Attend Before you Act: Leveraging human visual attention for continual learning

Khimya Khetarpal and Doina Precup

Mila-McGill University, Montréal (QC) Canada

Motivation

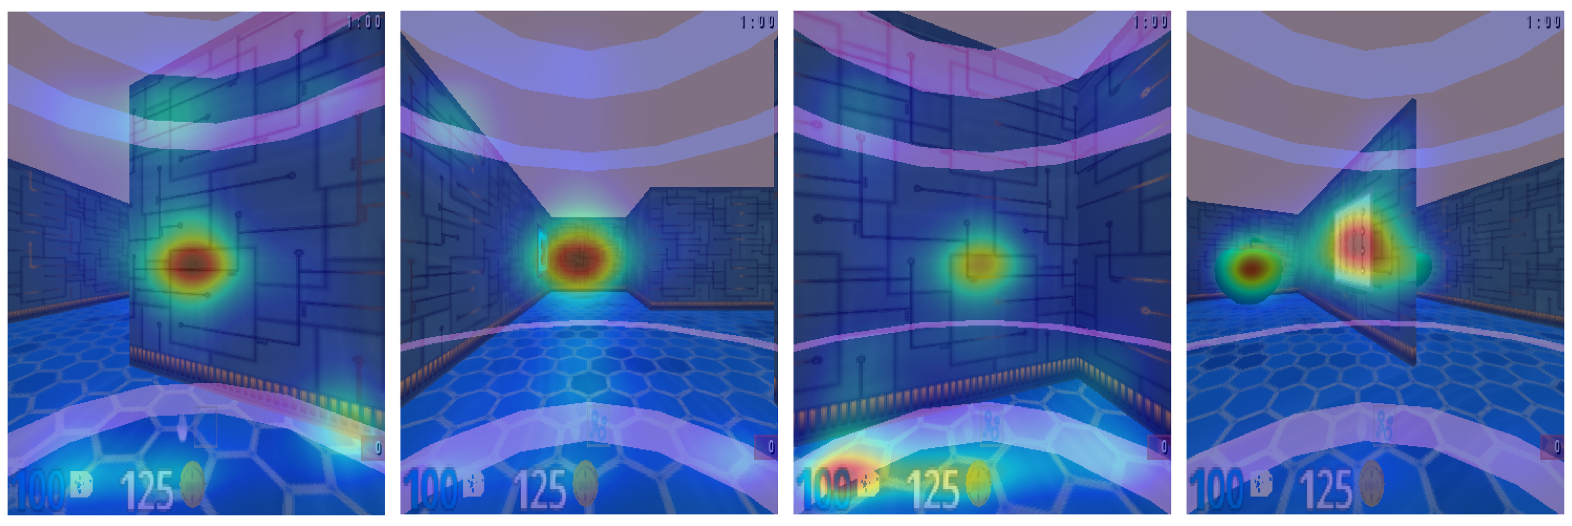

- **Hypothesis:** Knowing where to look in a task aids continual learning across tasks.
- Where do humans look while navigating in a $3D$ maze environment? Does foveating around the regions where humans look helps the reinforcement learning process in the context of continual learning?
- Explore leveraging where humans look in an image as an implicit indication of what is salient for decision making.

Algorithmic Approach

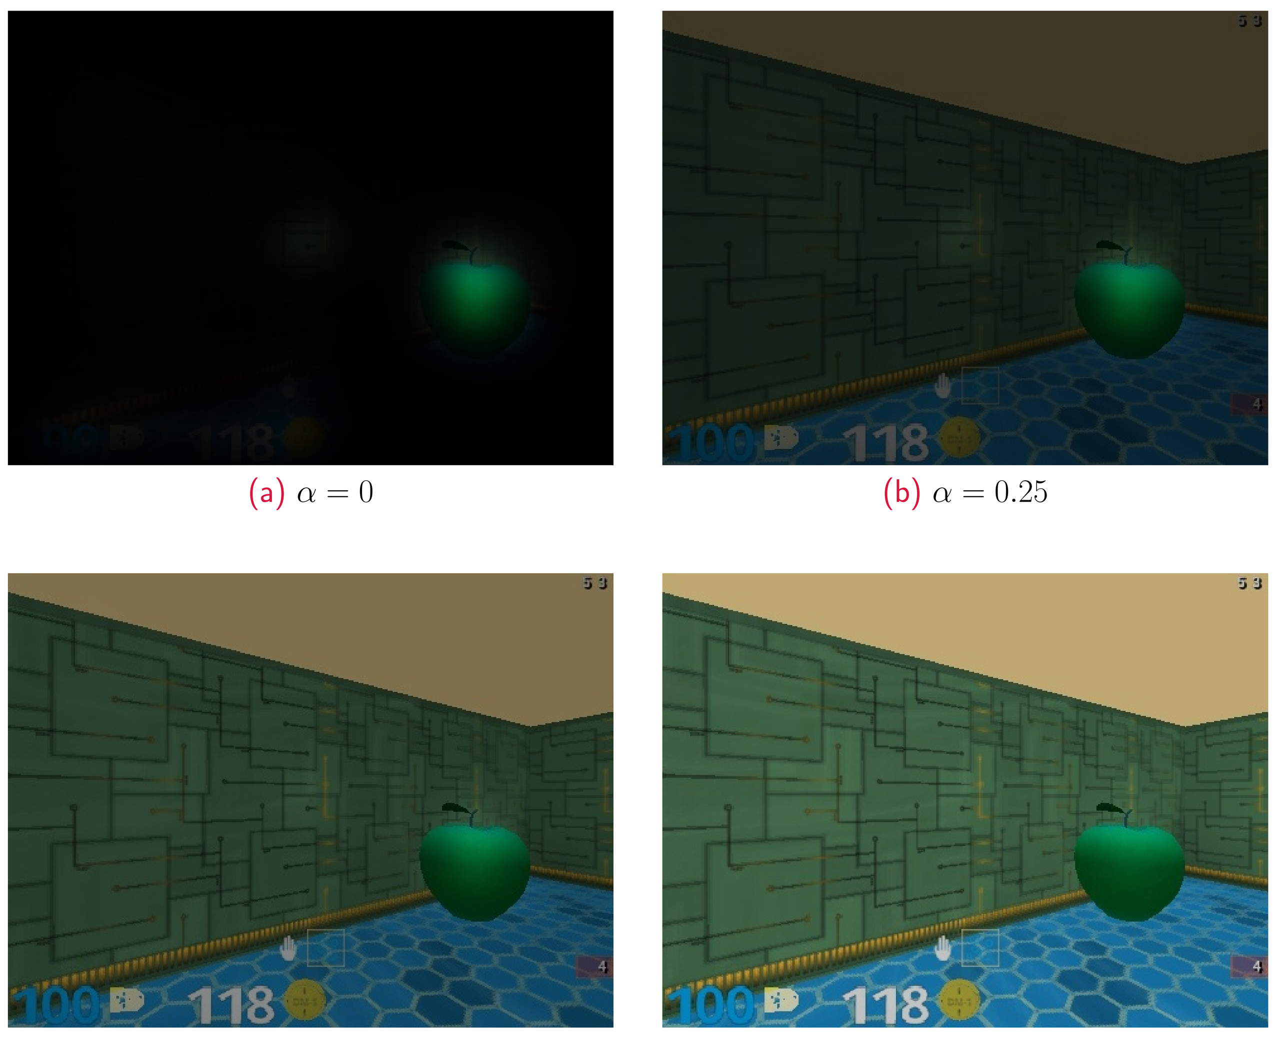

(a) $\alpha = 0$ (b) $\alpha = 0.25$

(c) $\alpha = 0.50$ (d) $\alpha = 0.75$

$$\boxed{I(x,y) = I(x,y) + \Big(S(x,y) + \alpha(1 - S(x,y))\Big)} \quad (1)$$

Visually Attentive UNREAL Agent

Overview

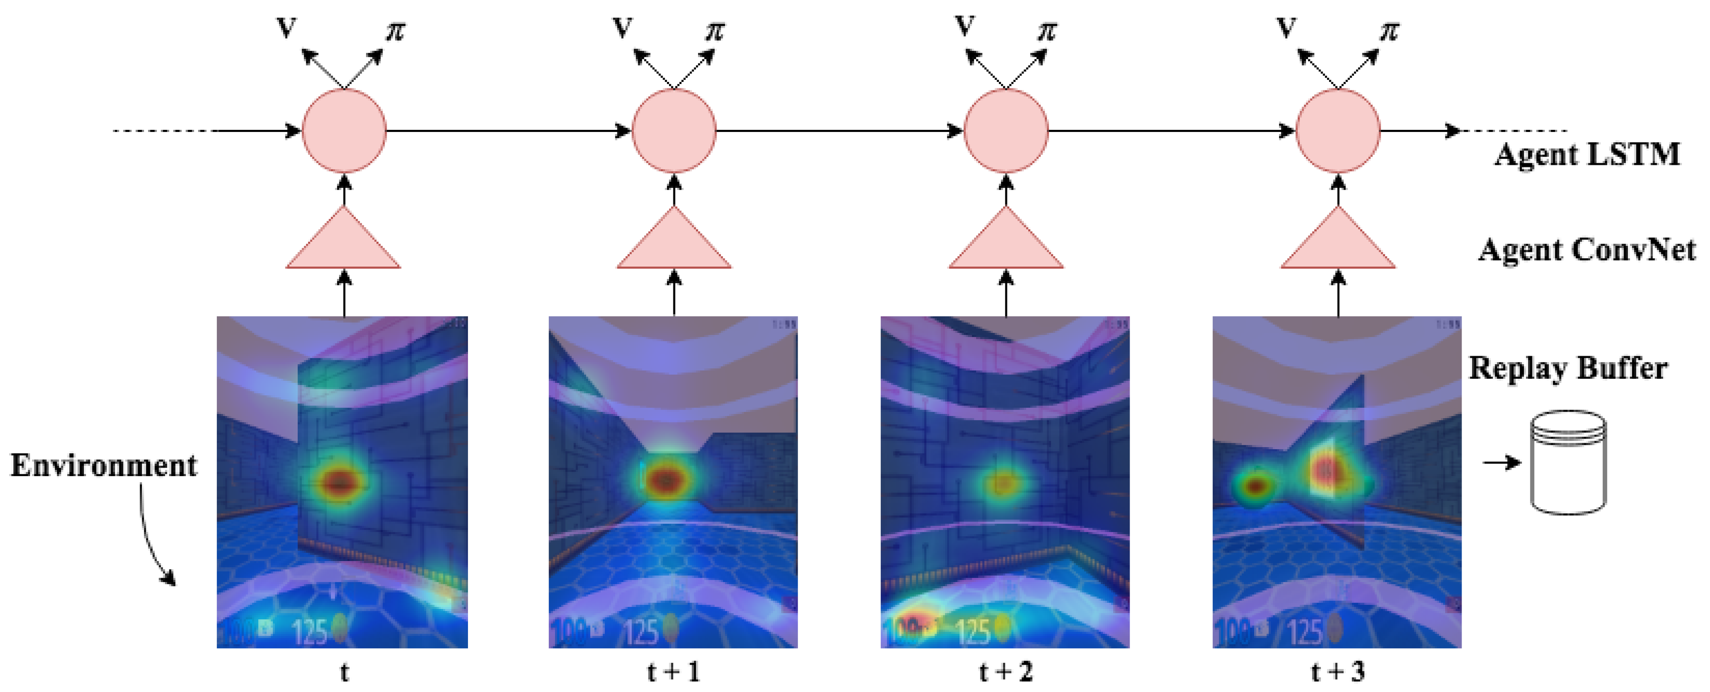

Figure: **Base A3C** is a CNN-LSTM where the agent foveates around where humans look in the incoming scenes from the environment. The degree of foveation is decided by the parameter α. The network architecture is kept consistent with the UNREAL agent. Experience stored in the replay buffer is used for auxiliary tasks comprising of reward prediction, value function replay, and pixel control. Note: the overlaid heatmap images here are for visualization of attended regions. The agents focuses around salient regions based on different degrees of foveation as shown above.

Key Idea

- **Baseline:** A CNN-LSTM base *UNREAL* agent [1] trained on-policy with A3C [2] loss.
- **Our approach:** Introduce the Visually-Attentive UNREAL agent by foveating around the salient regions in each image.
- **Saliency Method:** Real time Spectral Residual method [3]. Preliminary analysis using pre-trained Saliency Attentive Modelling method (SAM) [4]

Algorithm 1 Visually Attentive UNREAL Agent

α is factor controlling the foveation
$I \leftarrow$ Obtain original Input Image of $360 * 480$ from the Lab environment
$S \leftarrow$ SpectralSaliencyMethod (I)
$FoveatedImage \leftarrow$ SaliencyOverlay (I, S, α)
Process Base A3C CNN-LSTM (Foveated Image)
Process Auxiliary Tasks (Foveated Image)

Experiments

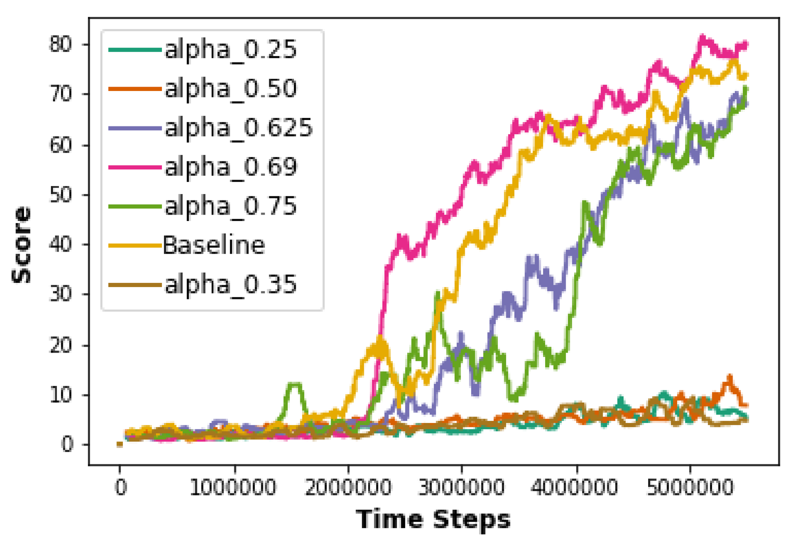

Figure: **Learning with varying degrees of visual attention** to navigate a $3D$ maze environment. Here $\alpha = 0.69$ speeds up the learning as compared to other settings for this instance of runs.

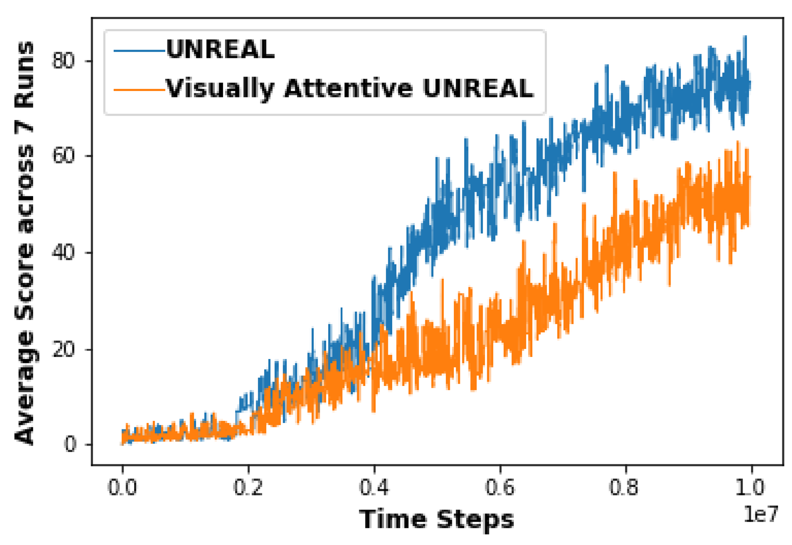

Figure: **Learning curves** in the $3D$ Navigation Maze Static. On an average, UNREAL agent learns better than Visually-Attentive UNREAL agent during the training phase.

Transfer Learning

- To evaluate the trained models for continual learning, we introduce three types of perturbations in the input frames namely; Gaussian noise, tinting of images at random with the same hue, and tinting of images at random with different hues

Agent	Testing	Continual Learning		
		Easy	Moderate	Difficult
UNREAL	96.92(8.08)	101.96(9.656)	92.64(12.35)	39.16(11.44)
Visually-Attentive UNREAL	95.92(10.88)	96.96(9.39)	83.52(10.09)	40.52(14.67)

Table: **Average Performance** over $k = 25$ games once training is completed. UNREAL agent and Visually-Attentive UNREAL agent are evaluated once training is stopped for transfer learning. Scores here are averaged for 25 games with standard deviation across these games in the brackets.

Wrap up

- Upon encountering flickering in frames at random,the **Visually-Attentive UNREAL agent is still able to perform as well as the baseline and is relatively more robust to distractors** in both easy and moderate categories of evaluation.
- Our approach can be used as a wrapper around *any* saliency model, so it would be easy to try better approaches.
- **Future work** should be done to study a setting where the agent can actively learn to control where to attend, rather than using a static attention model.

References

- Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937, 2016.
- Xiaodi Hou and Liqing Zhang. Saliency detection: A spectral residual approach. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007.
- Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. Predicting human eye fixations via an lstm-based saliency attentive model. *arXiv preprint arXiv:1611.09571*, 2016.

Scan me