselines, such as POPLIN and PETS, and (b) our method provides large gains in complex environments like Walker.

Performance of TD3 algorithms (R2). "R2: *The experiments on state-based environments are not convincing ... performance of TD3 is quite bad, compared to TD3 paper.*" We emphasize that the version of OpenAI Gym environments is different from TD3 paper (we used the setups of POPLIN, which is published in ICLR 2020), and we took the best reported performance for TD3 in POPLIN. We also checked out that similar scores can be reproduced by the official codebase from the POPLIN paper (e.g., 3273.4 on Cheetah and -447.3 on Walker using 10 random seeds). To clarify this concern, we will update the scores on all environments using 10 random seeds.

Application specific (R5). "R5: *More a engineering work for a nice application*" We show state-of-the-art results on common RL benchmarks like DeepMind control, ProcGen, and OpenAI Gym, which are standard benchmarks used to study RL by a suite of other general purpose RL algorithms (CURL, SLAC, Dreamer, PlaNet, PPO, POPLIN, PETS).