

How to Stay Curious while avoiding Noisy TVs using Aleatoric Uncertainty Estimation

A. N. Mavor-Parker¹,

K. A. Young^{2,3}, C. Barry^{2,†},

L. D. Griffin^{4,†}

^{1,2,4} UCL, ³Boston University

† equal contribution

a.mavor-parker@cs.ucl.ac.uk

self-supervisor.github.io



Creative Machine Learning

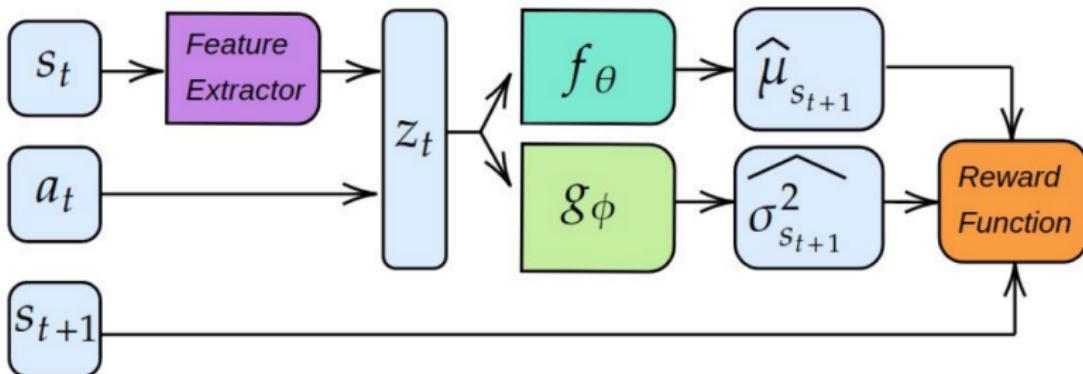
- ▶ When the environment does not provide us with rewards we must generate our own **intrinsic** rewards¹
- ▶ Curiosity driven learning¹ uses **surprise** as intrinsic rewards



- ▶ Things are often **impossible** to predict that doesn't mean they are meaningfully surprising...²
- ▶ These are "**Stochastic Traps**" - Shyam et al.³ or 'Noisy TVs' - Schmidhuber et al.⁴



- ▶ Our intrinsic reward computation reduces surprise when there is high **aleatoric uncertainty**⁵
- ▶ s and a are states and actions, f and g are prediction networks for the mean and aleatoric variance of the next state



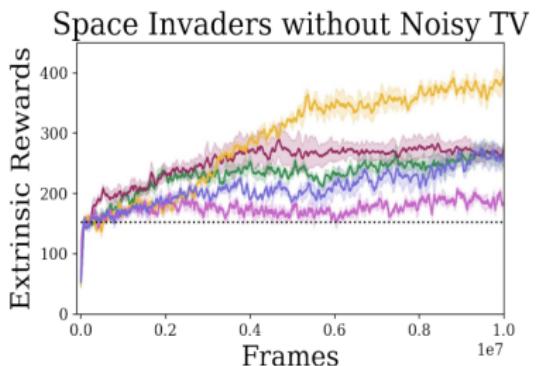
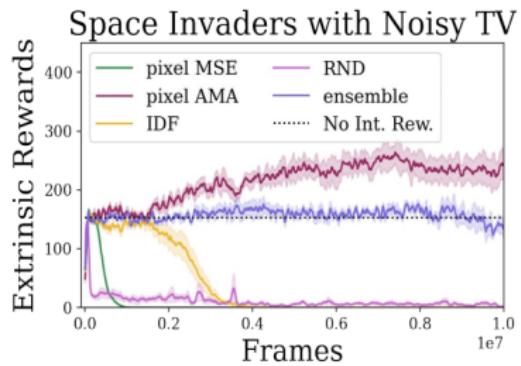
Play Retro Games or watch CIFAR-10 images?



Image from gym-super-mario-bros documentation



Example CIFAR images tiled across game screen



- ▶ AMAs can avoid **Noisy TVs** in curiosity driven learning
- ▶ We tested on artificial noisy TVs and isolated a natural noisy TV in Bank Heist
- ▶ This suggests a possible role for acetylcholine coding for **“expected” aleatoric uncertainties**⁷—future neuroscience research should compare the AMA model with **biological data**
- ▶ Future RL research should integrate AMA further into state of the art exploration algorithms and investigate how frequently Noisy TVs appear in real world applications

- 1 - Pathak, D., Agrawal, P., Efros, A.A. and Darrell, T., 2017. Curiosity-driven exploration by self-supervised prediction. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (pp. 16-17).
- 2 - Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T. and Efros, A.A., 2018. Large-scale study of curiosity-driven learning. In International Conference on Learning Representations. 2018.
- 3 - Shyam, P., Jaśkowski, W. and Gomez, F., 2019, Model-based active exploration. In International Conference on Machine Learning (pp. 5779-5788). PMLR.
- 4 - Overview on Schmidhuber's website: <http://people.idsia.ch/~juergen/interest.html>
- 5 - Kendall, A. and Gal, Y., 2017. What uncertainties do we need in bayesian deep learning for computer vision?. In Advances in neural information processing systems (pp. 5574-5584).
- 6 - Pathak, D., Gandhi, D. and Gupta, A., 2019. Self-supervised exploration via disagreement. In International conference on machine learning (pp. 5062-5071). PMLR.
- 7 - Angela, J.Y. and Dayan, P., 2005. Uncertainty, neuromodulation, and attention. *Neuron*, 46(4), pp.681-692.