Attend Before you Act
Leveraging human visual attention for continual learning

Khimya Khetarpal, Doina Precup

Mila - McGill University

Motivation

Where do humans look in an image?

Where do humans look in an image?

Where do we look in an image?

Why does it matter?

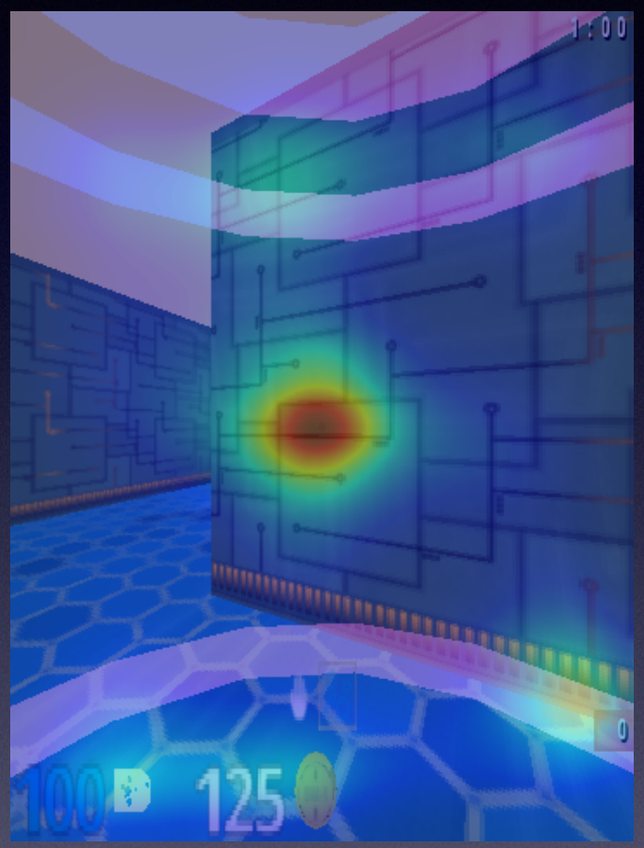

Time

Why does it matter?

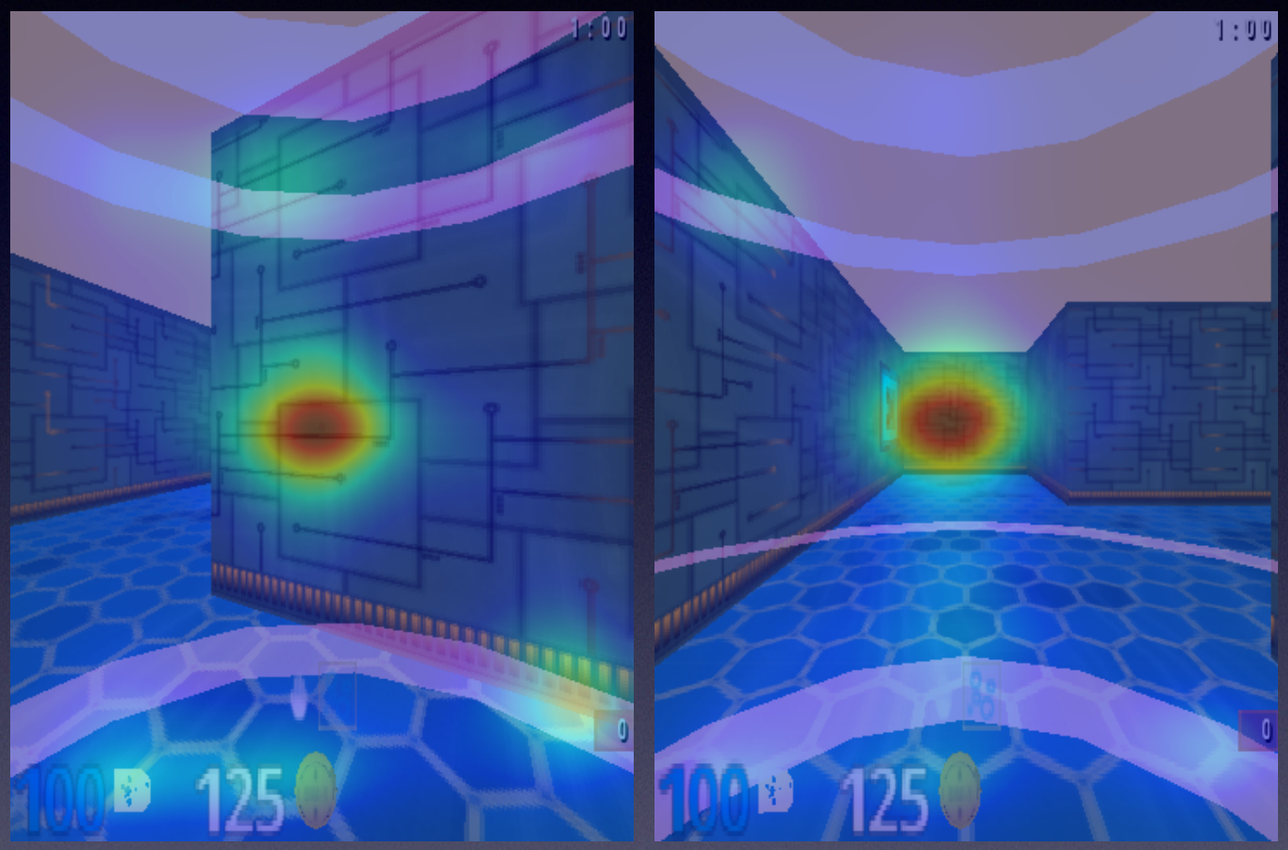

Time

Why does it matter?

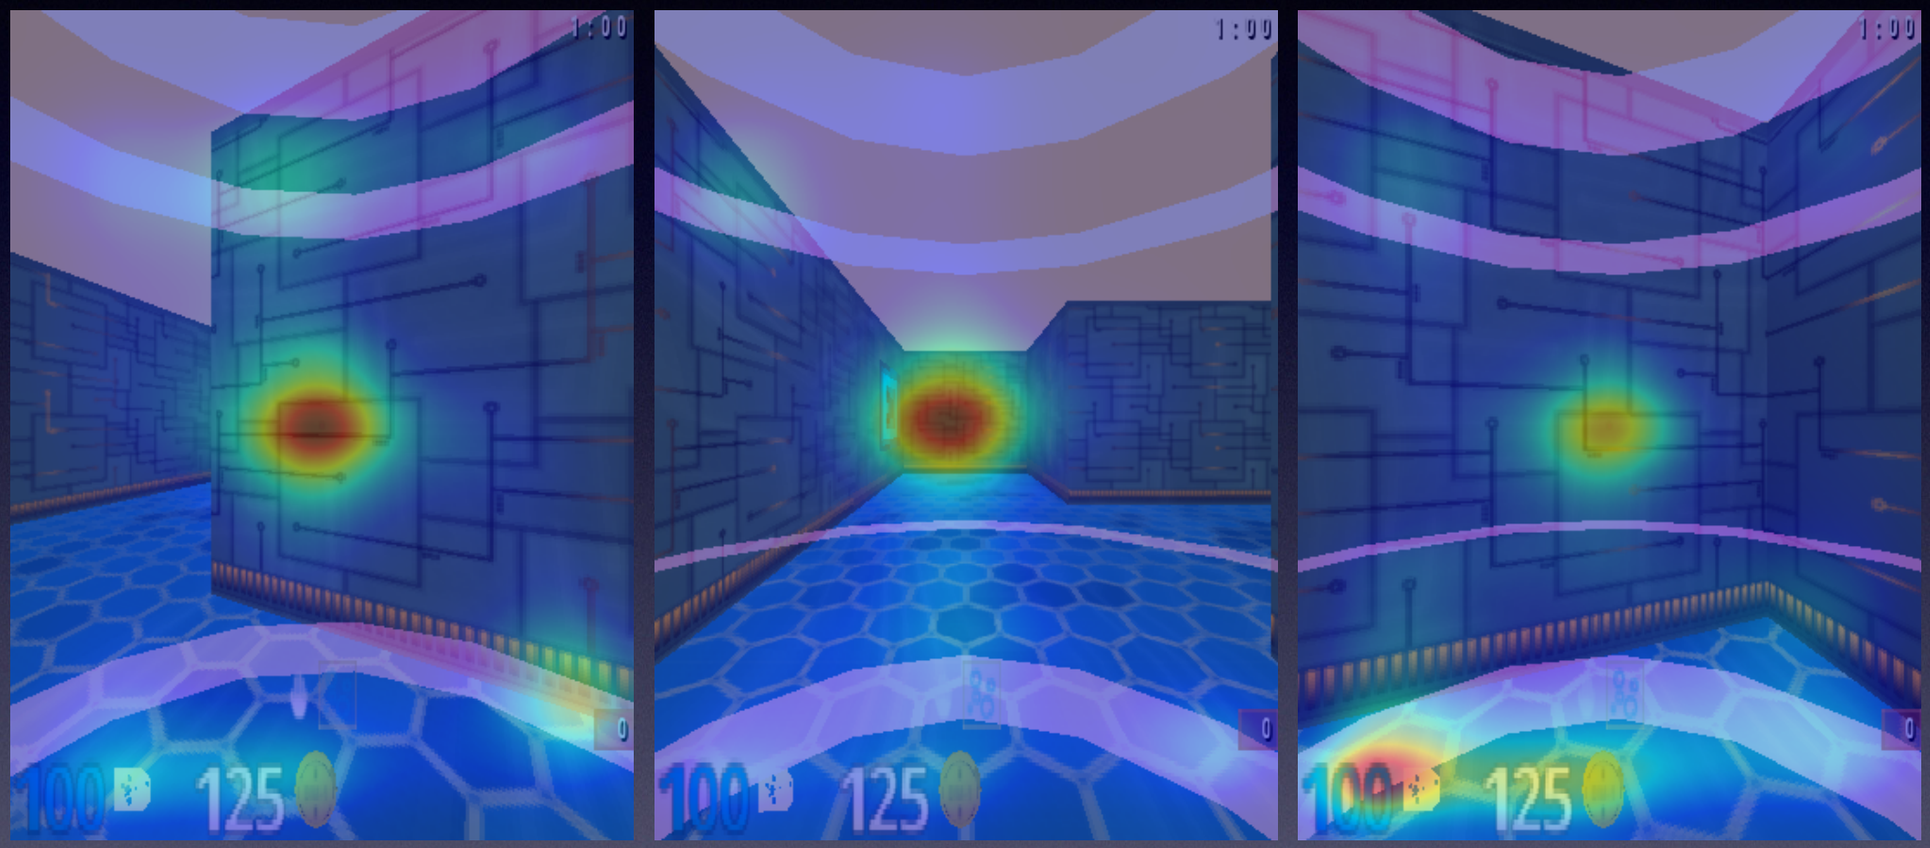

Time

Why does it matter?

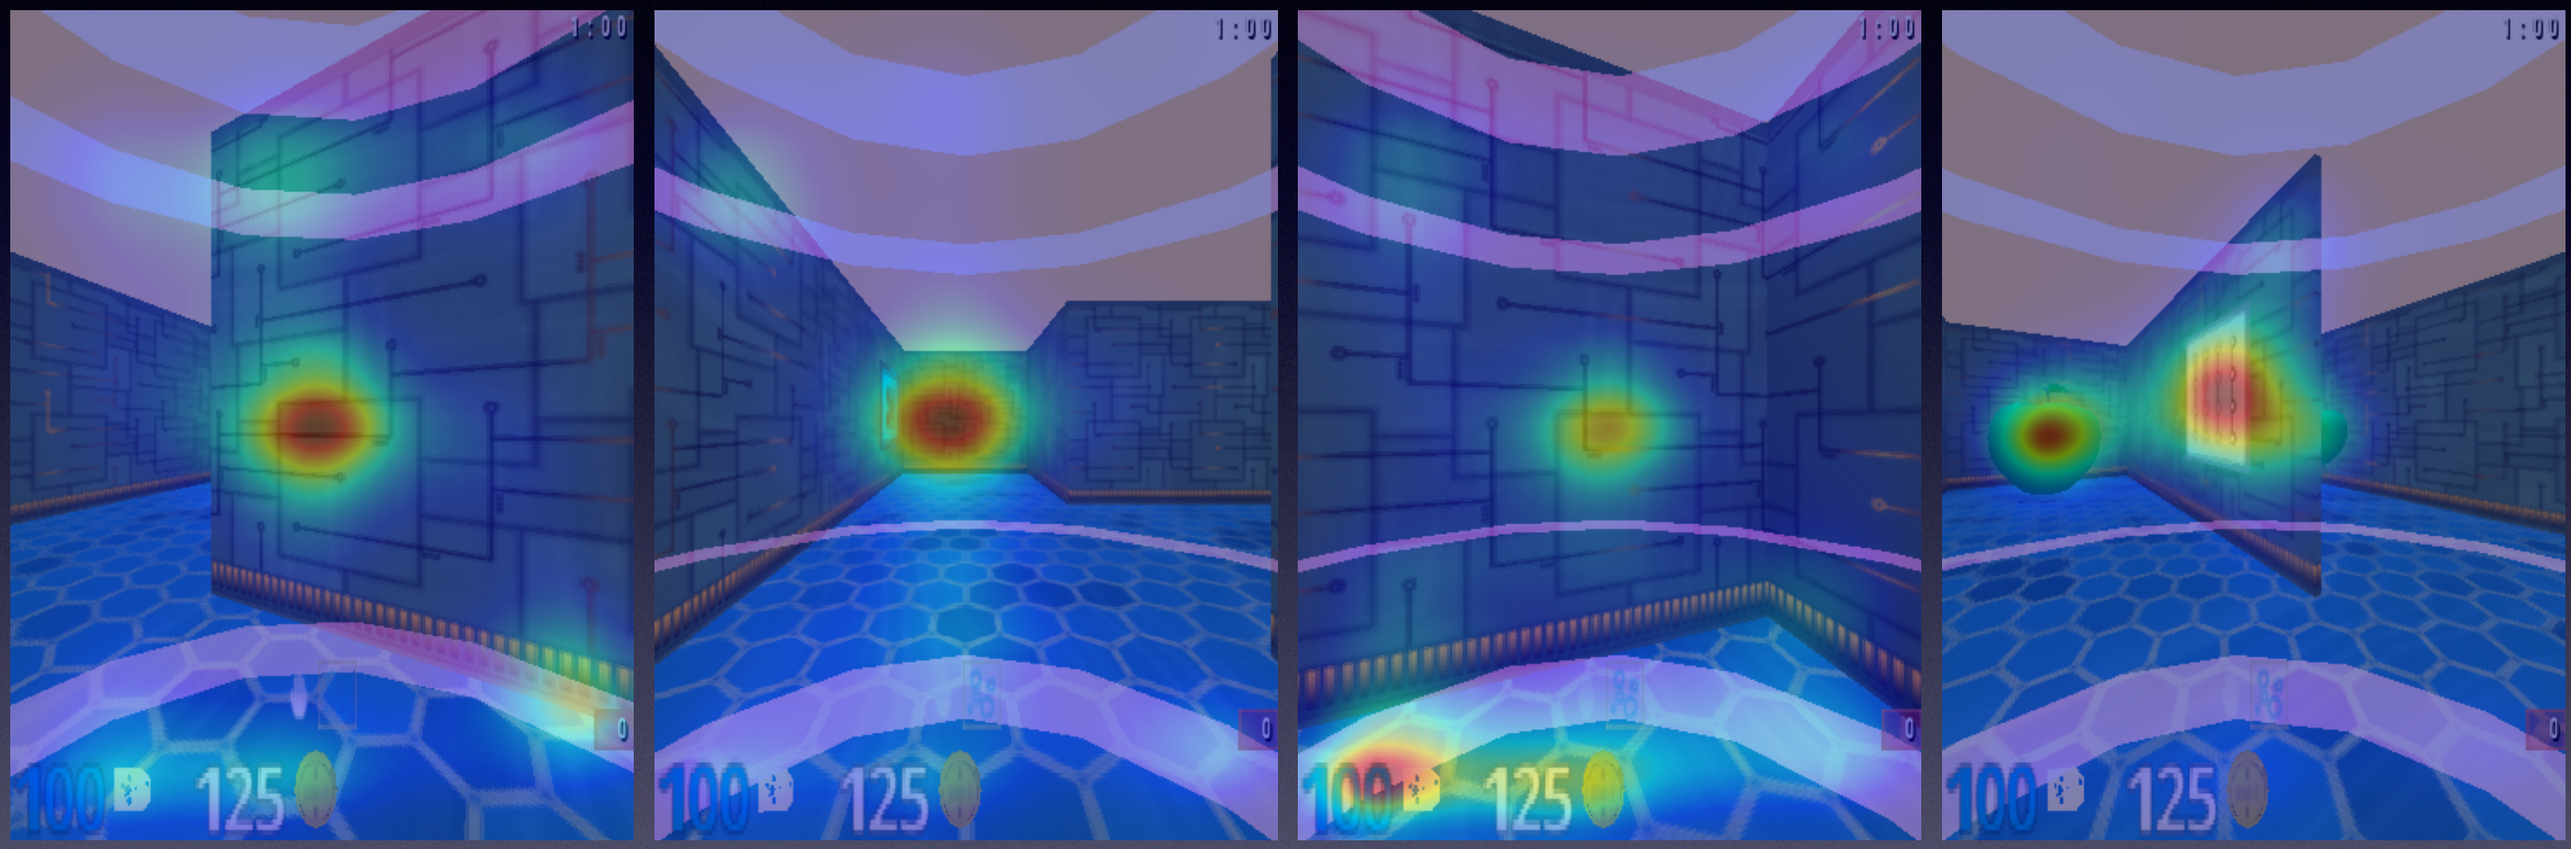

Time

This Work

- Does foveating around the regions where humans look helps the reinforcement learning process in the context of continual learning?

This Work

- 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.

This Work

- 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.

Hypothesize: *Knowing where to look in a task aids continual learning*

Varying Degrees of Foveation

Varying Degrees of Foveation

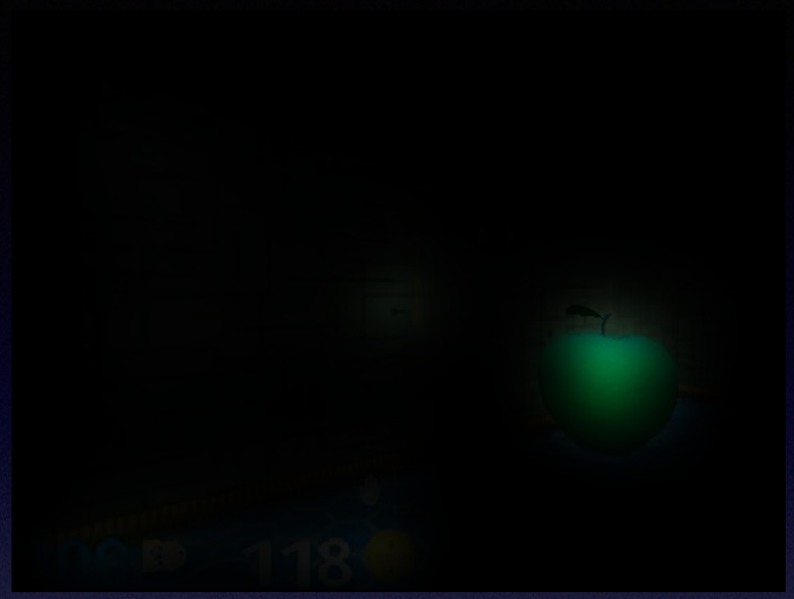

Varying Degrees of Foveation

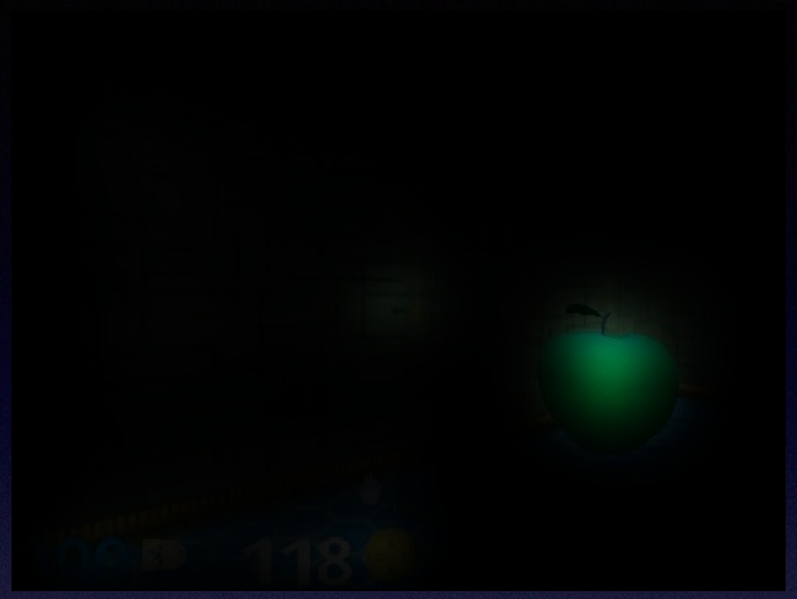

Varying Degrees of Foveation

Varying Degrees of Foveation

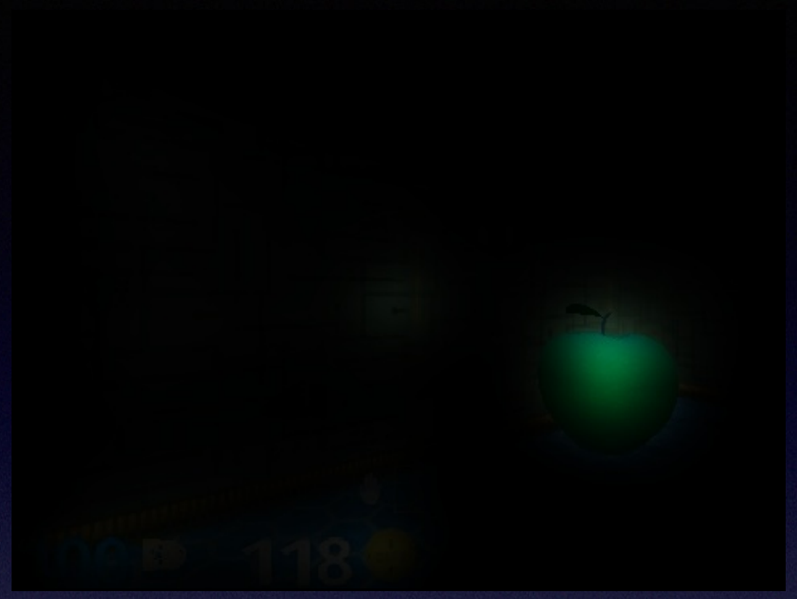

$$I(x,y) = I(x,y) + \Big(S(x,y) + \alpha(1 - S(x,y))\Big)$$

Baseline

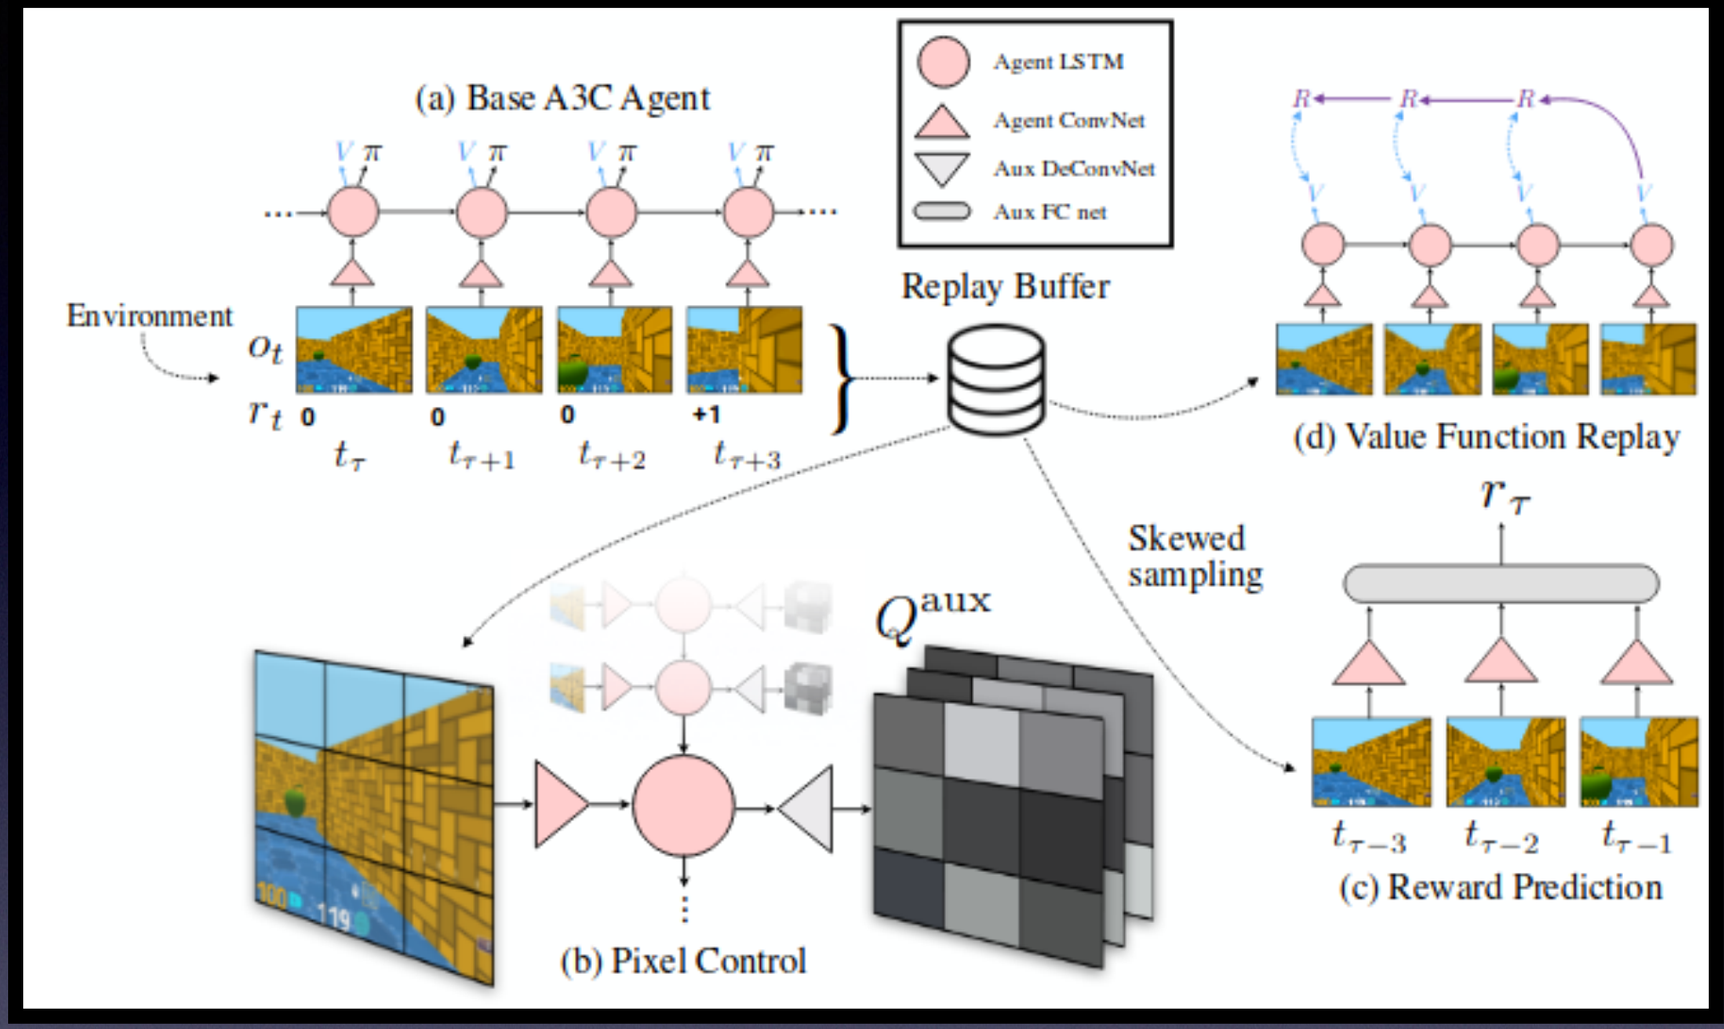

- **Baseline:** UNREAL agent [Jaderberg et al., 2016] [https://github.com/miyosuda/unreal]

Visually Attentive UNREAL agent

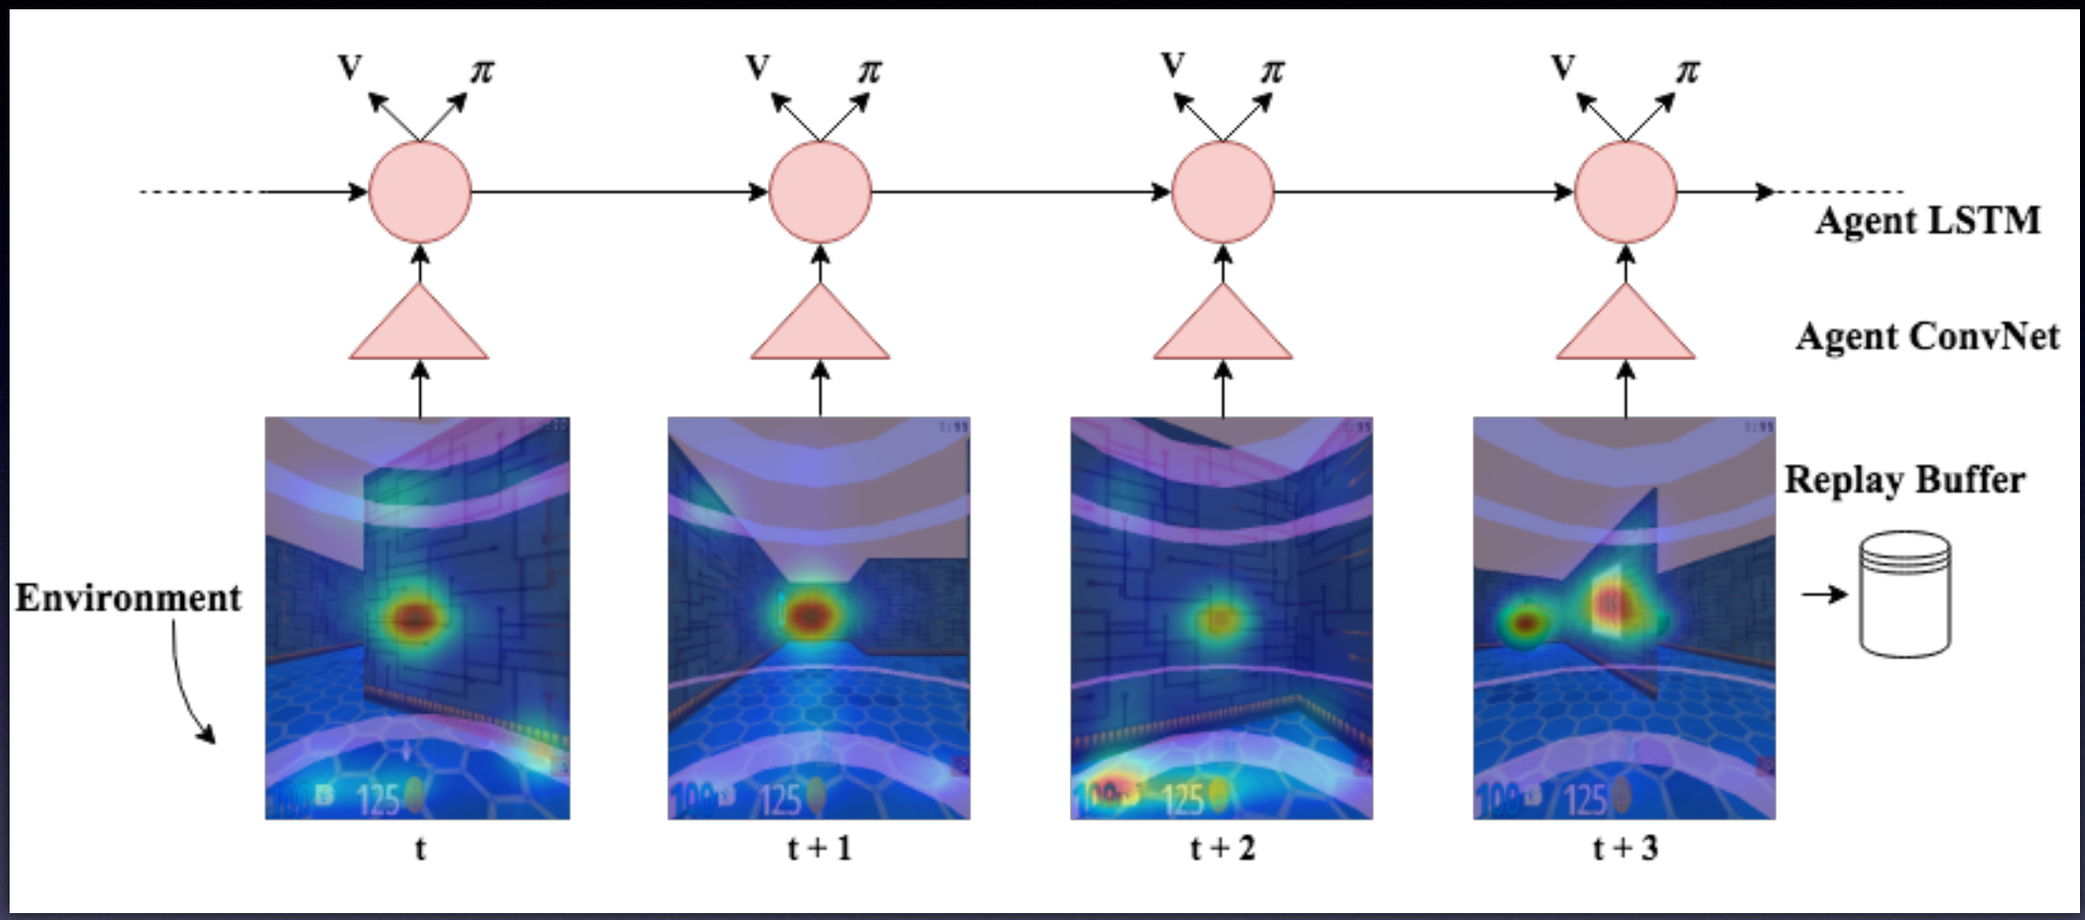

- The **Visually-Attentive UNREAL agent** attends around the salient regions in each image and then acts
- Real time Spectral Residual method [Hou & Zhang, 2007] for generating salient regions in frames

Visually Attentive UNREAL agent

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)

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

Learning with varying degrees of foveation

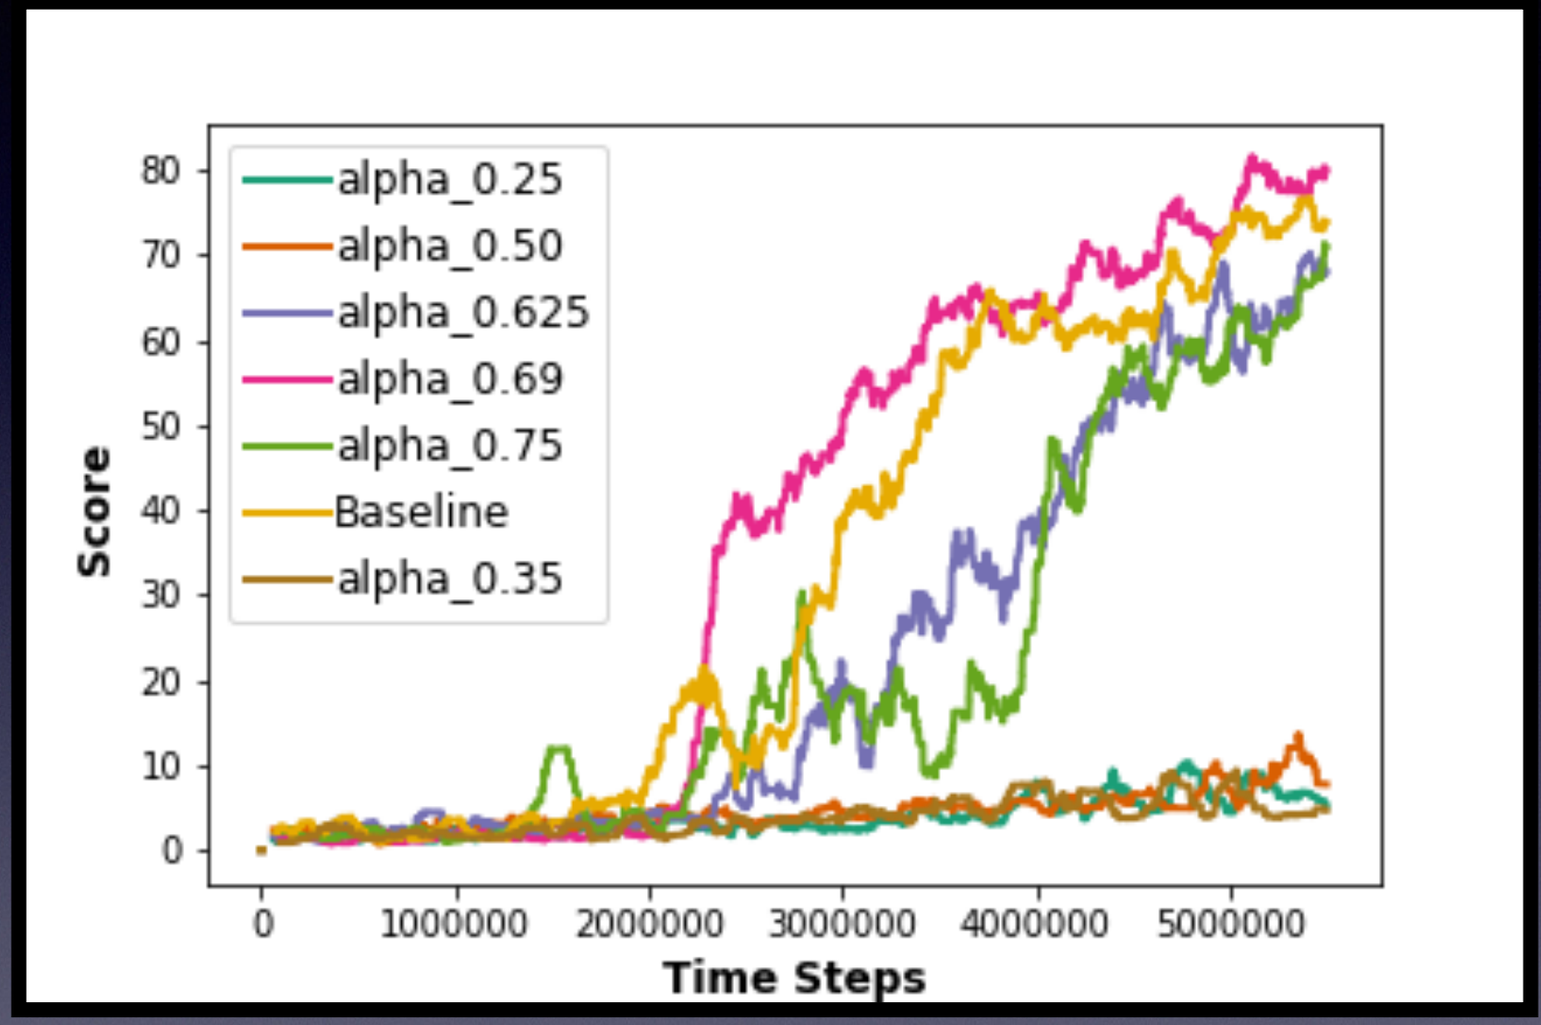

Best value of alpha vs baseline

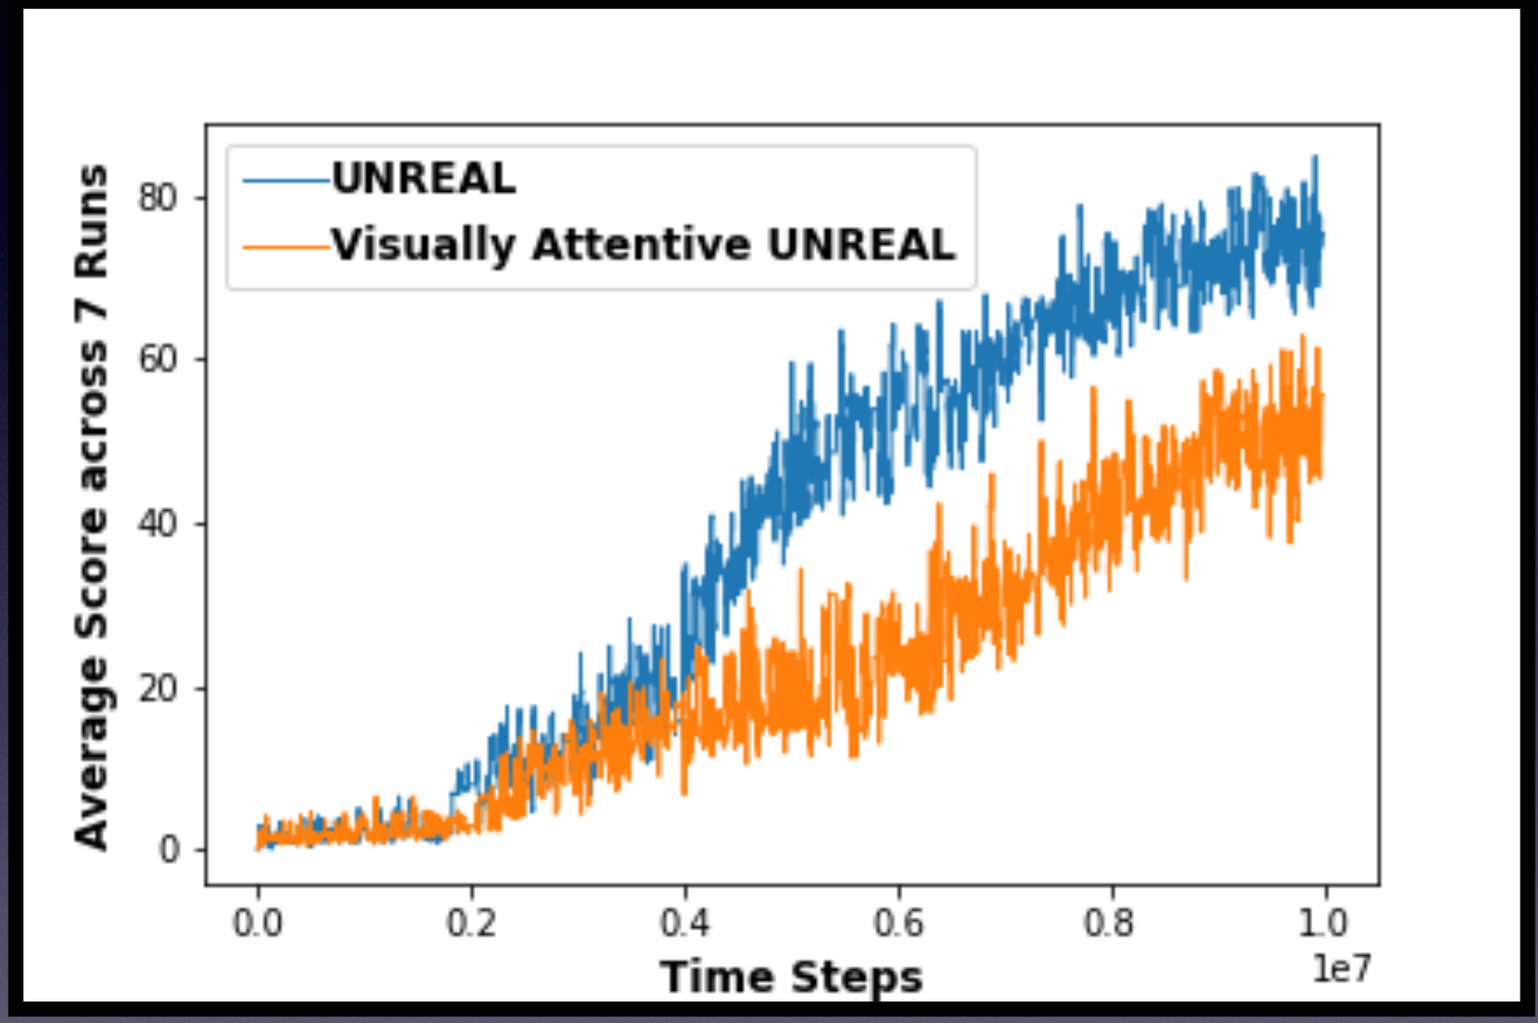

Continual Learning

Average performance over 25 games				
Agent	**Testing**	**Continual Learning**		
		Easy	**Moderate**	**Difficult**
UNREAL	96.92 (8.08)	101.96 (**9.65**)	92.64 (**12.35**)	39.16 (11.14)
Visually Attentive UNREAL Agent	95.92 (10.88)	96.96 (**9.39**)	83.52 (**10.09**)	40.52 (14.67)

- **Easy:** Gaussian noise
- **Moderate:** Tinting of images at random with same hue of 0.25
- **Difficult:** Tinting of images at random with different hue for each tinted frame

Continual Learning - Moderate

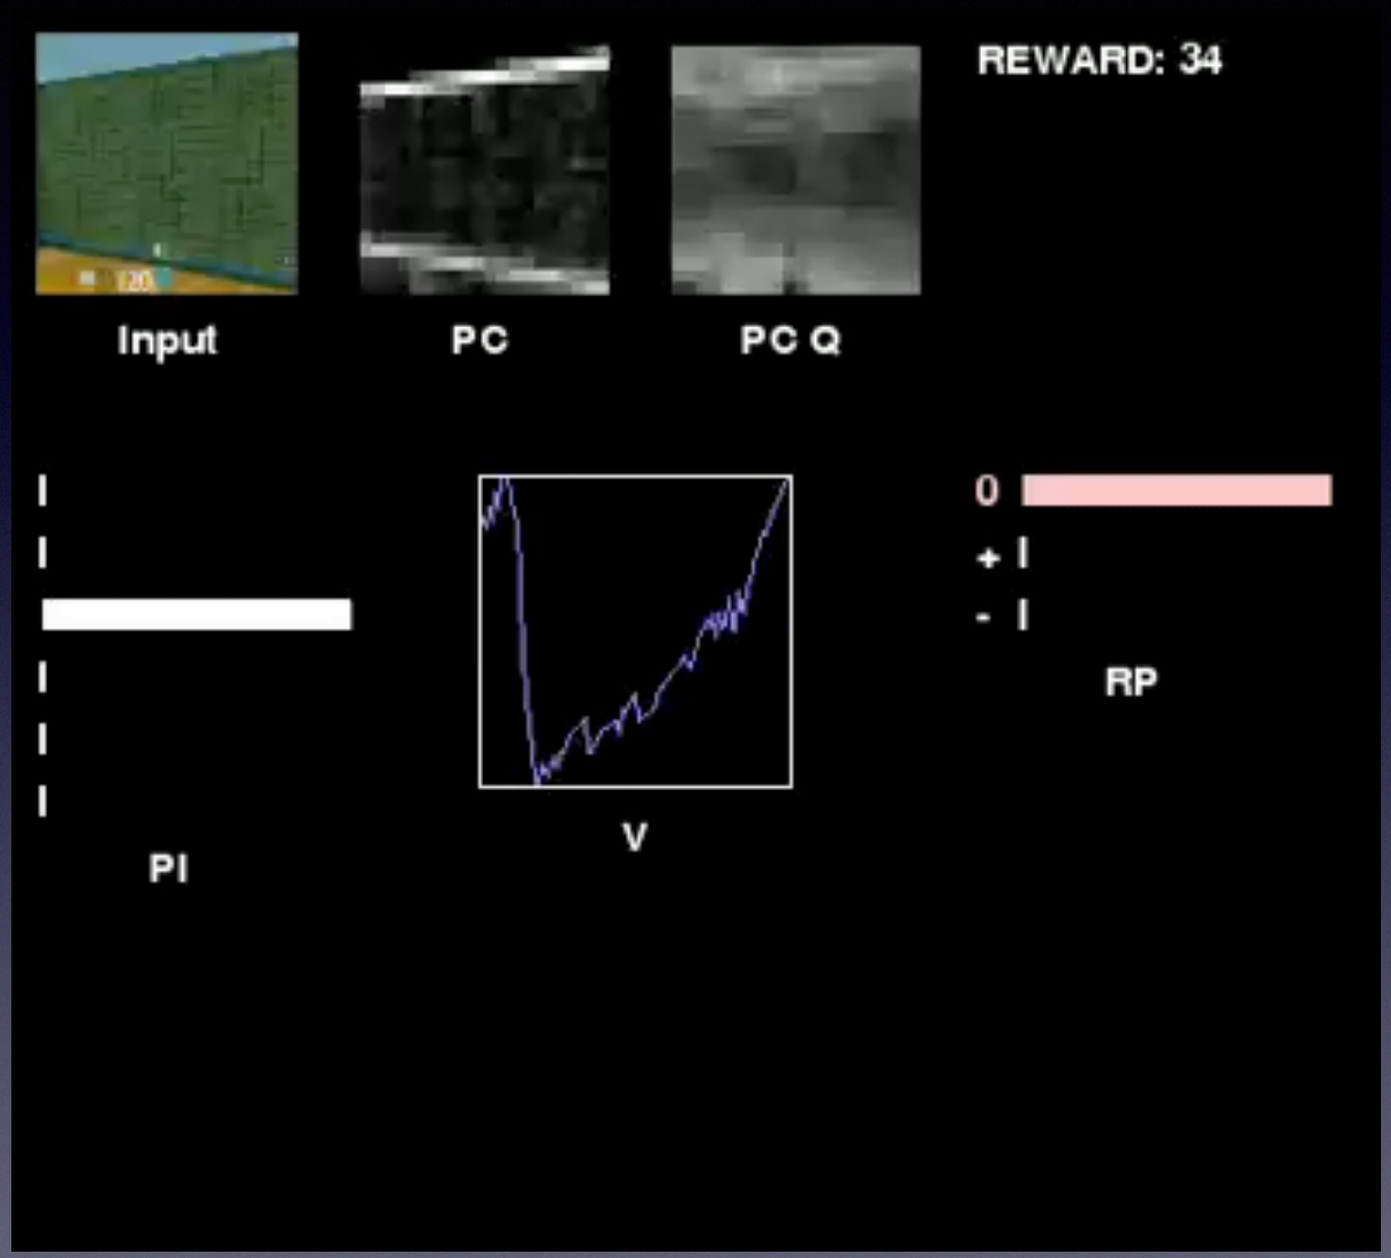

Continual Learning - Difficult

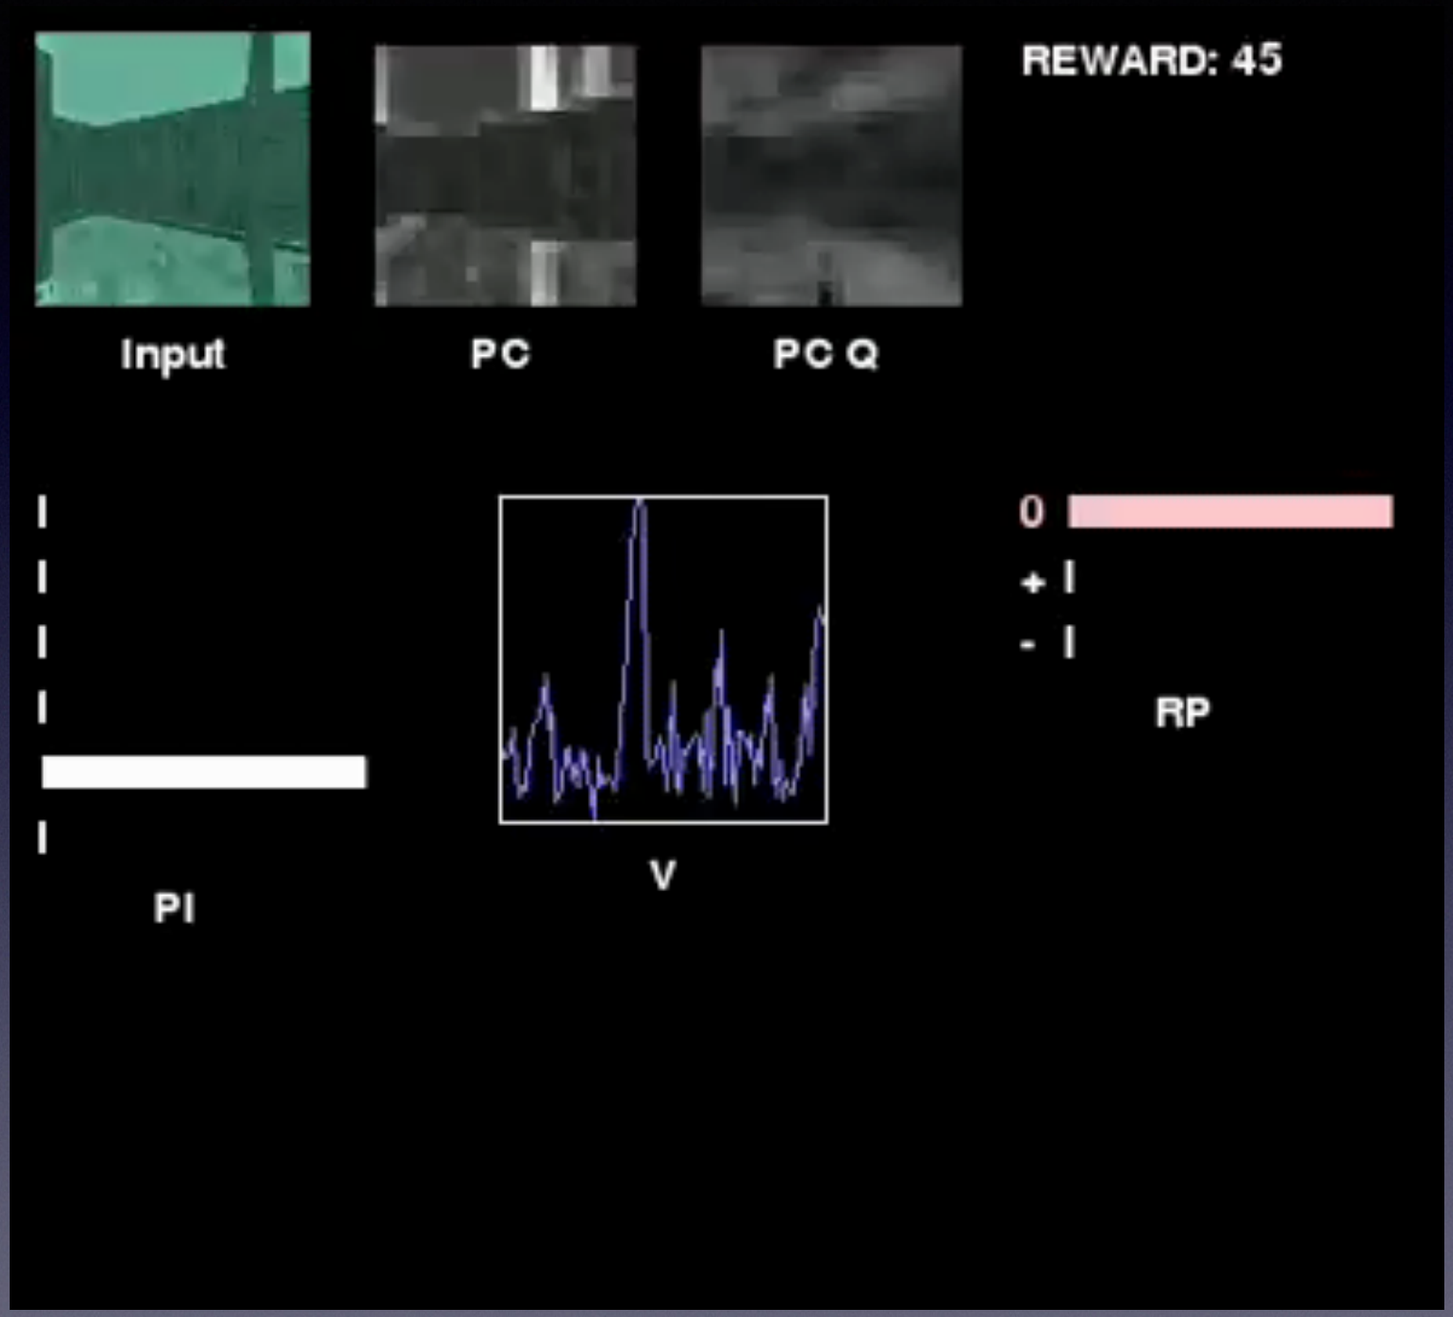

Findings

- *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.

Findings

- *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

Findings

- *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

- The performance evaluation on perturbations in the train setting demonstrate promising results for further analysis of continual learning with visual attention.

Human visual attention could be interpreted as

- A source of ***intrinsic motivation*** and ***curiosity***. [Ref: what we see is what we need, Kulkarni et al.]
- A source of ***subgoals*** humans target while at a complex task.

Future Work

- Study a setting where agents ***actively learn to control where to attend***, rather than using a static attention model.

- Start with an ***initial belief of where humans look across n tasks***, and work towards learning with those limited labels about concepts, relations and world models in an interactive way

- Connections to ***specialization*** in terms of ***states of interest*** in a hierarchical reinforcement learning framework

Recent Related Work

- Greydanus, Sam, et al. "Visualizing and Understanding Atari Agents." *arXiv preprint arXiv:1711.00138* (2017).

- Dubey, Rachit, et al. "Investigating Human Priors for Playing Video Games." *arXiv preprint arXiv:1802.10217* (2018).

Thank you

Questions / Feedback ?

https://sites.google.com/view/attendbeforeyouact

Extra Slides

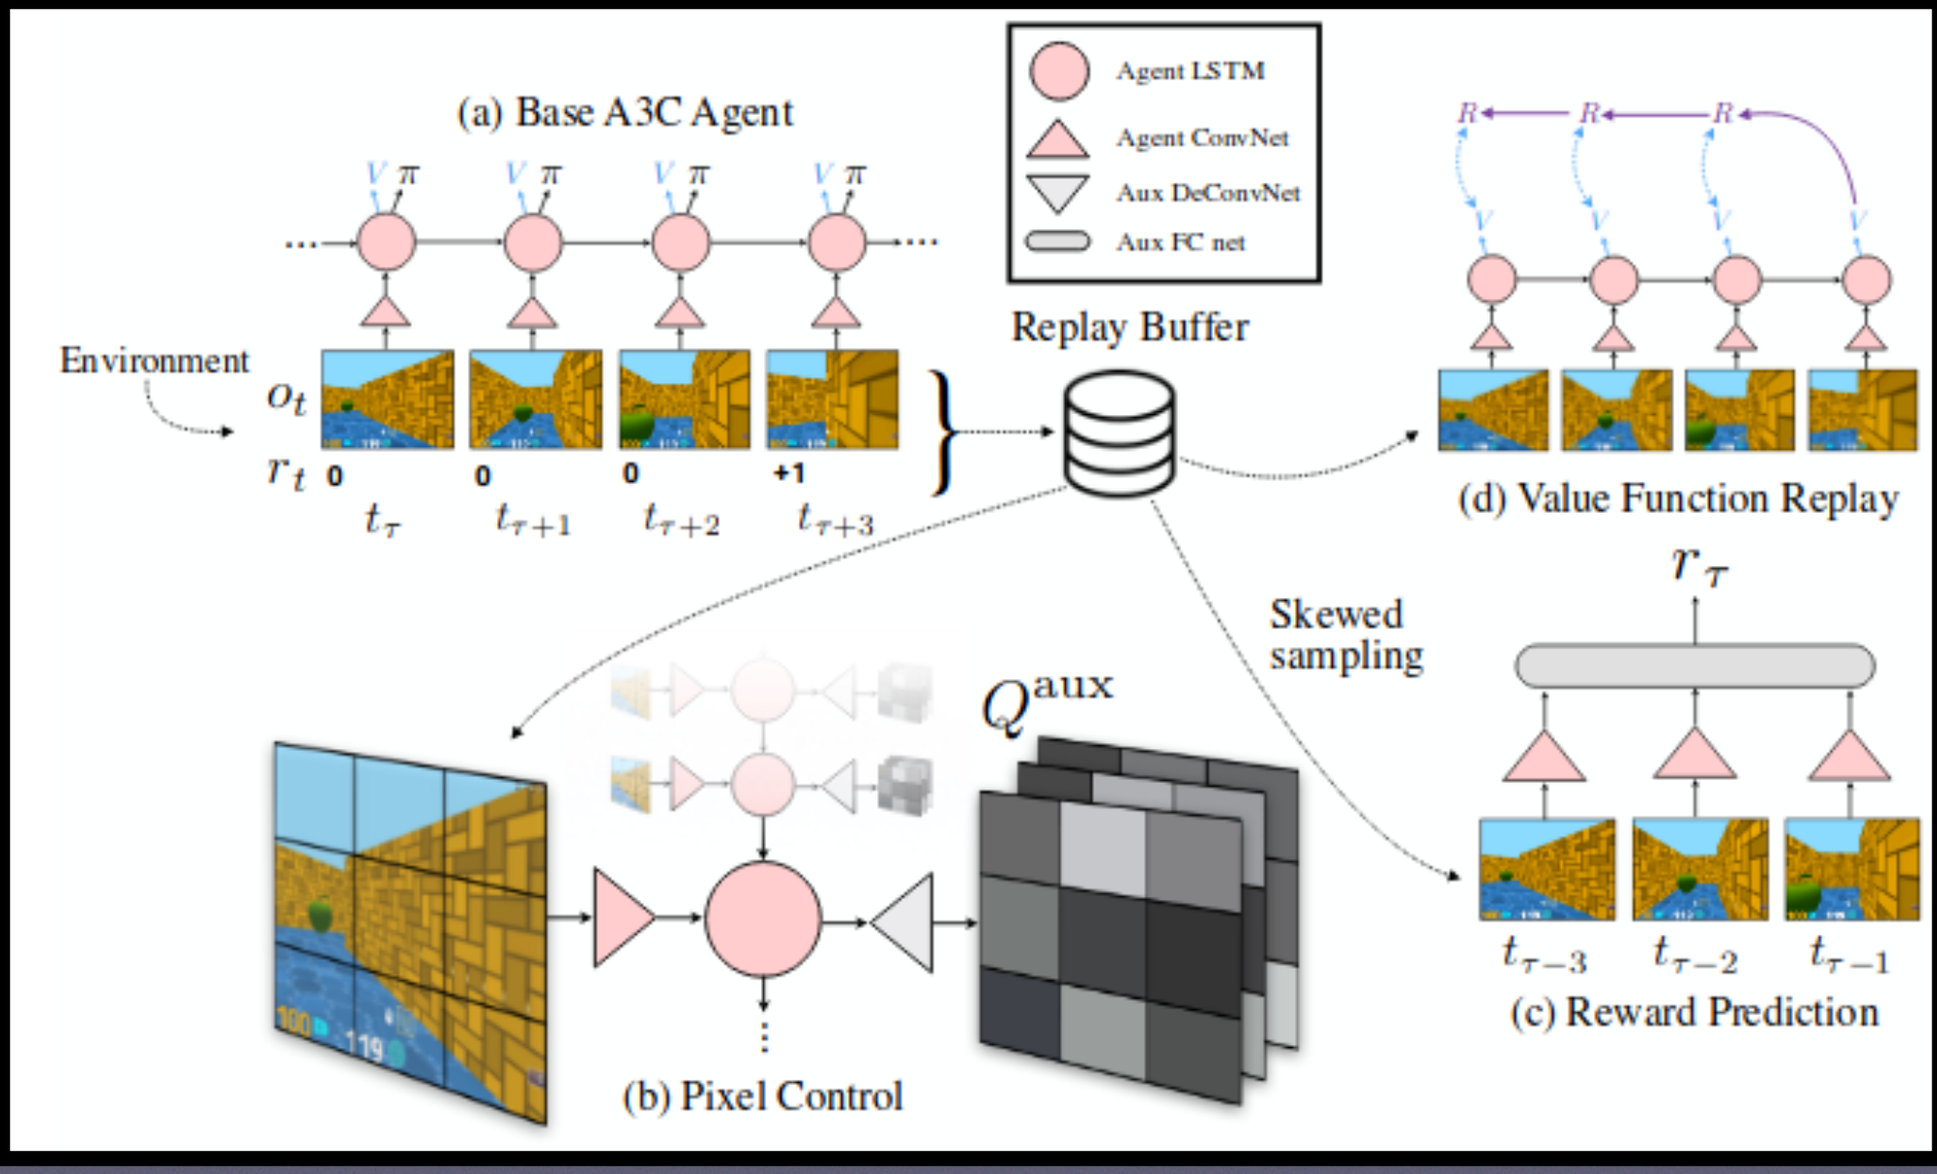

Extra Slides

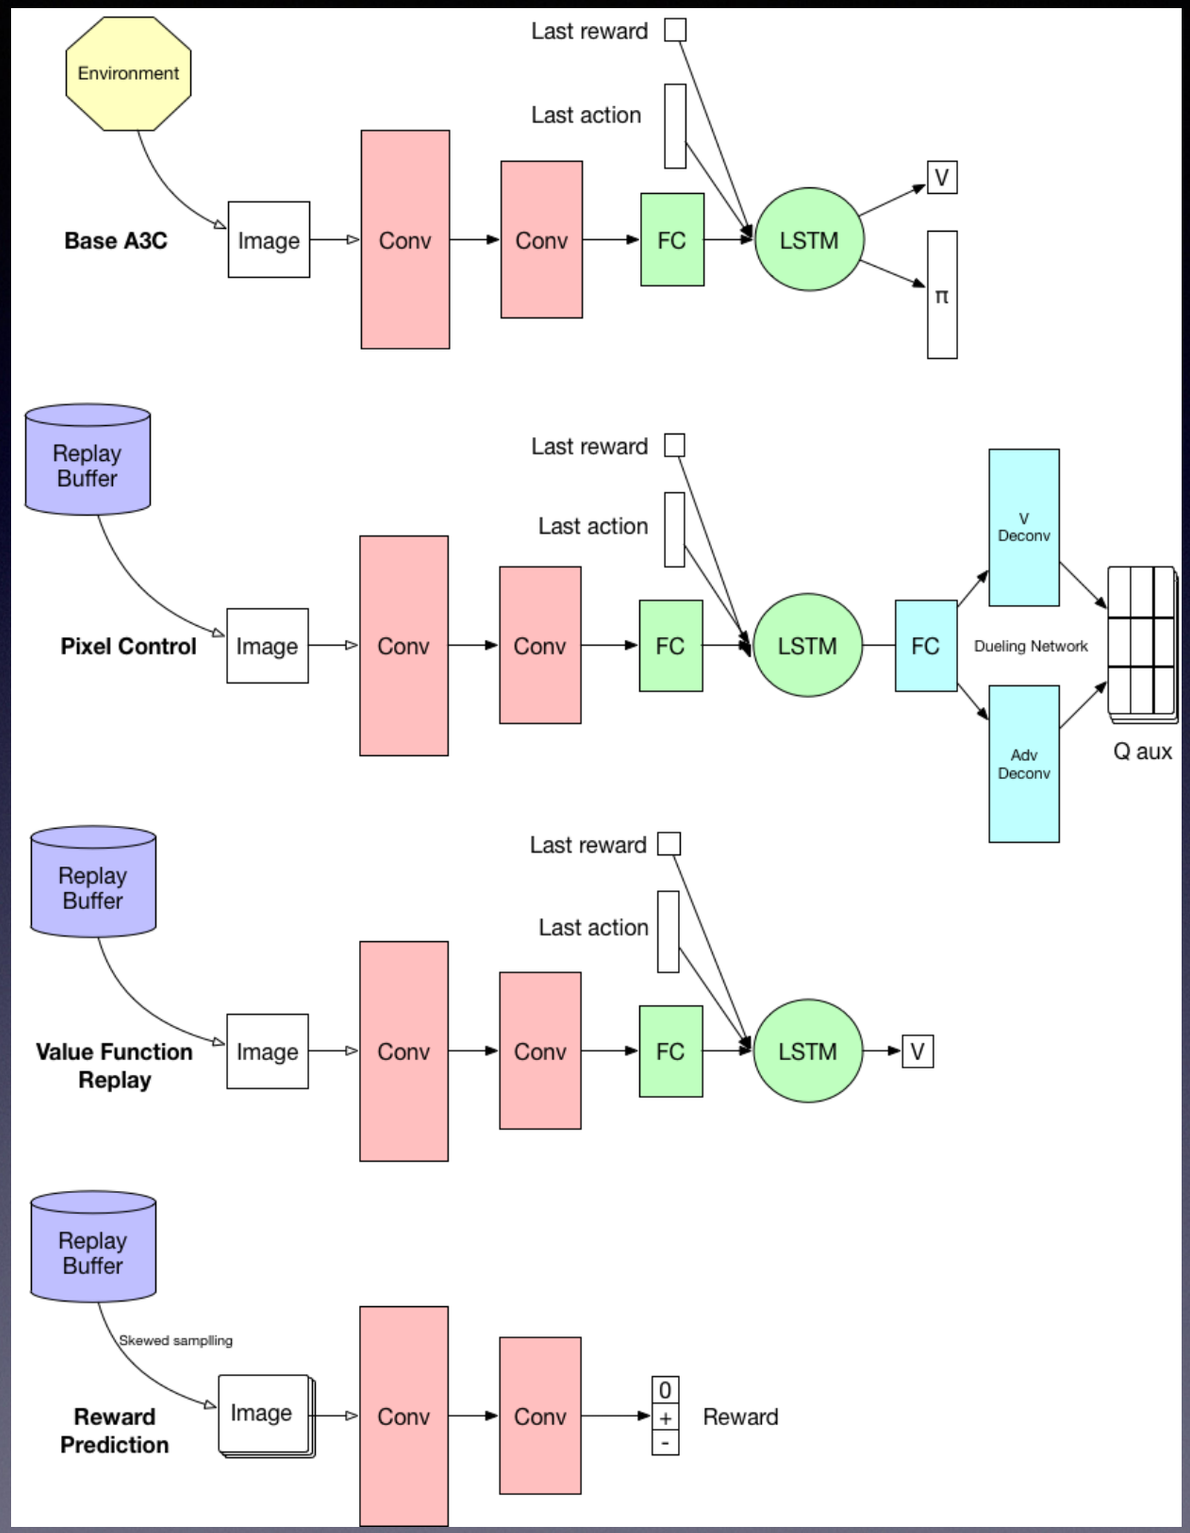

Extra Slides

- Compute the average frequency domain and subtract it from a specific image domain to obtain the spectral residual

- The log spectrum of each image is analyzed to obtain the spectral residual

- This is then transformed to a spatial domain with the location of the porto-objects

- Proto objects are pre-attentive structures with limited spatial and temporal coherence with thin a visual stimuli which generate the perception of an object when attended to