Understanding the Effects of Random Shift Augmentations for Image-based Reinforcement Learning

David Klee, Robin Walters, Rob Platt
Northeastern University

Introduction

Reinforcement learning using image observations is challenging. Most methods struggle to learn a representation that is useful for policy learning, resulting in poor performance or high sample complexity. Data augmentation, in the form of random pixel shifts, has been shown to overcome these issues for certain continuous control tasks [1, 2]. The success of shift augmentation is surprising since conventional understanding [3] suggests the shift invariance of the convolutional encoder should be adequate.

Method

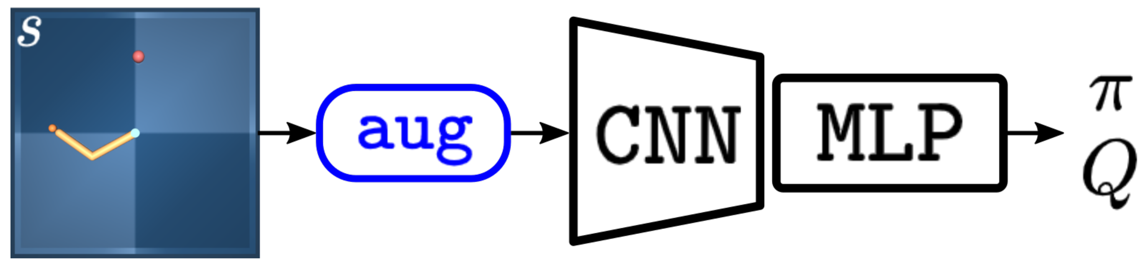

- Actor-critic method is trained for 100k environment steps on DMControl Suite environments, using a shared convolutional encoder that provides latent vector to actor and critic networks
- Random translations of ± 16 px are performed on s_t and s_{t+1} when computing TD loss

Experiments

- Ablation study of shift augmentation to identify what aspects of the model benefit from augmentation
- Comparison of shift augmentation to traditional regularization techniques
- Investigation into effects of augmentation on invariance properties of the learned model

Results

- **Augmentation helps fully-connected layers**: augmenting the final feature map achieves similar performance to augmenting the input image
- **Restoring shift invariance is not the key benefit of augmentation**: augmentation is effective even if the convolutional layers are perfectly shift invariant to the transformations
- **Learned representations have spatial continuity**: shift augmentations cause convolutional layers to exhibit shift invariance *and* equivariance
- **Augmentation outperforms traditional regularization techniques**

Shift augmentation gives you more than **shift invariance**, it provides **spatially continuous** representations.

Take a picture to download the full paper

Ablating Pixel Shift Augmentations

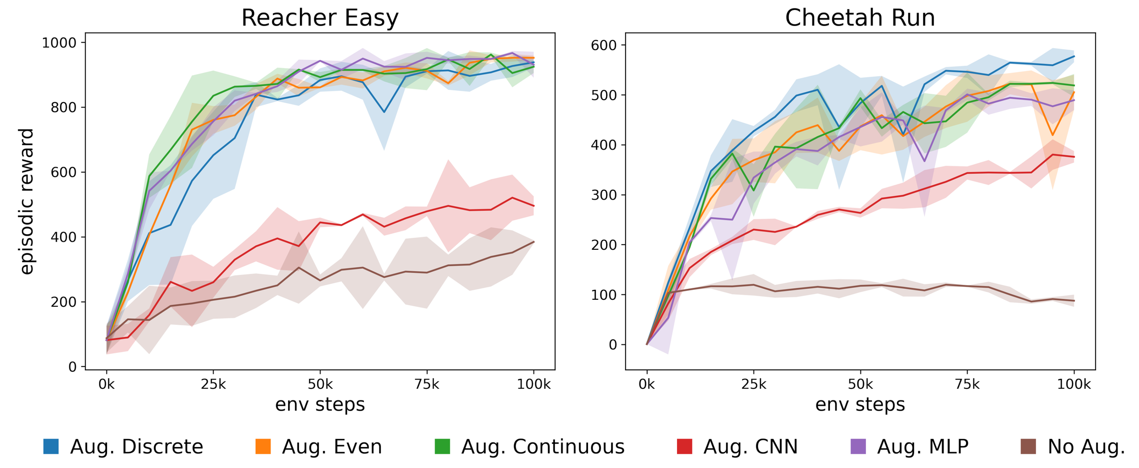

Comparing performance with variations on the set of pixel shifts (`Discrete`, `Even`, `Continuous`) and where the shift occurs (before `CNN` or before `MLP`).

Visualizing Learned Feature Maps

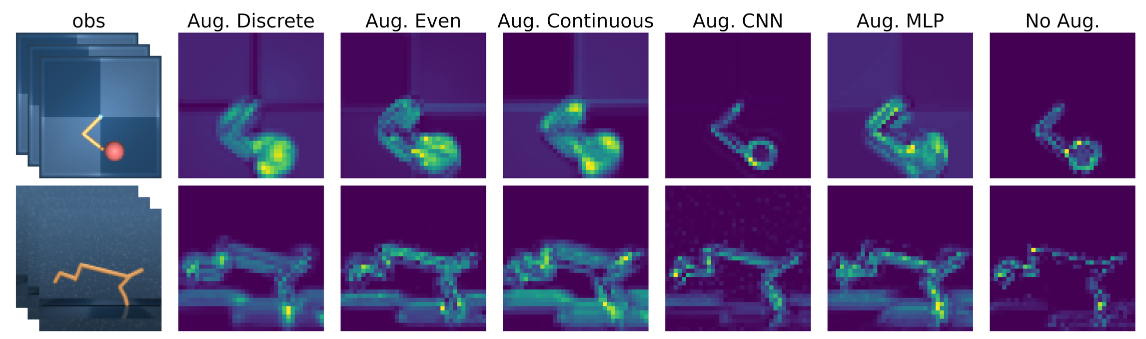

Spatial attention maps of encoders shown in learning curves above. Methods that achieve higher returns have spatially continuous encodings.

Shift Augmentation vs. Regularization

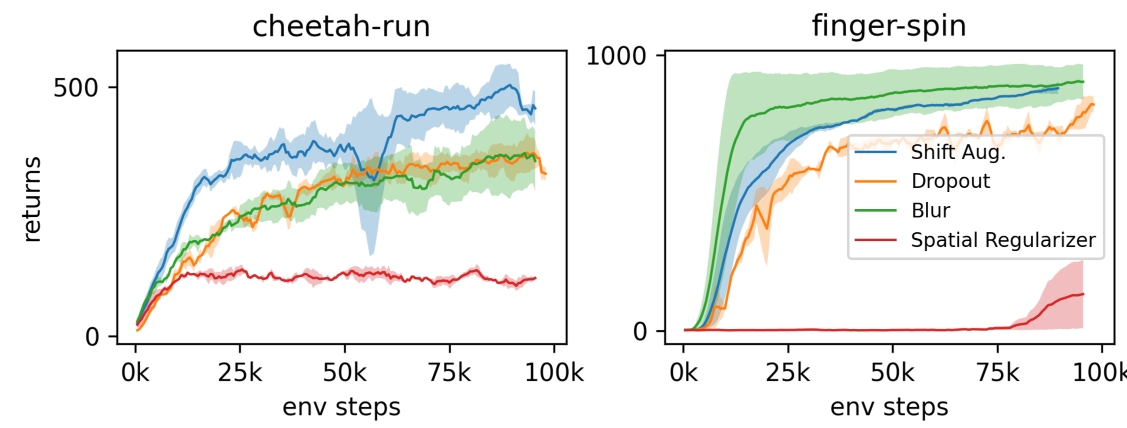

Applying spatial dropout or Gaussian blur to the final feature map can be beneficial, but are not as consistently effective as shift augmentation.

Shift Augmentation vs. Shift Invariance

A function f is equivariant to transformation $\mathcal{T}$ if:

$$\mathcal{T}' \cdot f(x) = f(\mathcal{T} \cdot x)$$

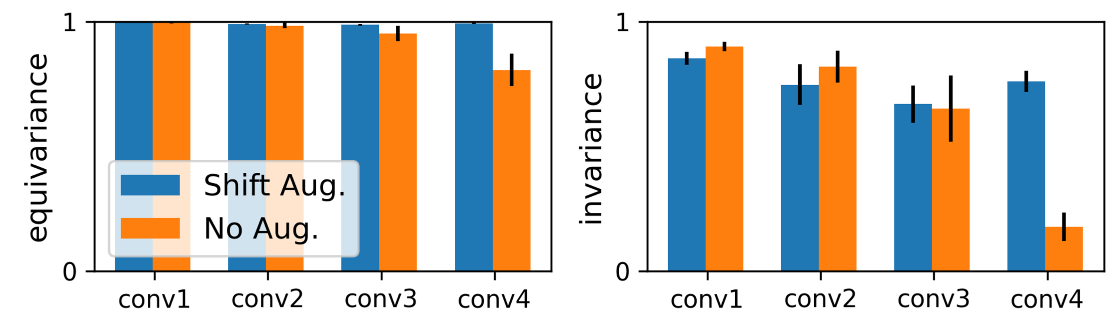

[1] Michael Laskin et al. "Reinforcement learning with augmented data". In: *arXiv preprint arXiv:2004.14990* (2020).
[2] Ilya Kostrikov, Denis Yarats, and Rob Fergus. "Image augmentation is all you need: Regularizing deep reinforcement learning from pixels". In: *arXiv preprint arXiv:2004.13649* (2020).
[3] Aharon Azulay and Yair Weiss. "Why do deep convolutional networks generalize so poorly to small image transformations?" In: *arXiv preprint arXiv:1805.12177* (2018).