SEED

UPPSALA UNIVERSITET

Improving Generalization in Game Agents with Data Augmentation in Imitation Learning

Derek Yadgaroff, Alessandro Sestini, Konrad Tollmar, Ayca Ozcelikkale, Linus Gisslén

WCCI 2024 - CEC Special Session on Games

Electronic Arts

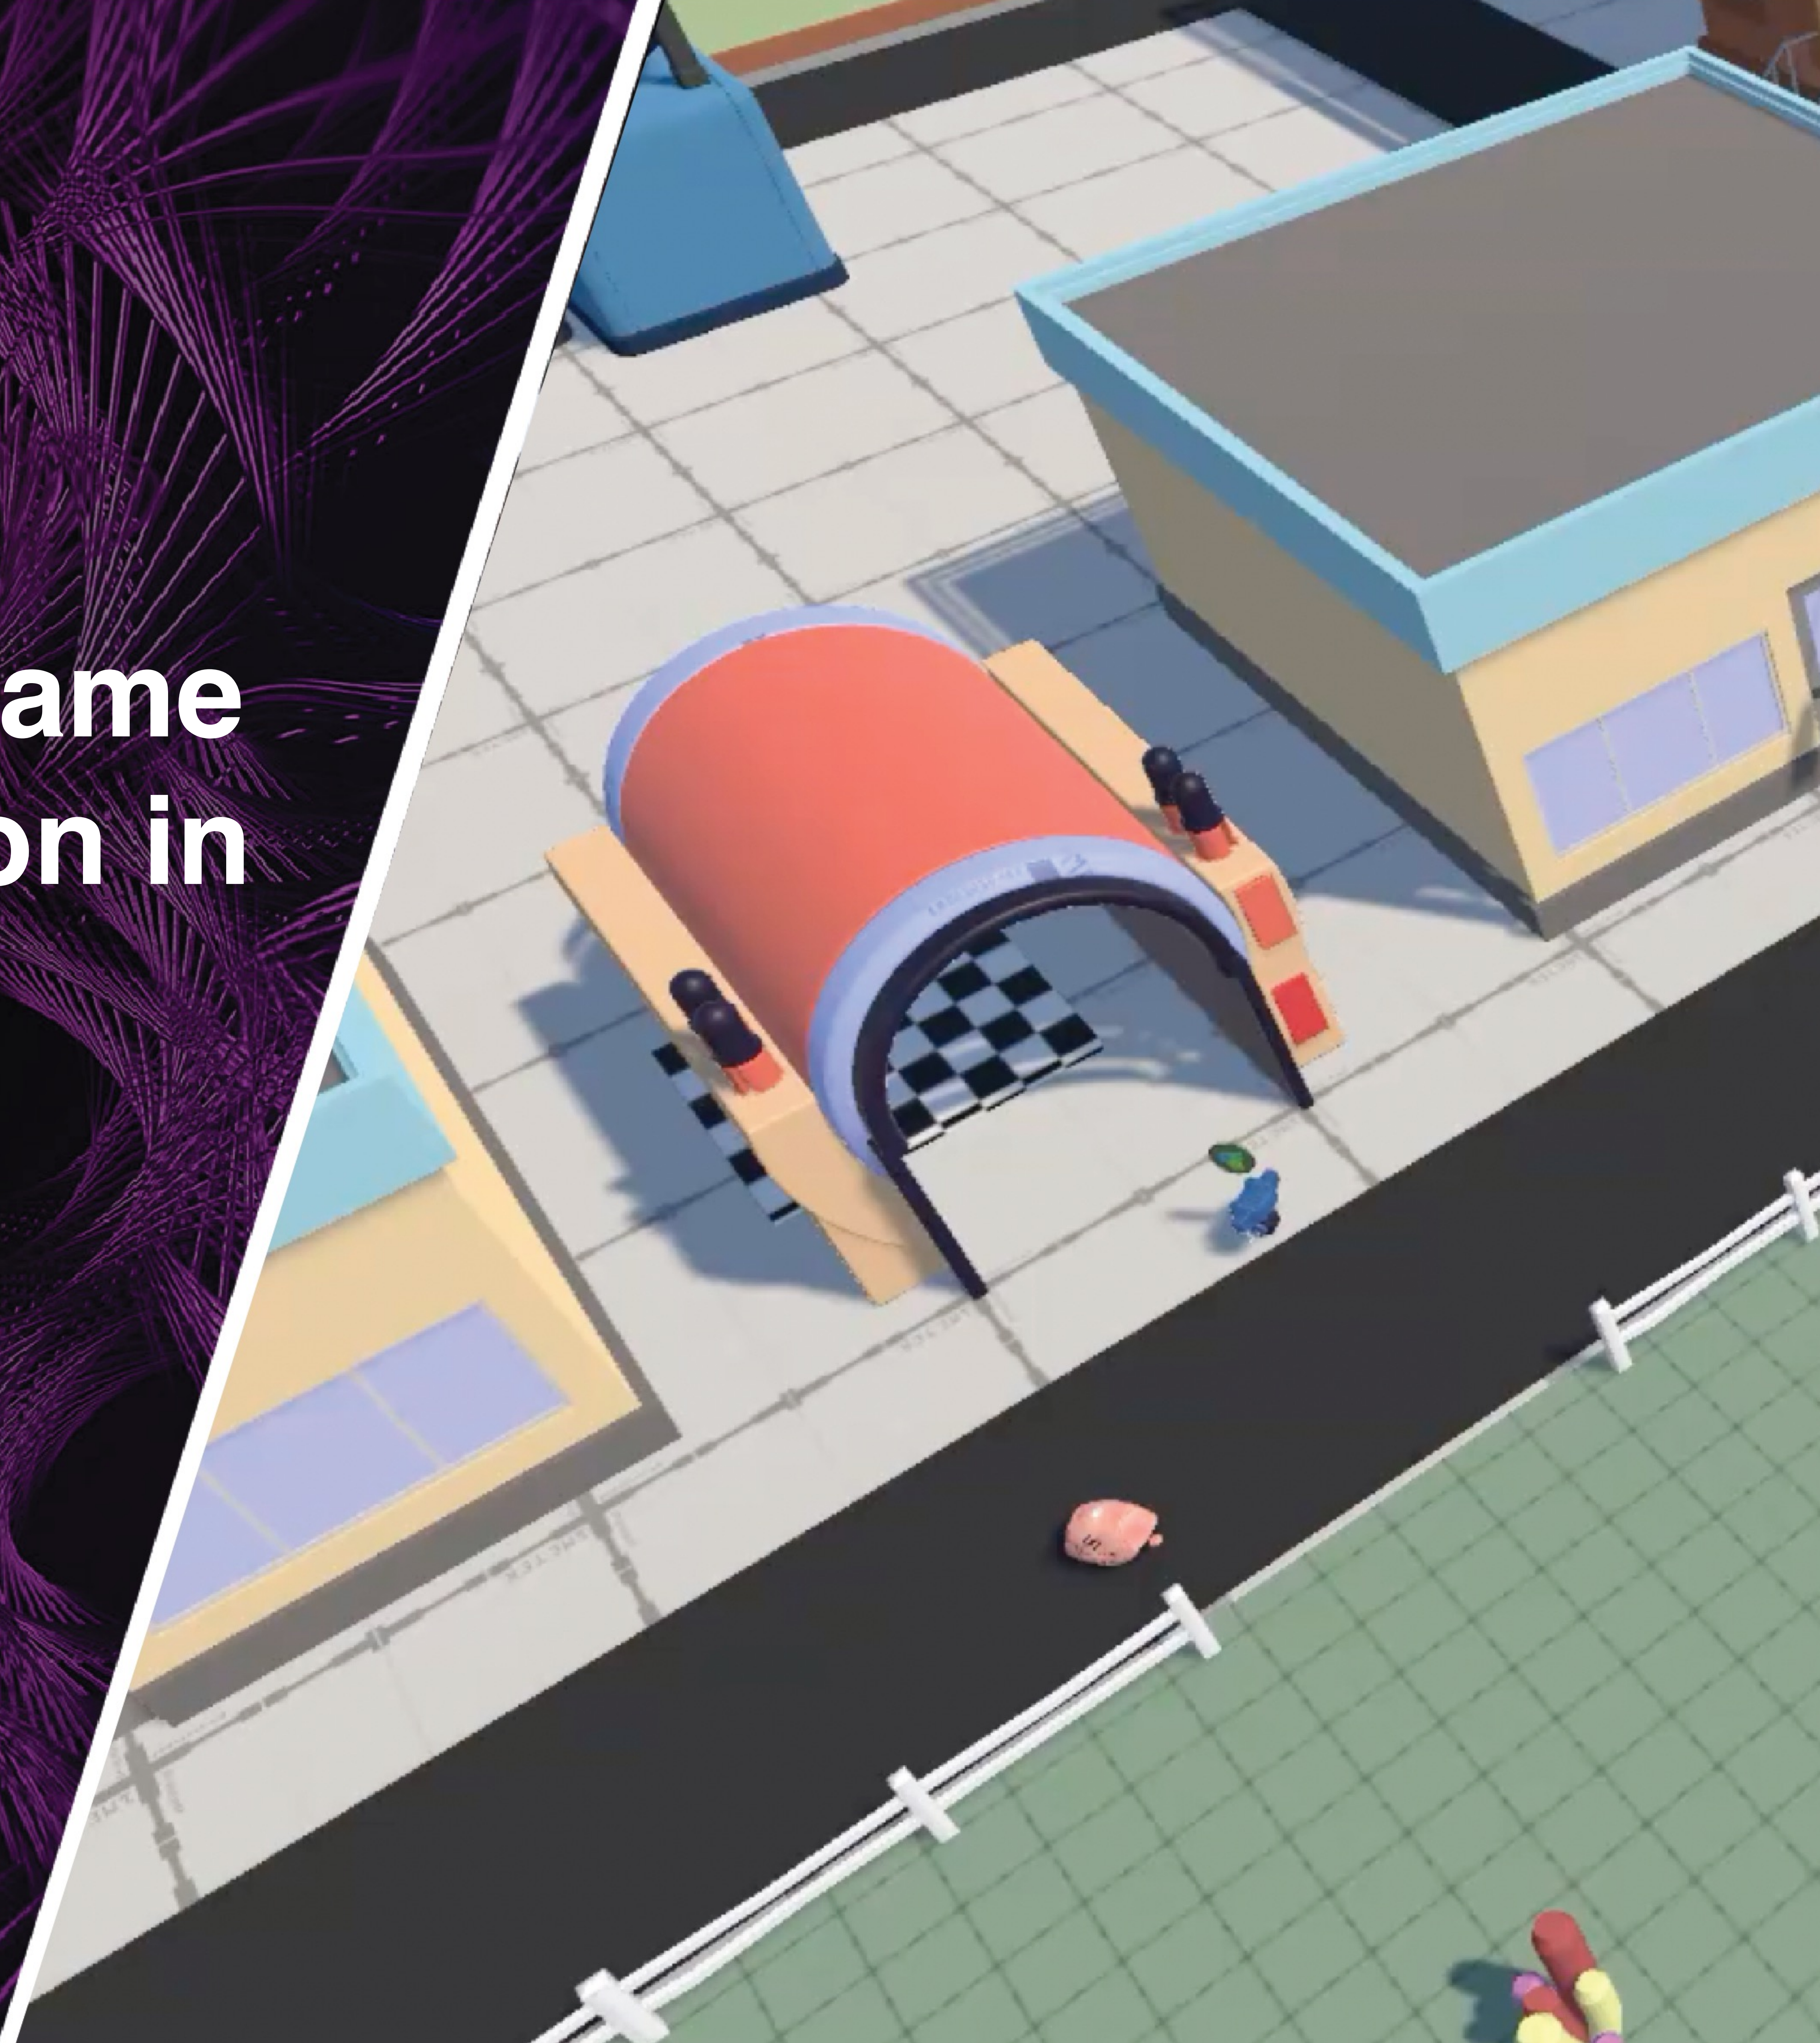

Automated Game Testing

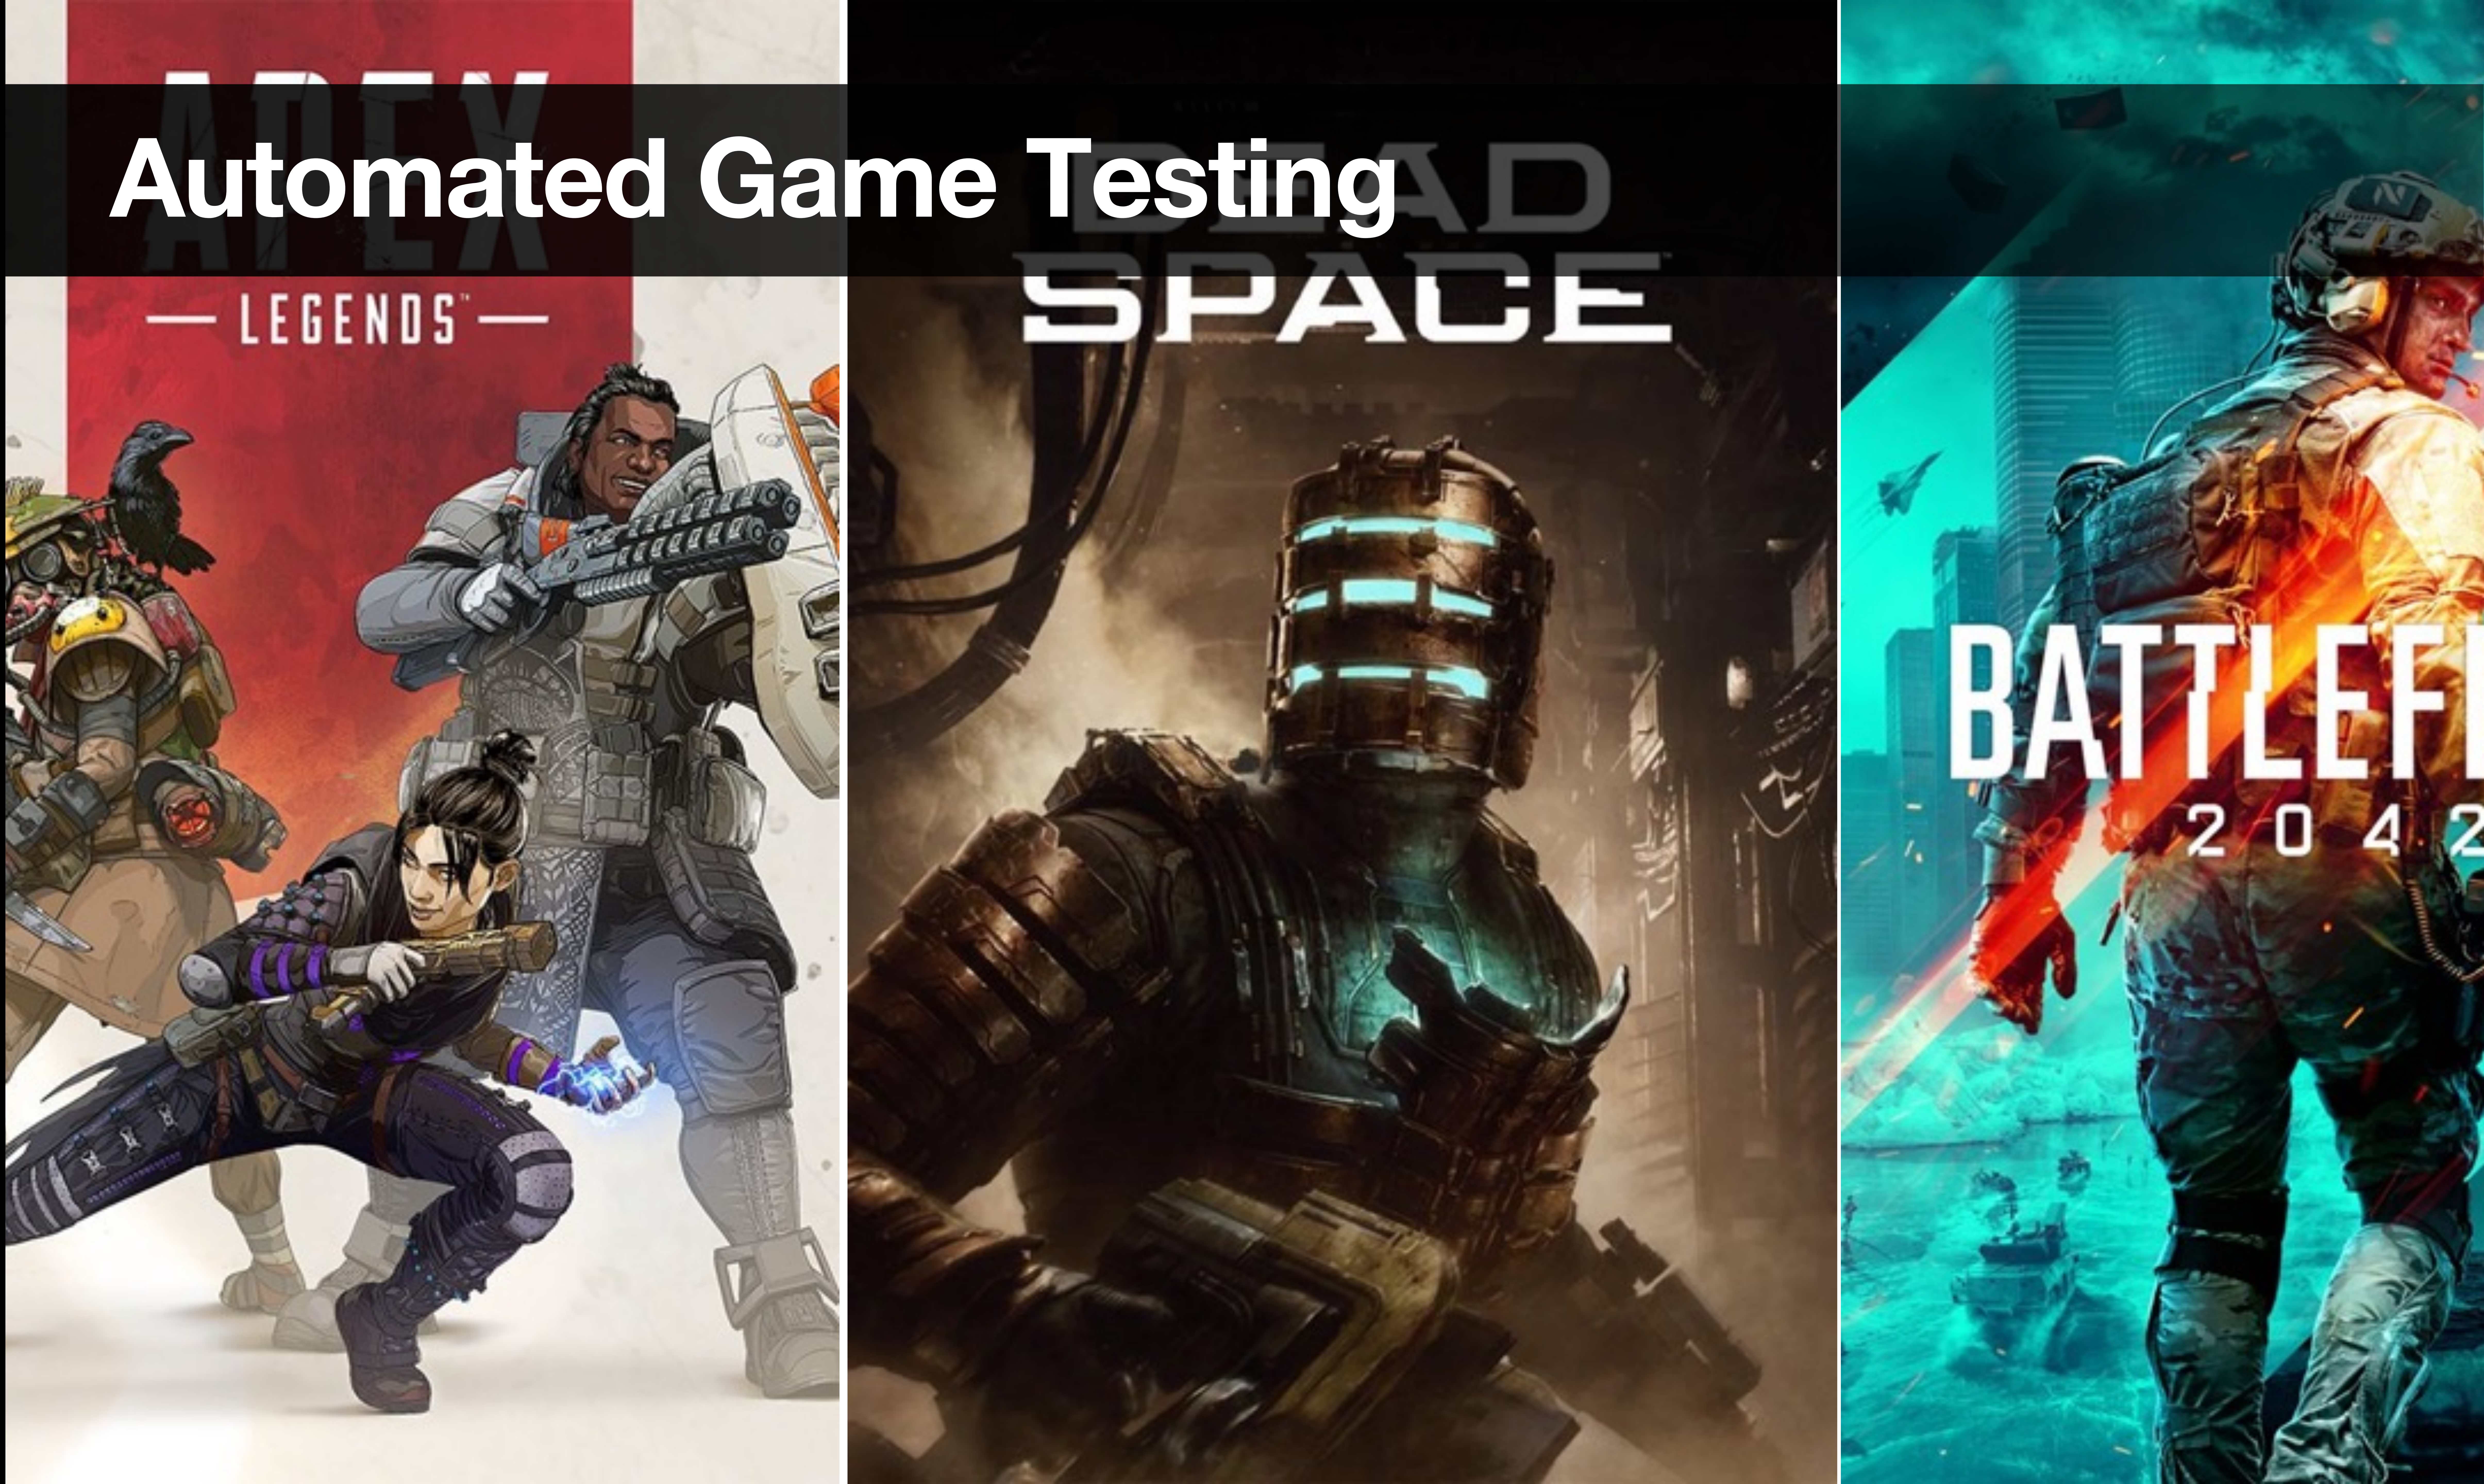

Imitation Learning for Game Testing[1]

Property	Scripting	RL	IL
Setup time	🕒🕒	🕒🕒🕒	🕒
Exploration	✗	✓	✗
Exploitation	✗	✓	✗
Controllability	✓	✗	✓
Generalization	✗	✓	✓
ML knowledge required	○○○○○	●●●●○	●●○○○
Programming needed	●●●●○	●●○○○	●○○○○

[1] *Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning*, Sestini et al., 2023

Automated Game Testing - An Example

Problem

- Designers can use **prior knowledge** to guide the agent towards its goal.
- For an agent to deal **out-of-distribution** data we need significant number of datasamples.
- We investigate how to improve generalization reducing data need via **data augmentation**.

Data Augmentation

Data Augmentation

Data Augmentation





Data Augmentation with Agents[2]

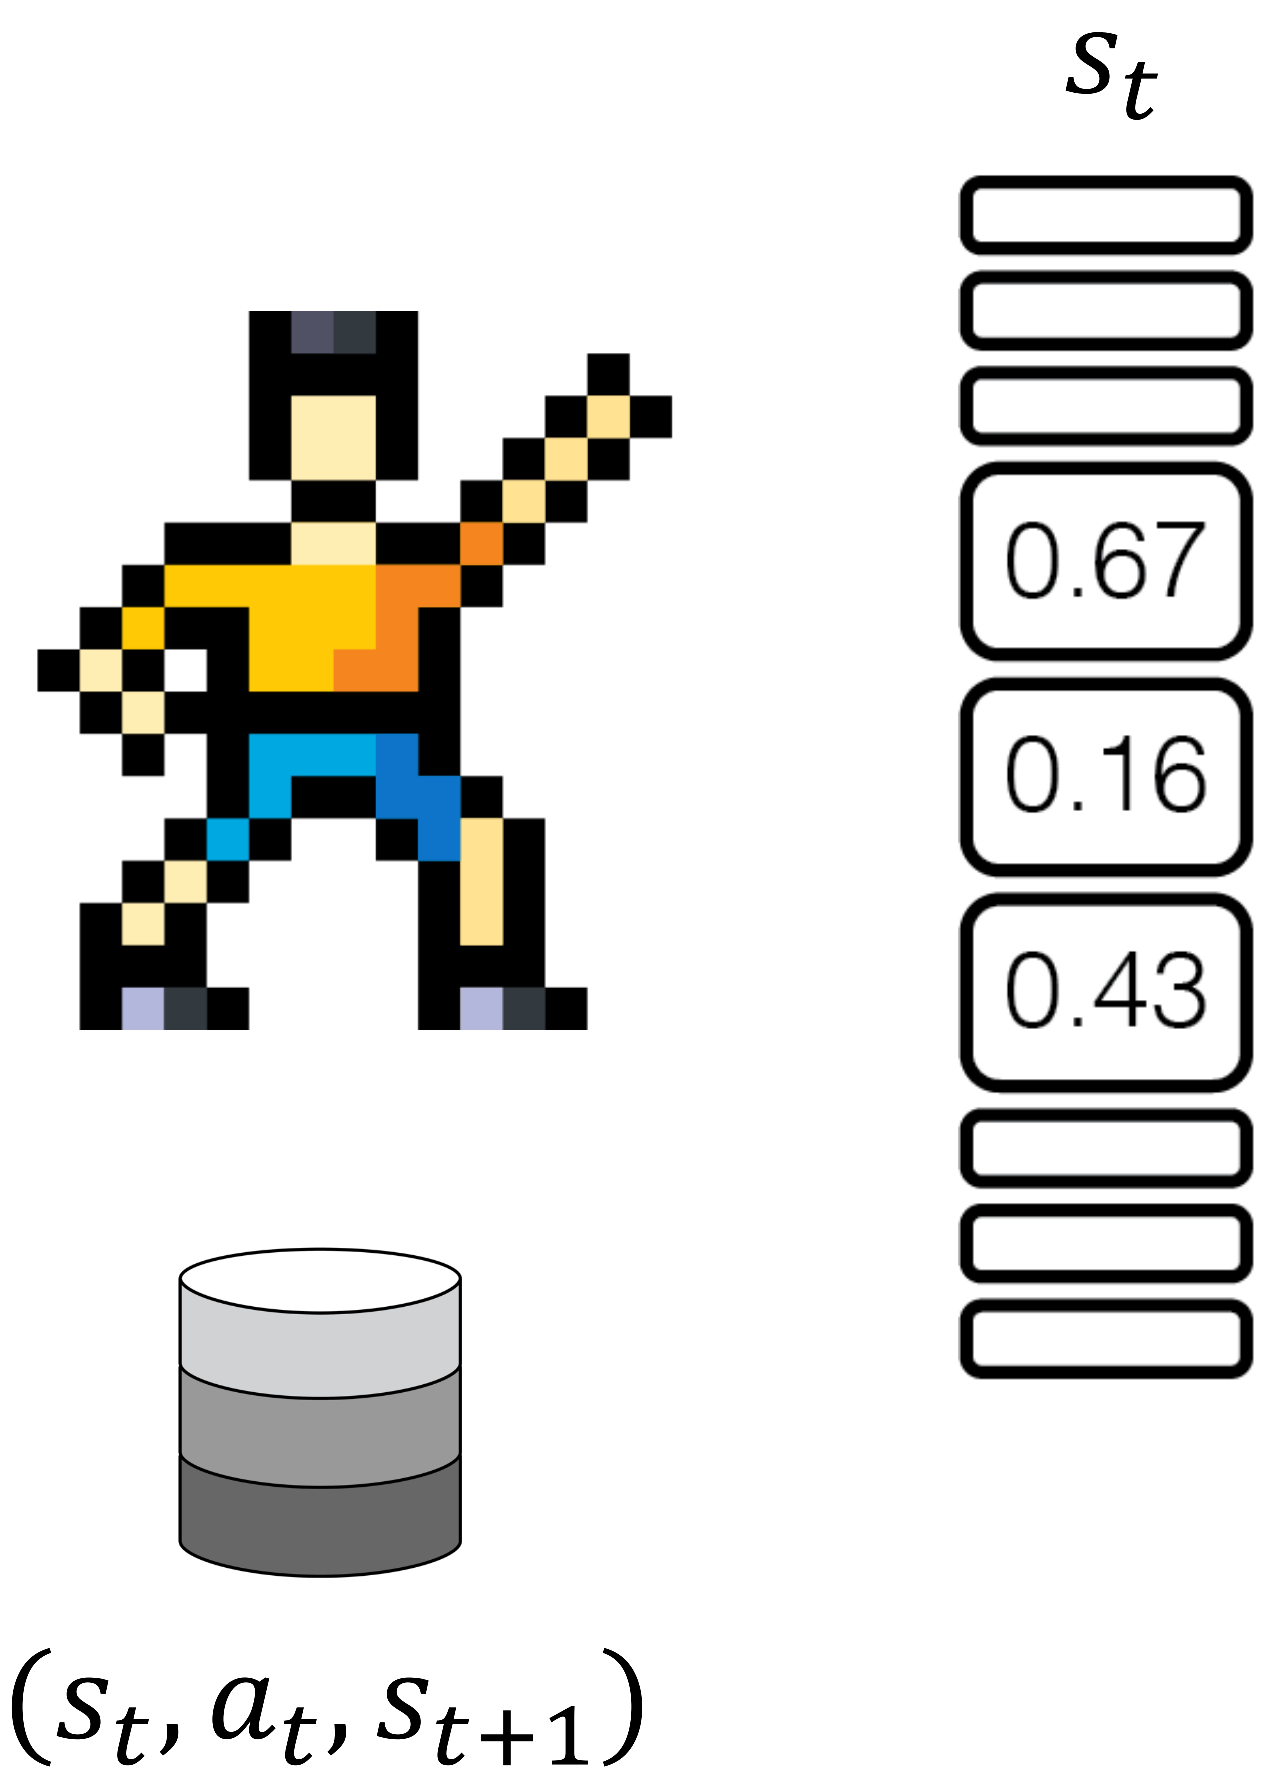

[2] S4RL: Surprisingly Simple Self-Supervision for Offline Reinforcement Learning, Sinha et al., 2021

Data Augmentation with Agents[2]

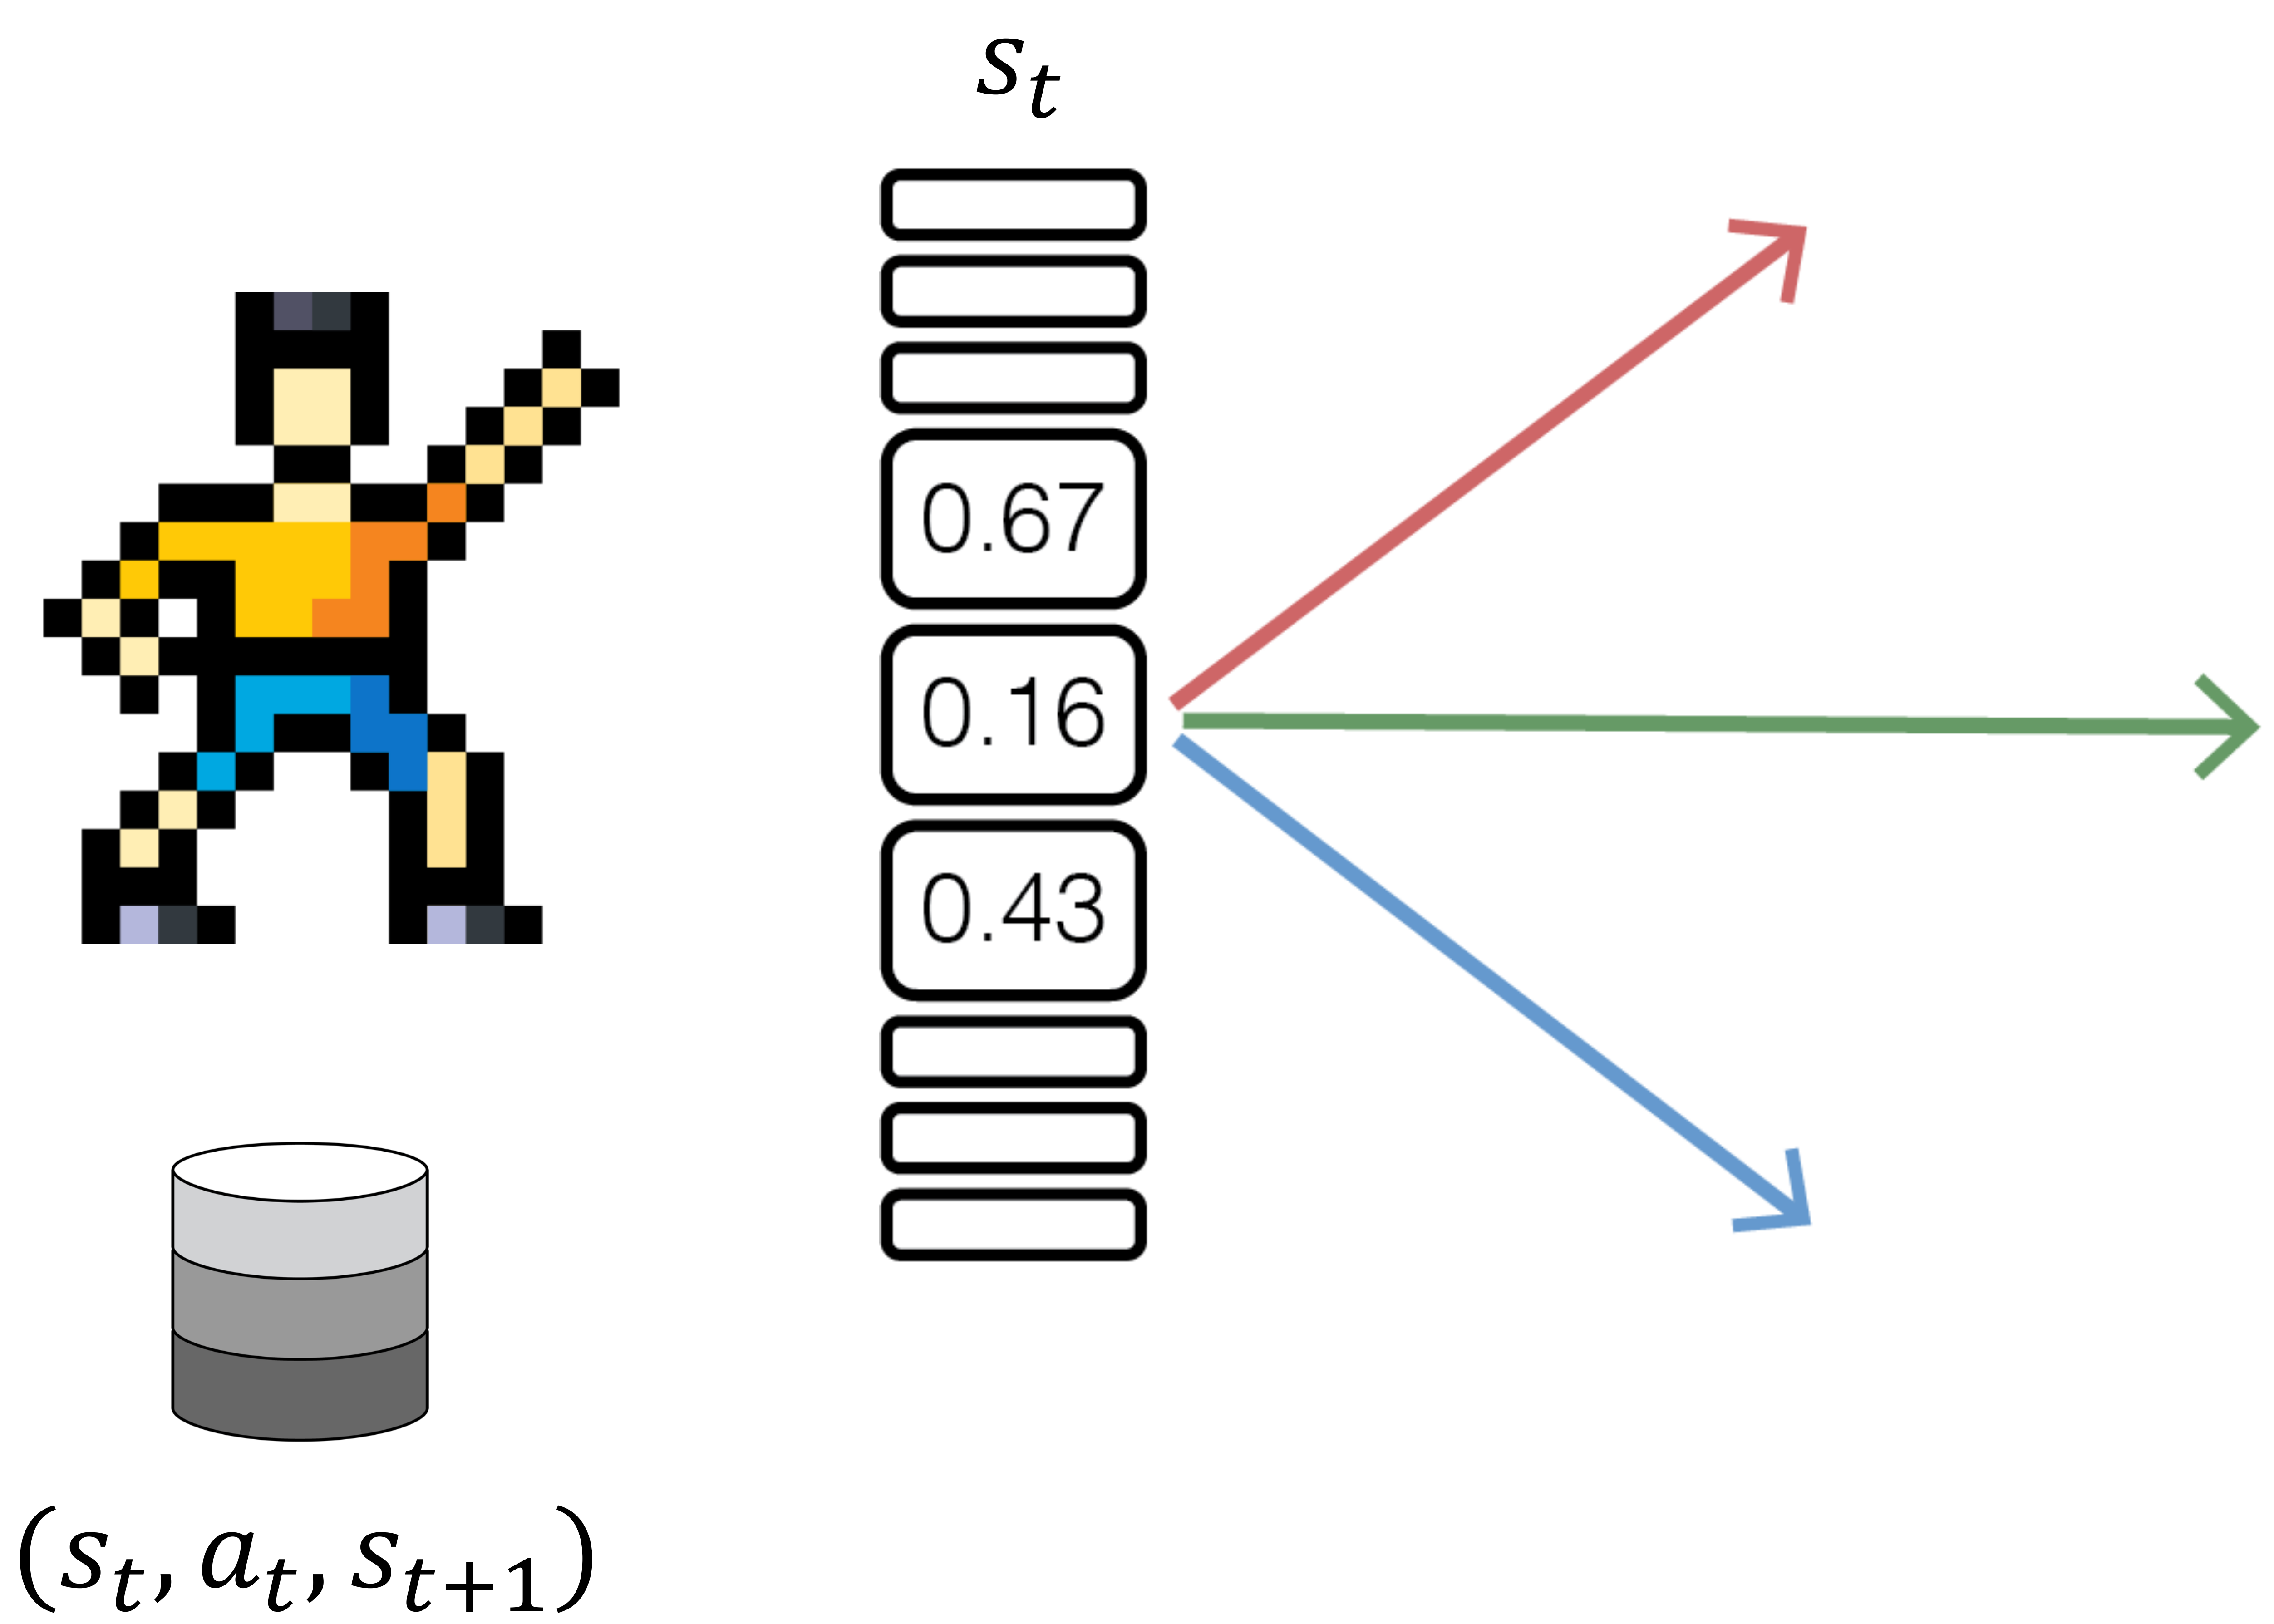

[2] S4RL: Surprisingly Simple Self-Supervision for Offline Reinforcement Learning, Sinha et al., 2021

Data Augmentation with Agents[2]

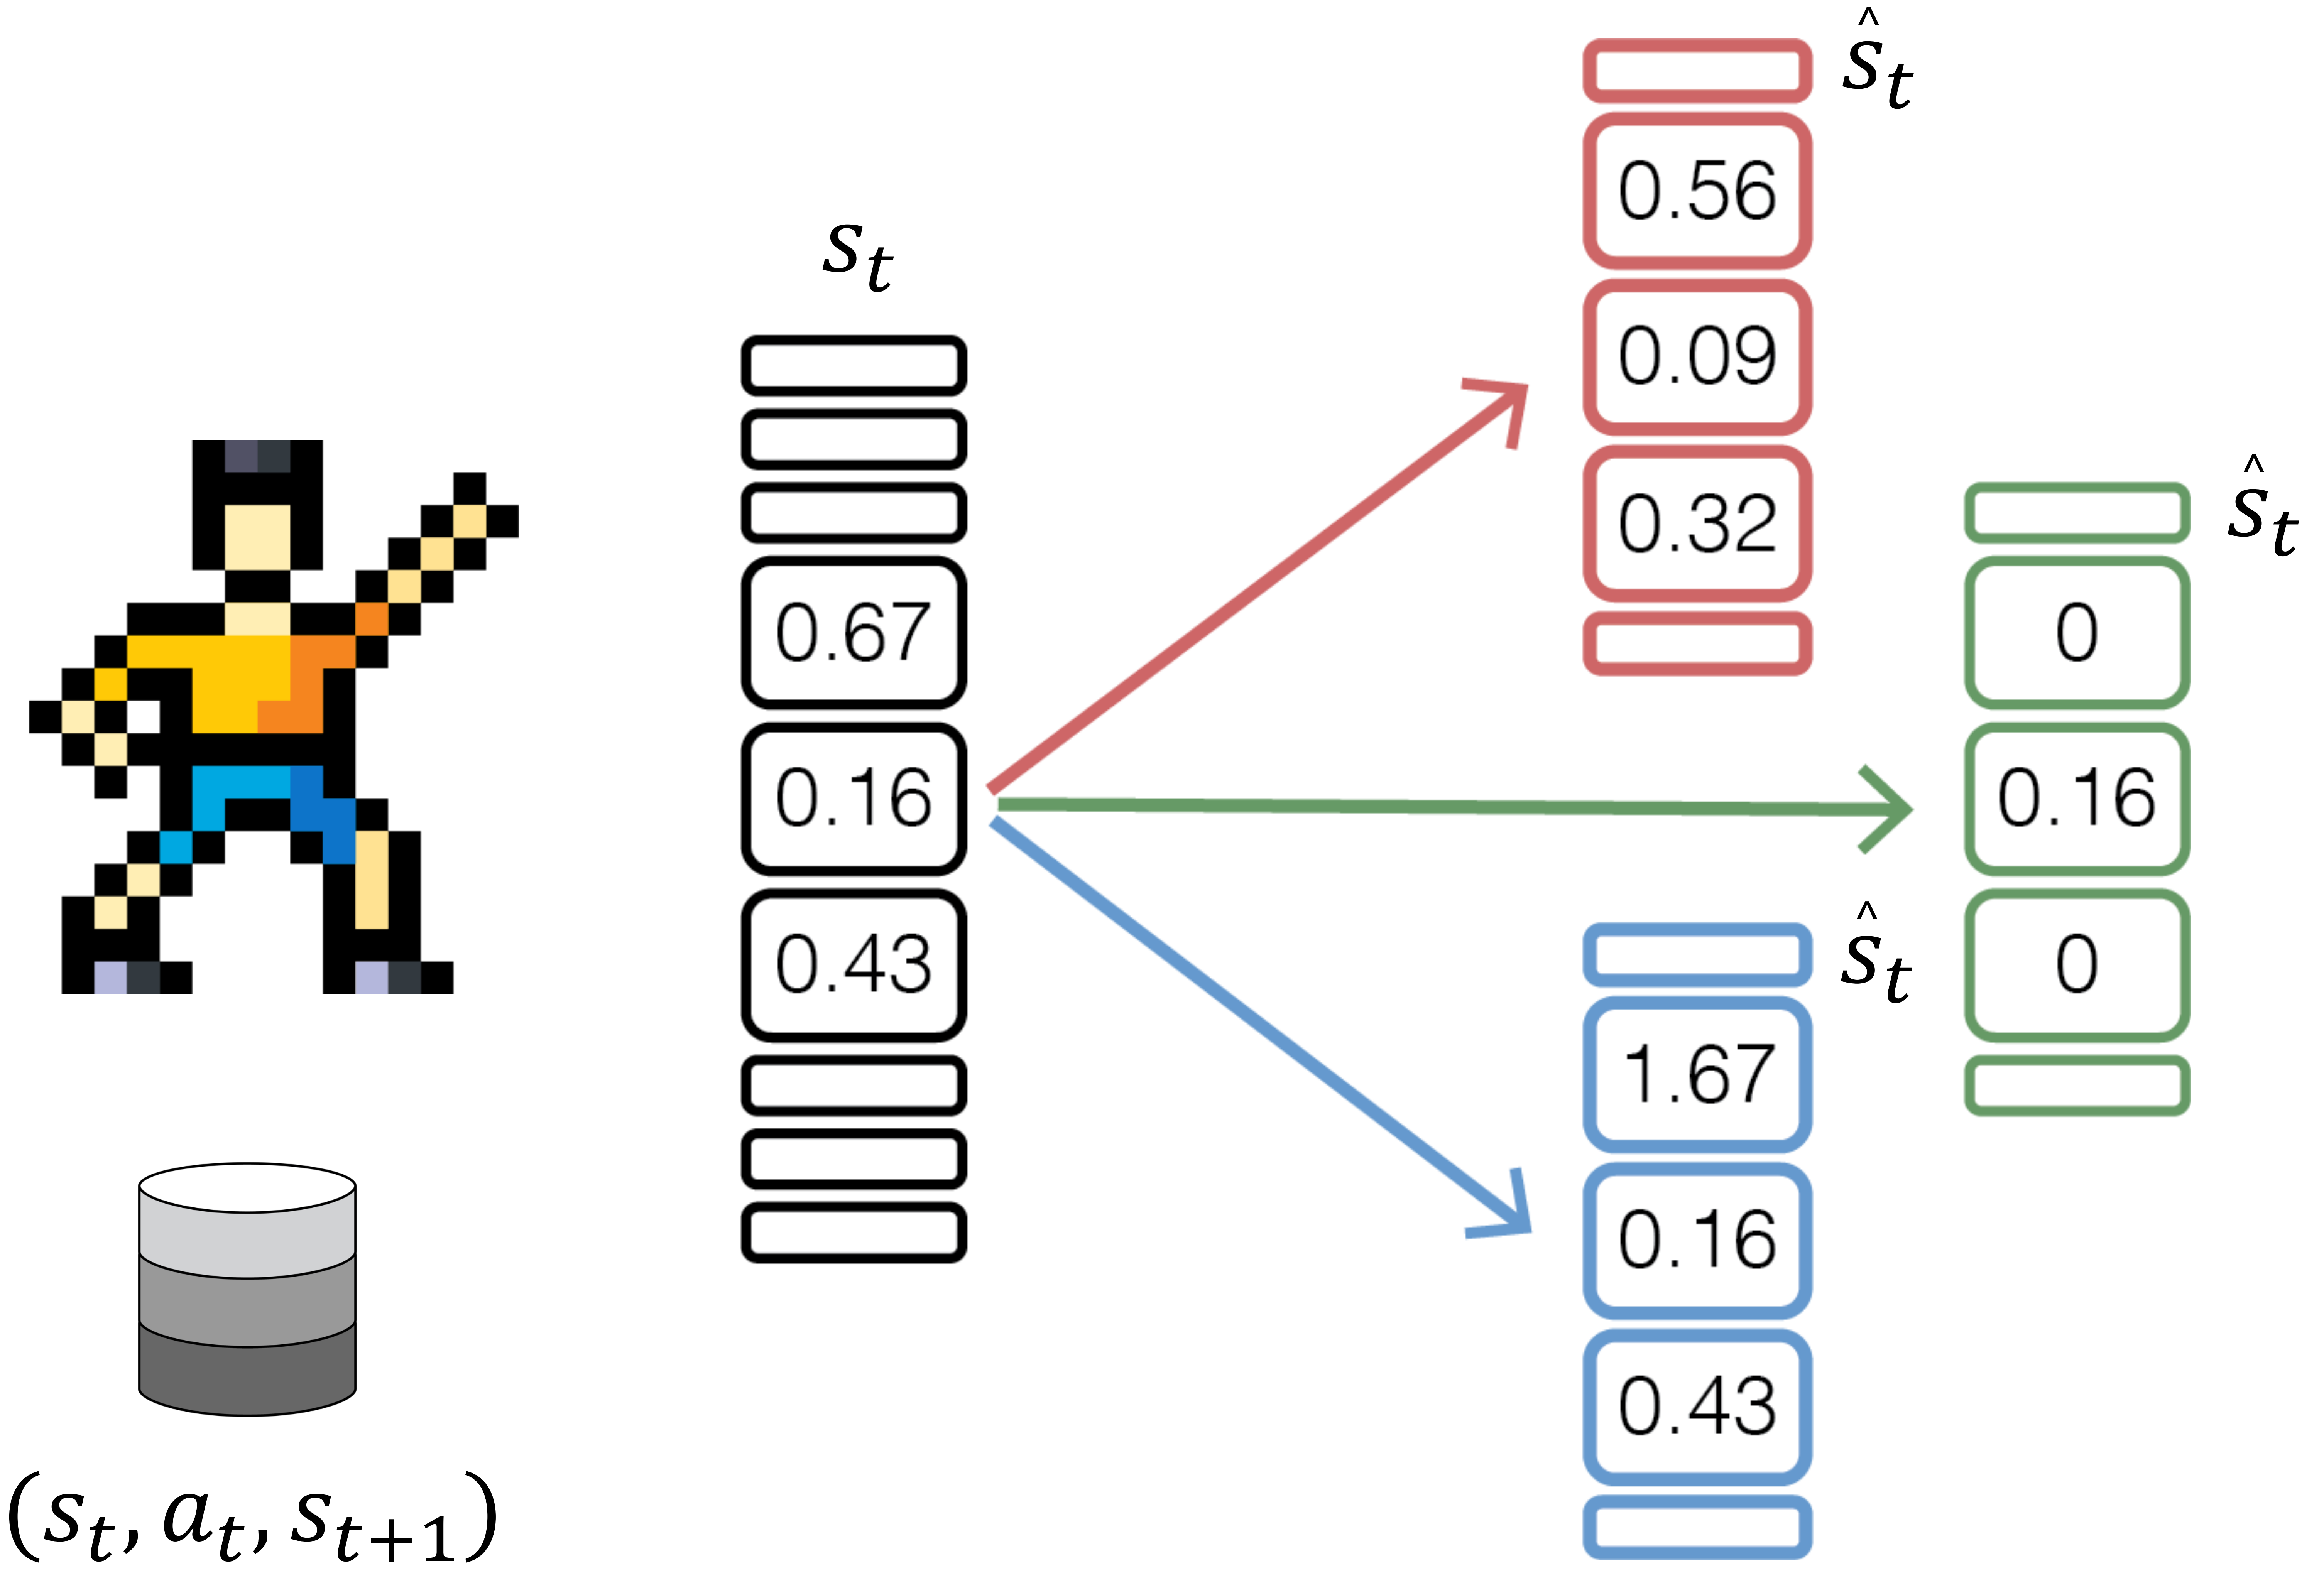

[2] S4RL: Surprisingly Simple Self-Supervision for Offline Reinforcement Learning, Sinha et al., 2021

Data Augmentation with Agents[2]

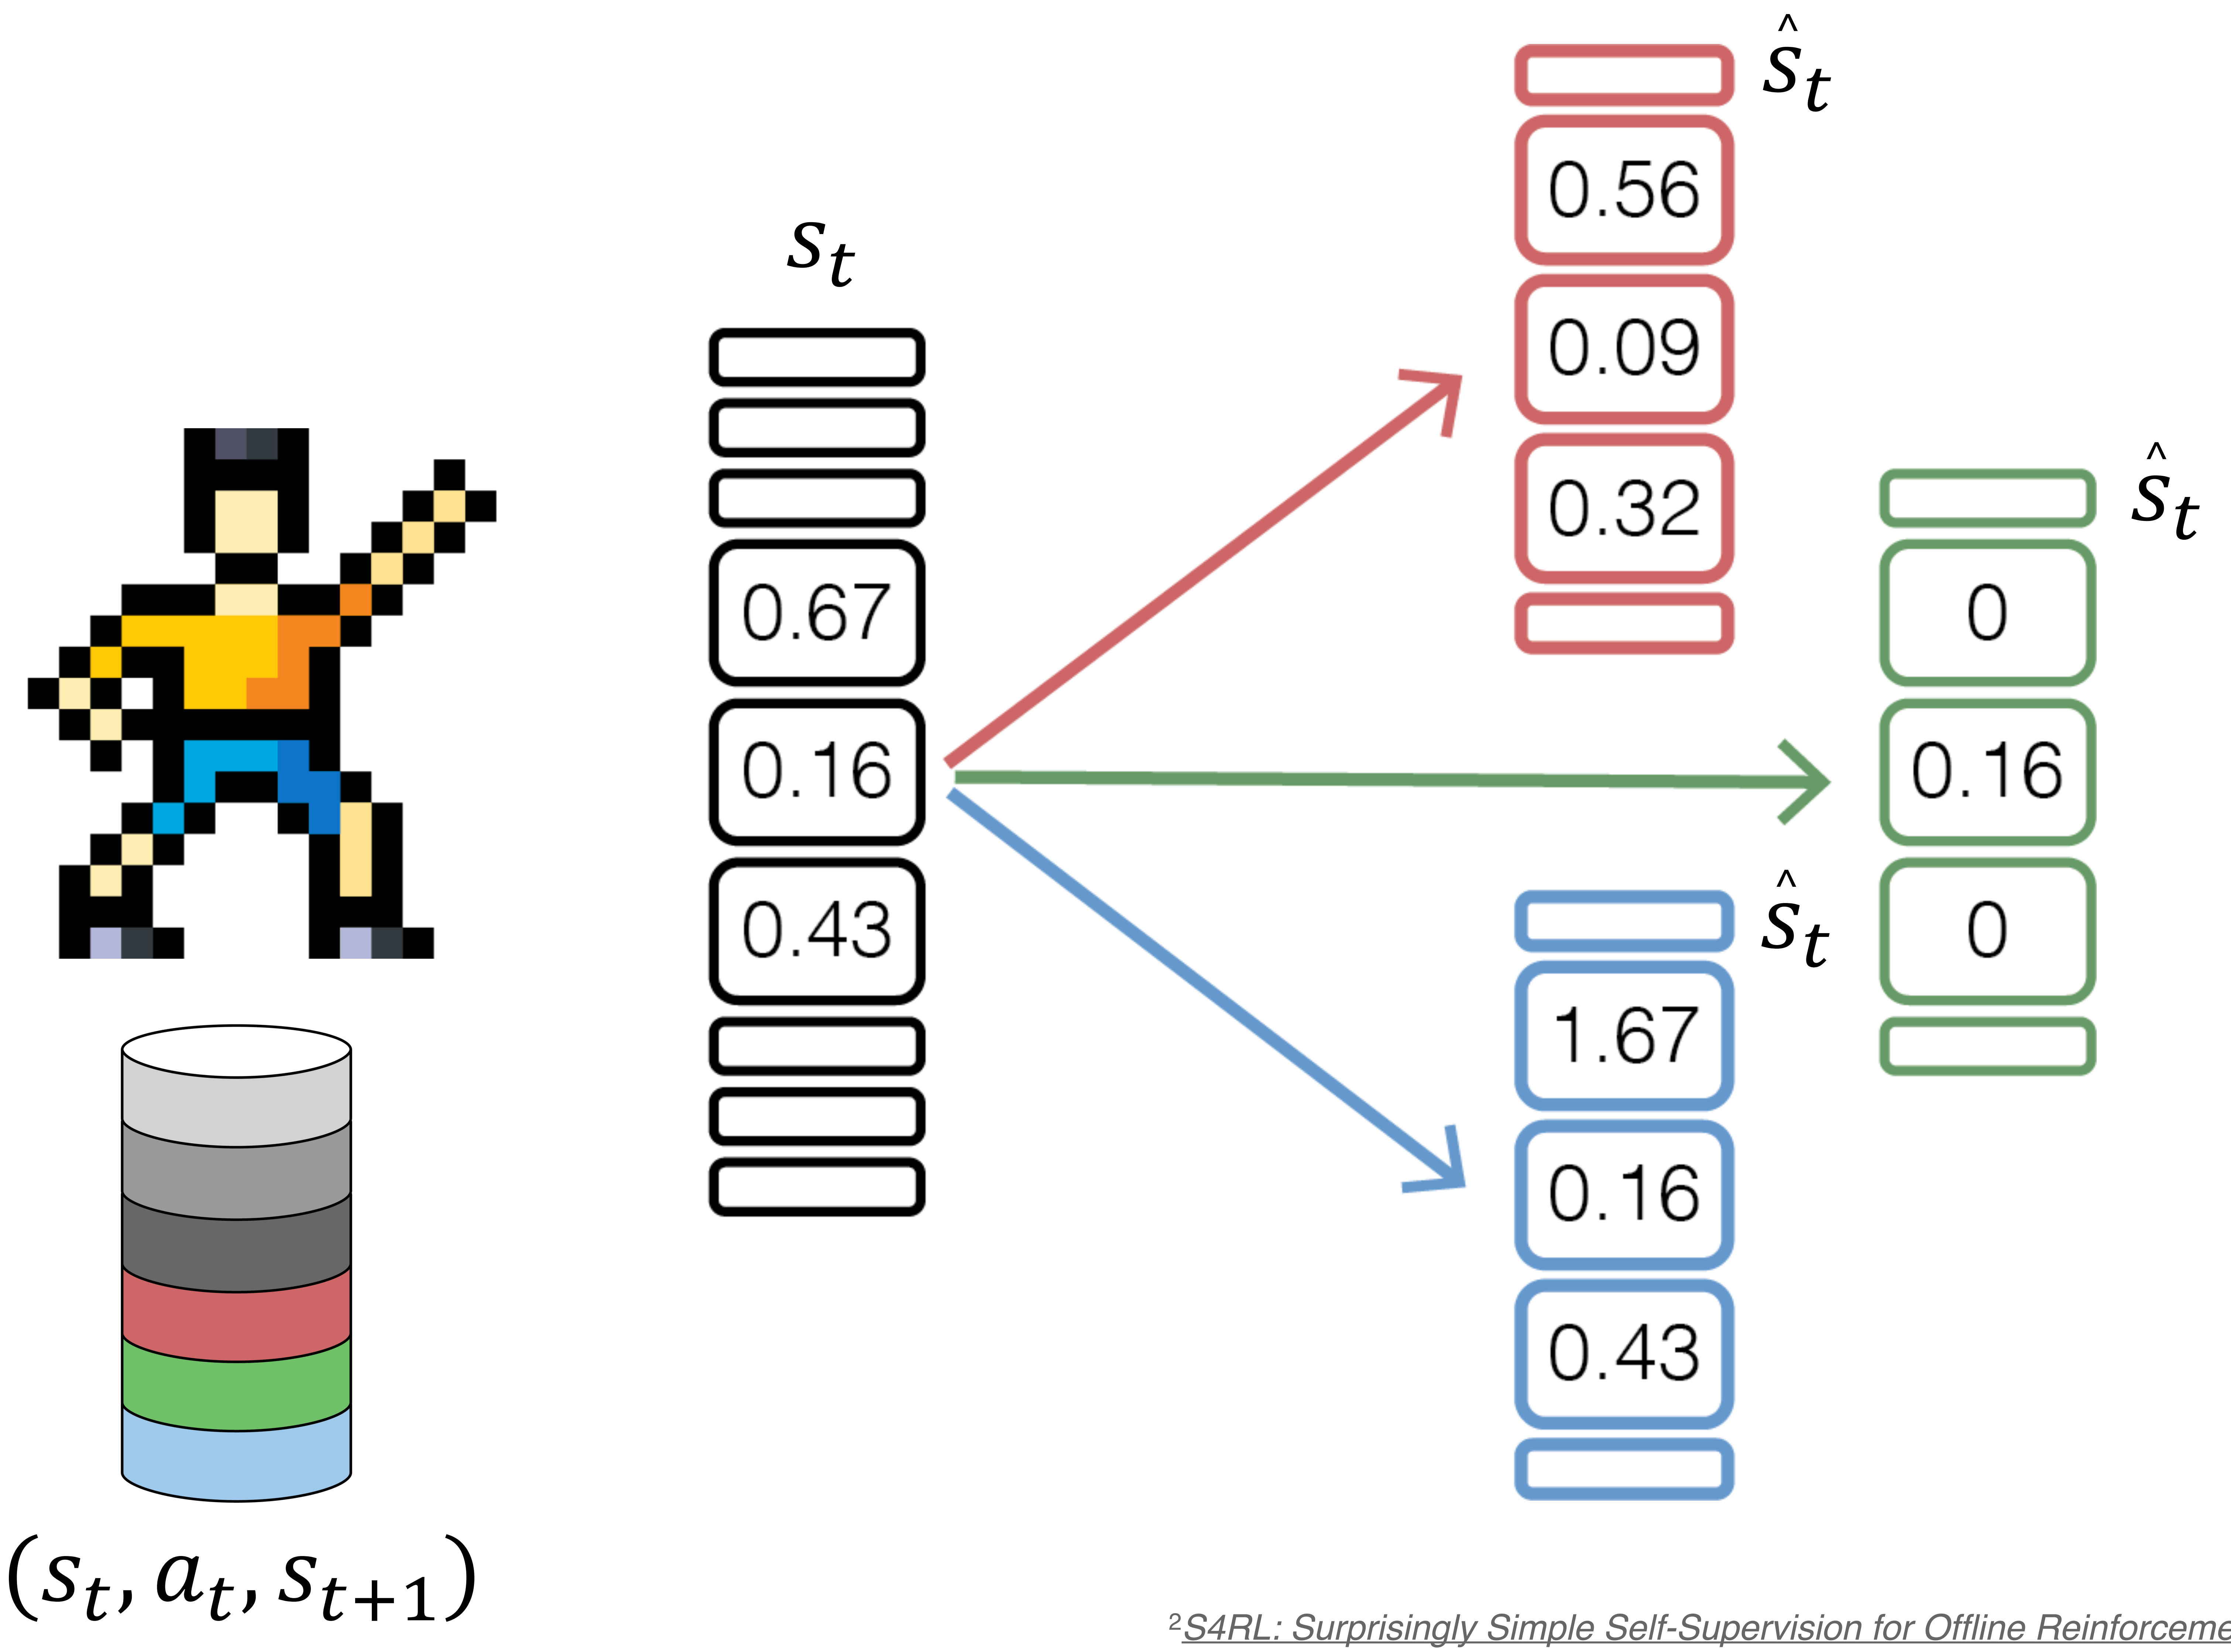

[2] S4RL: Surprisingly Simple Self-Supervision for Offline Reinforcement Learning, Sinha et al., 2021

Our Approach - Augmentations

Gaussian Noise: $\hat{s}_t = s_t + \epsilon \; where \; \epsilon \sim N(\mu, \sigma)$.

Uniform Noise: $\hat{s}_t = s_t + \epsilon \; where \; \epsilon \sim U(-\lambda, \lambda)$.

Scaling: $\hat{s}_t = s_t * \epsilon \; where \; \epsilon \sim U(\alpha, \beta)$.

State-MixUp: $\hat{s}_t = s_t * \epsilon + s_{t+1} * (1 - \epsilon) \; where \; \epsilon \sim \beta(\alpha, \alpha)$.

Our Approach - Augmentations

Continuous Dropout:

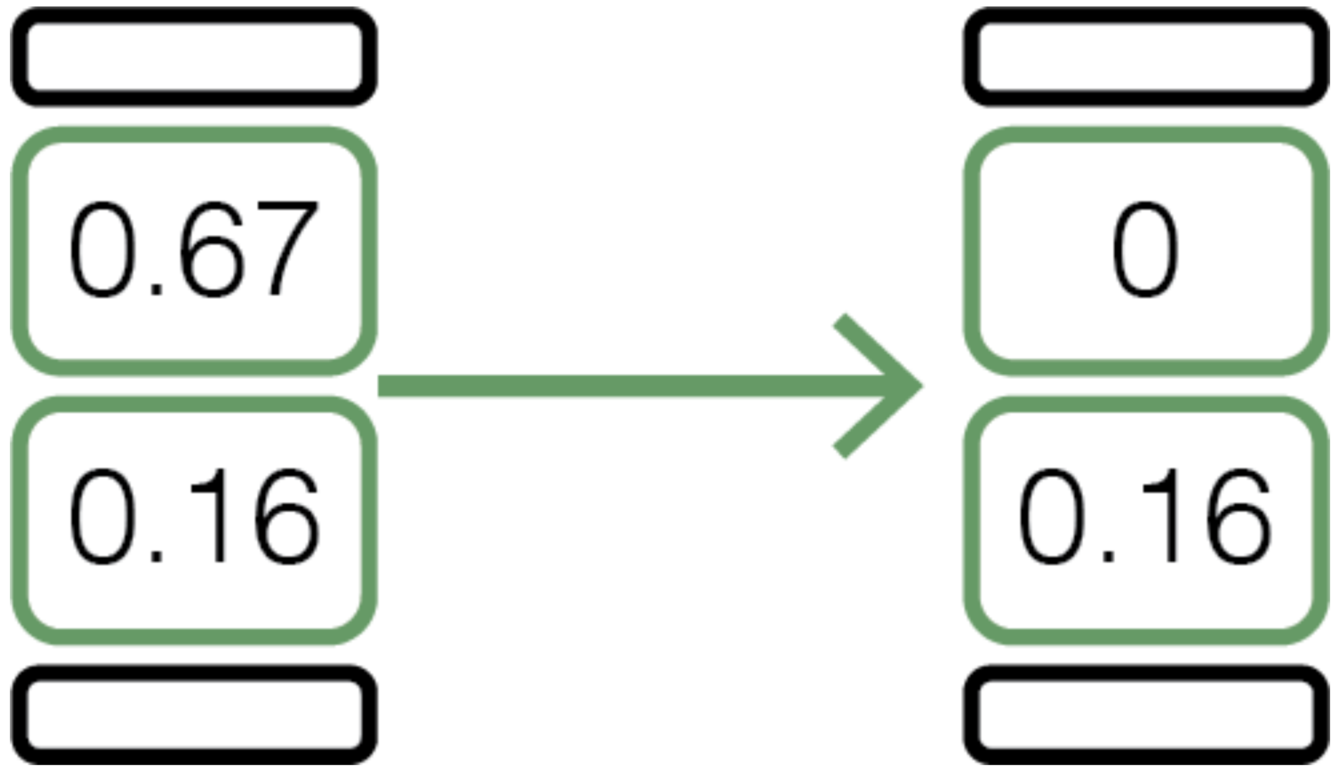

Semantic Dropout:

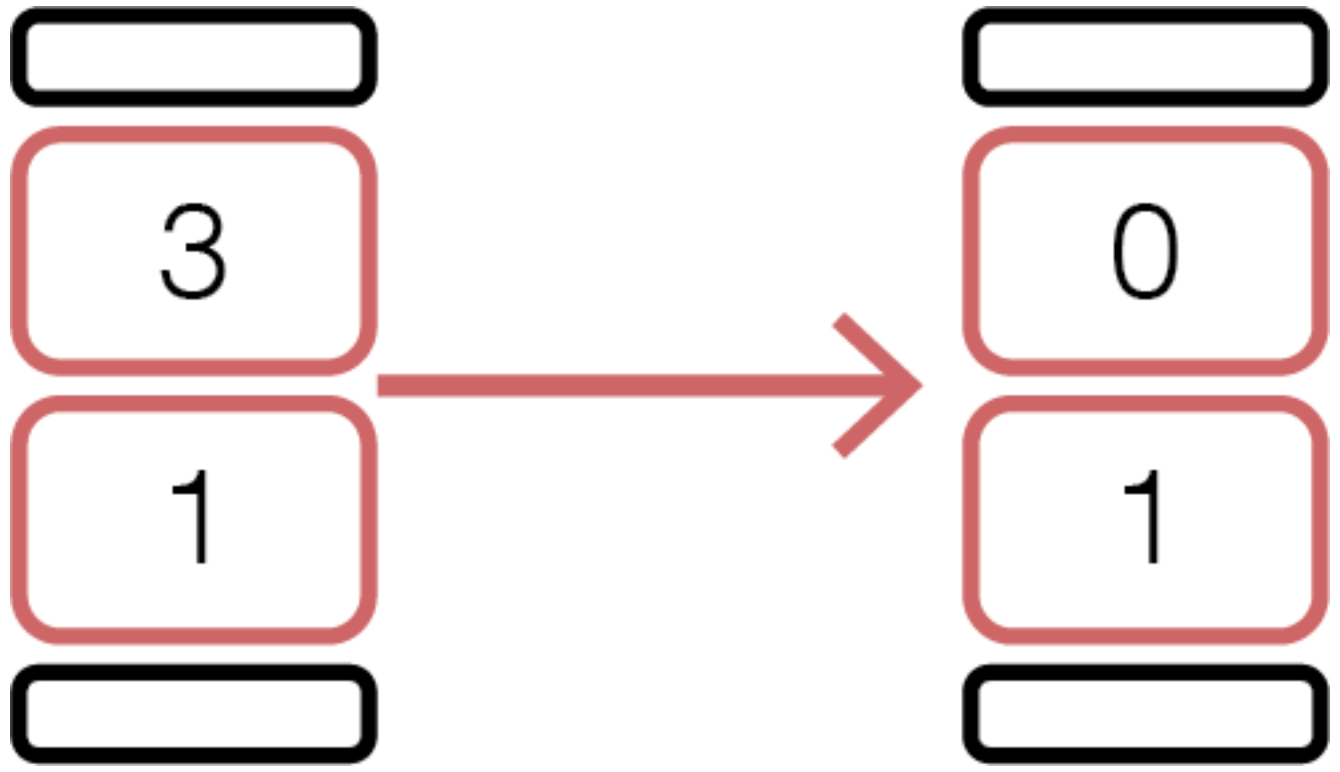

Our Approach - Training Algorithm

- Given a **demonstration** dataset of N trajectories τ_i:

$$D = \{\tau_i | \tau_i = (s_0^i, a_0^i, \ldots, s_T^i, a_T^i), i = 1, \ldots, N\},$$

- the objective aims to **mimic** the expert behavior which is represented by the dataset D:

$$L = \arg\max_{\theta} \mathbb{E}_{(s,a)\sim D}[\log \pi_\theta(a|s)].$$

Experiments - Environment[3]

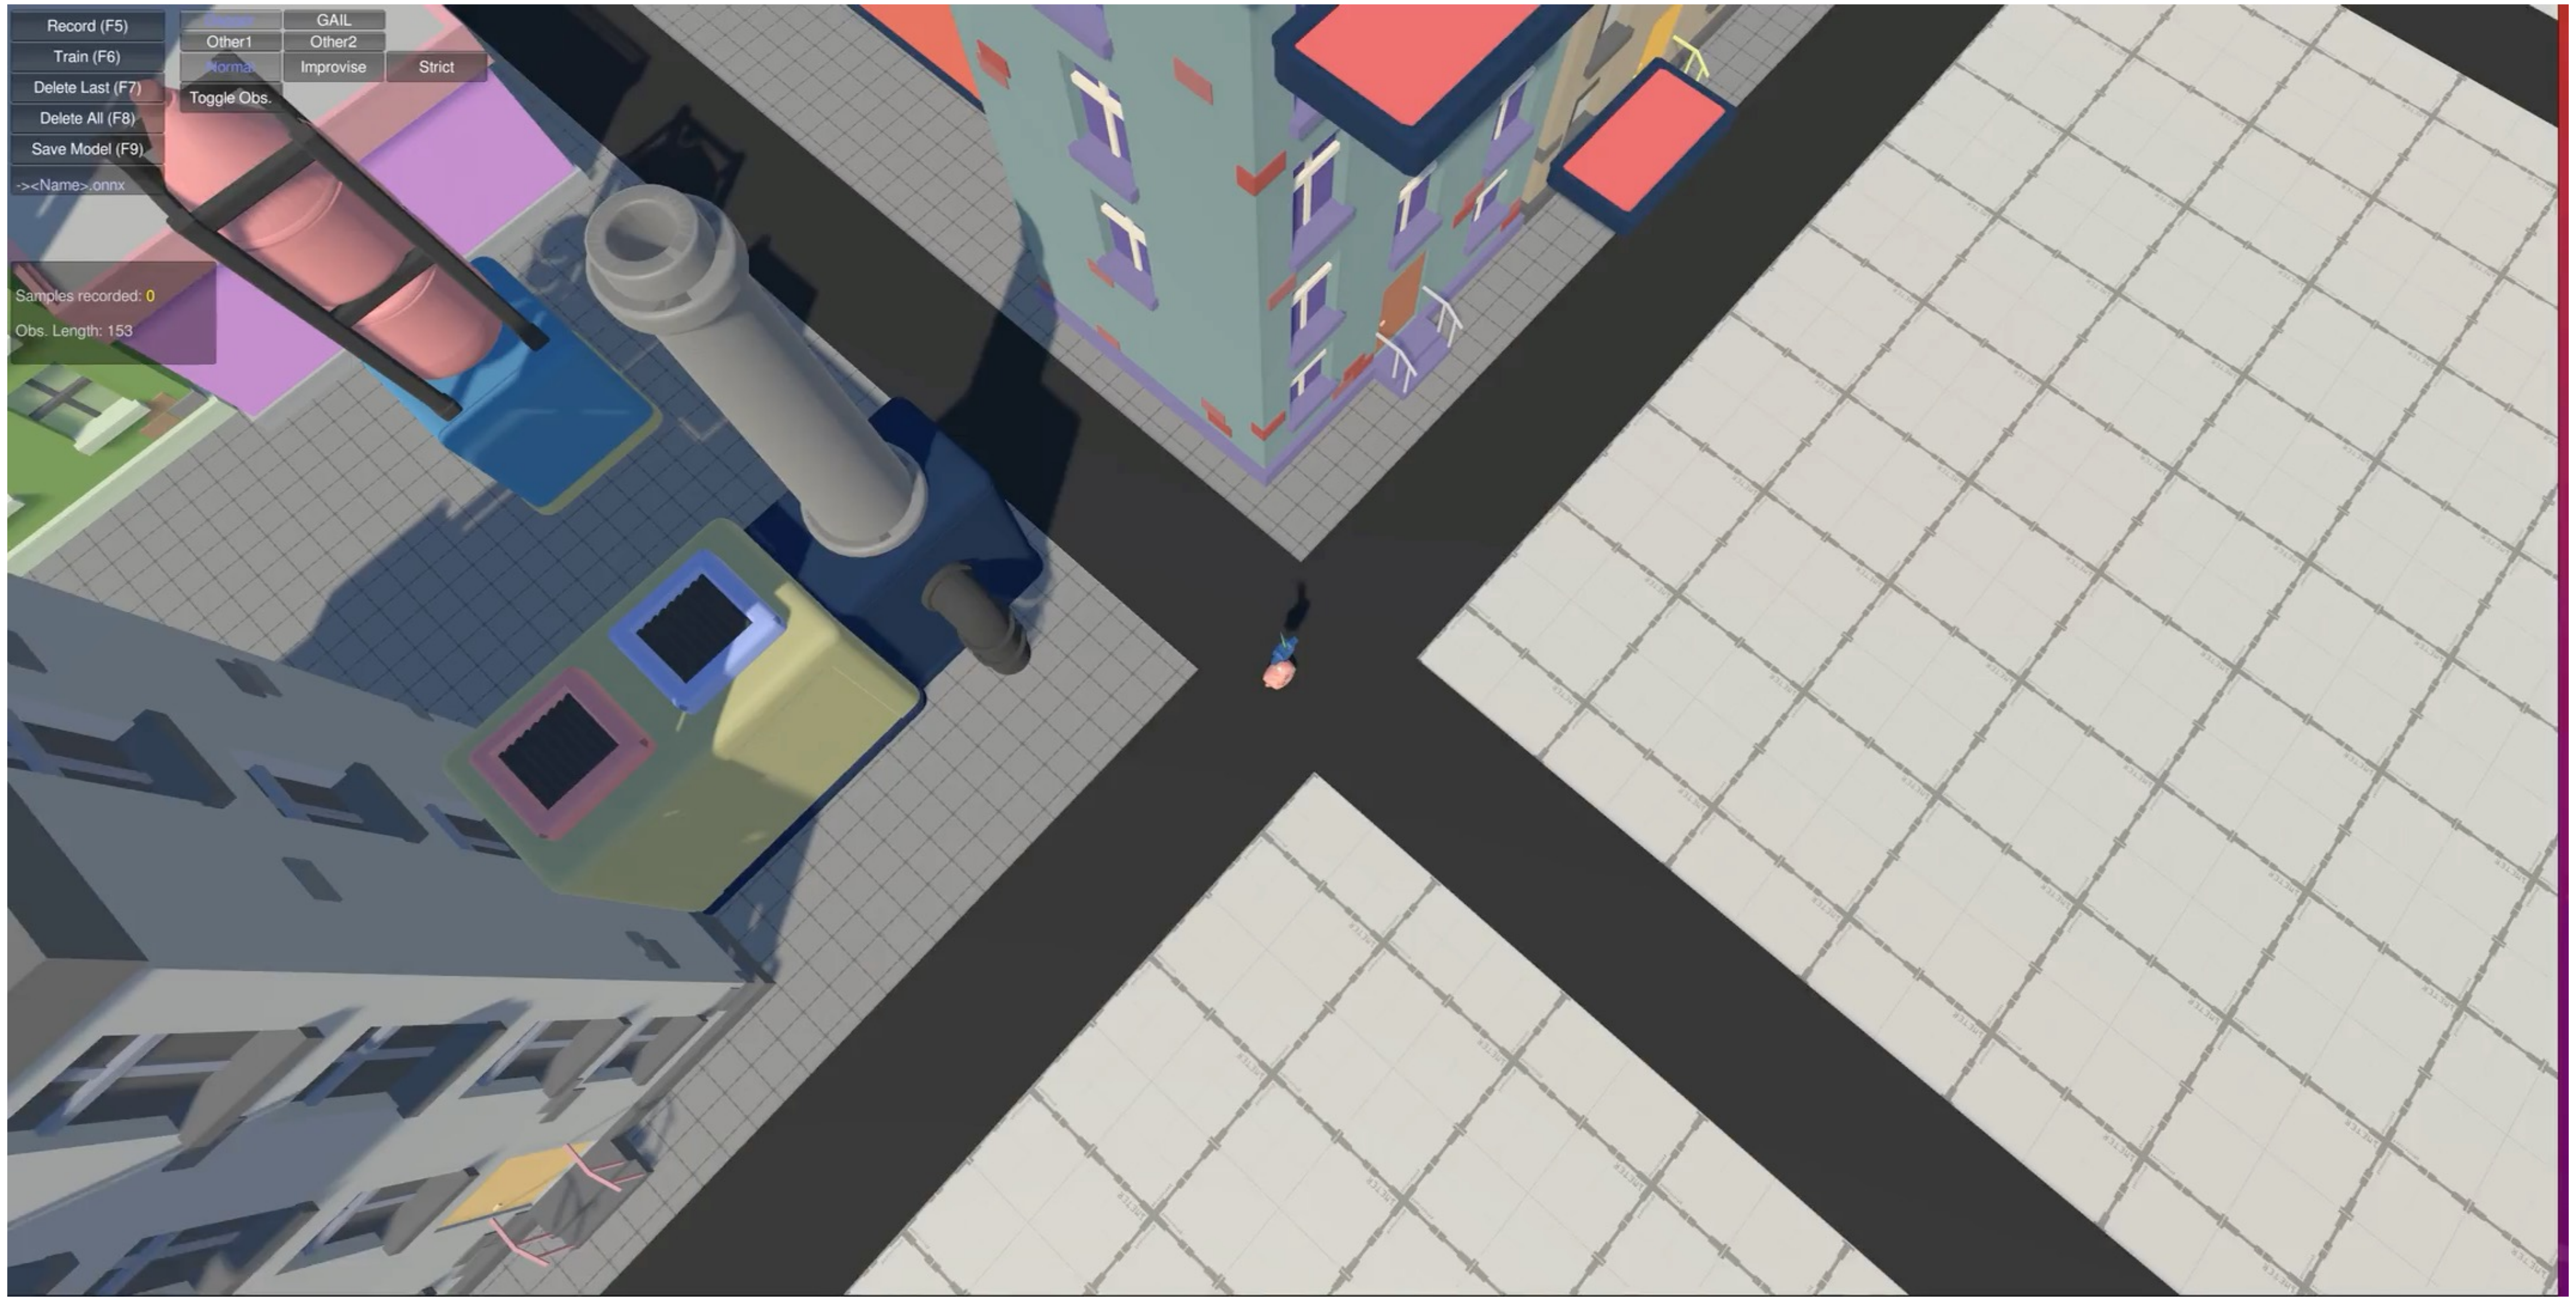

[3] *Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning*, Sestini et al., 2023

Experiments - Environment[3]

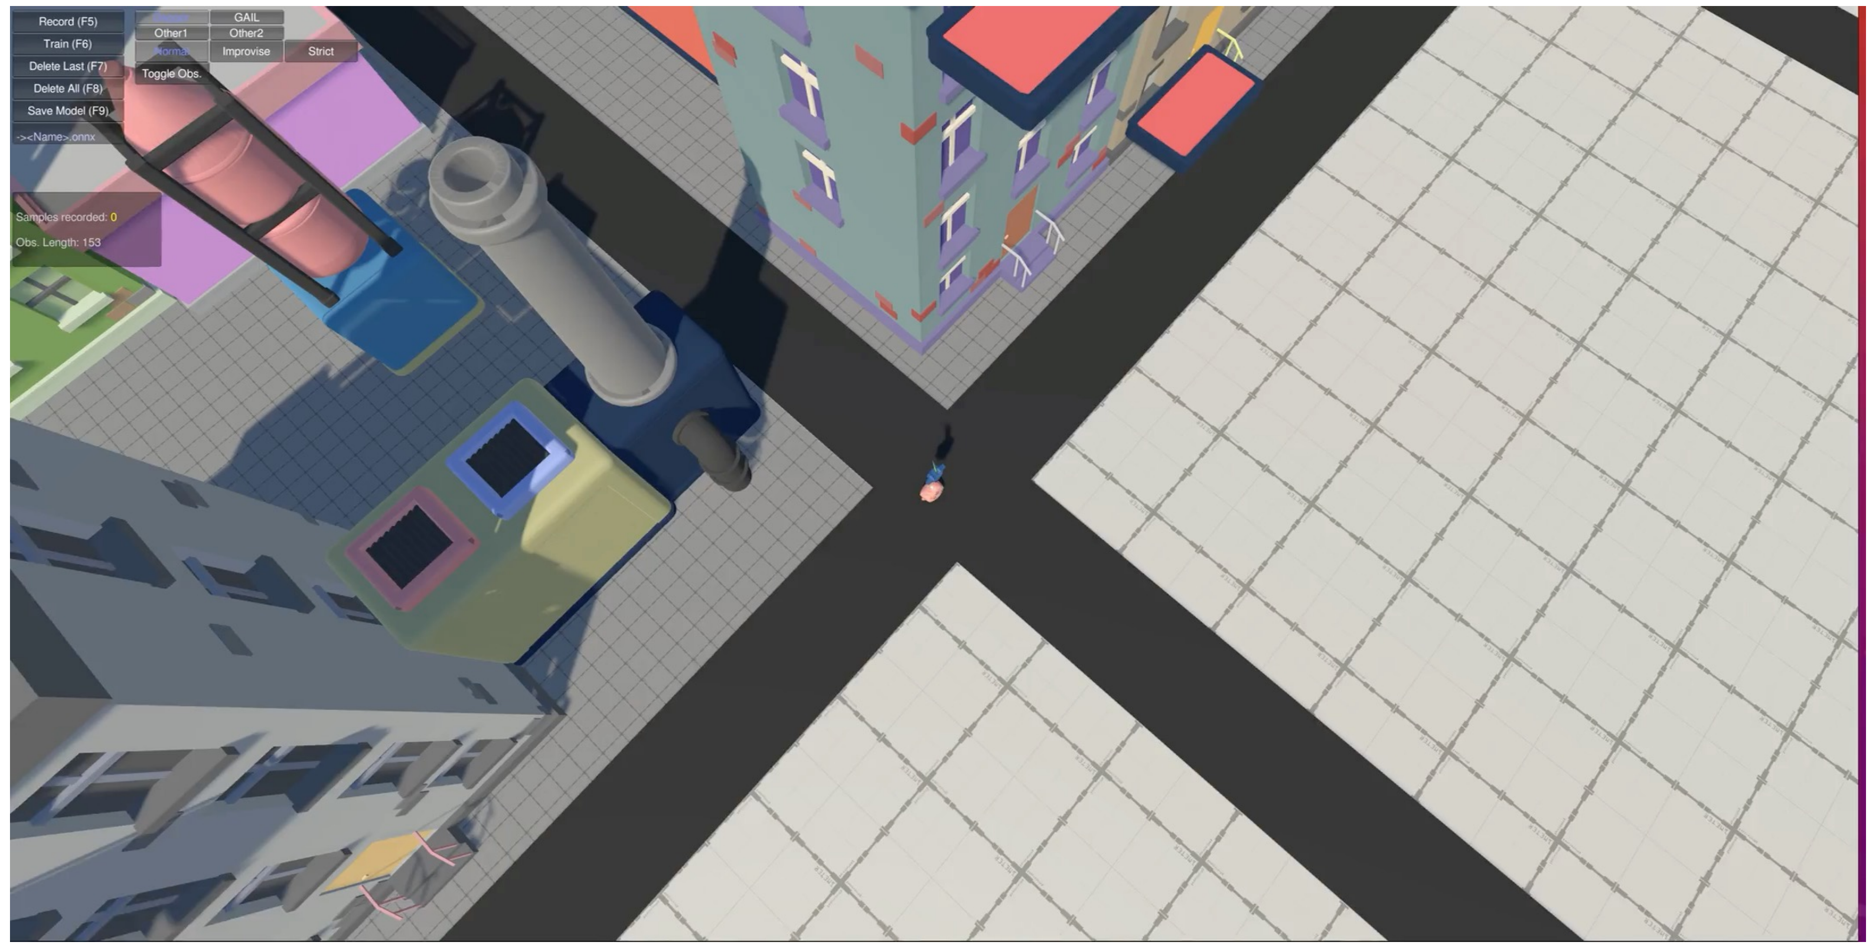

Training Environment

[3] *Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning*, Sestini et al., 2023

Experiments - Environment[3]

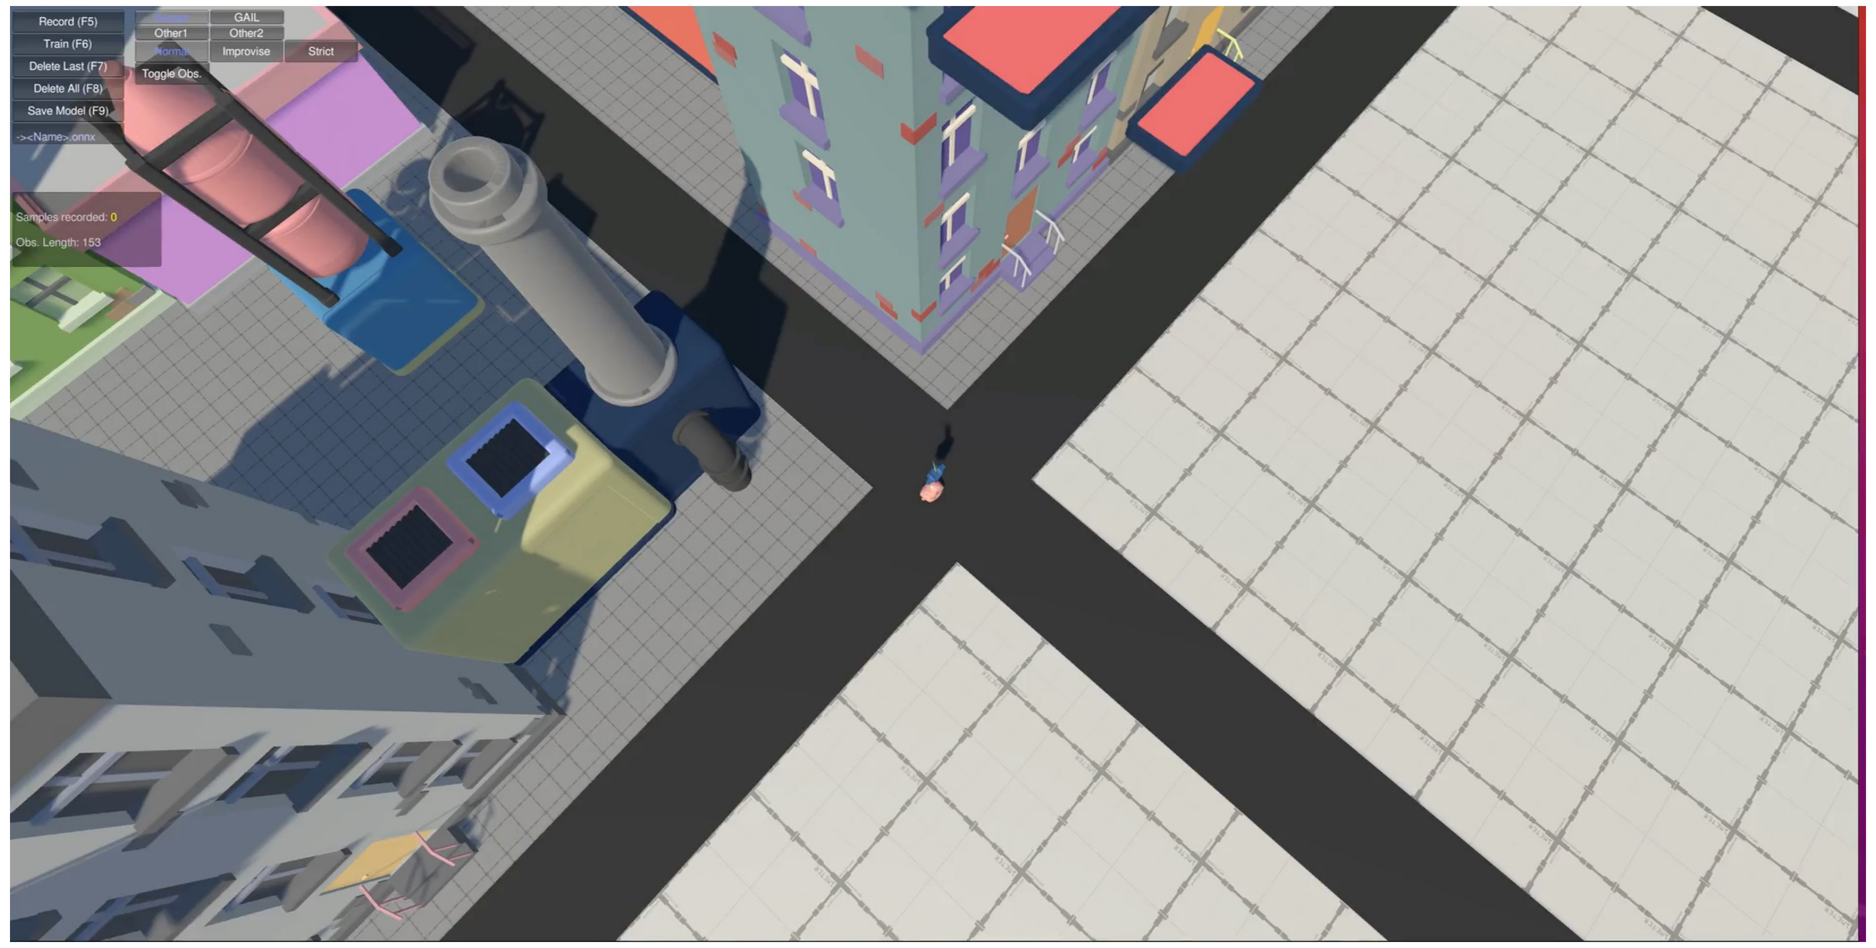

Training Environment

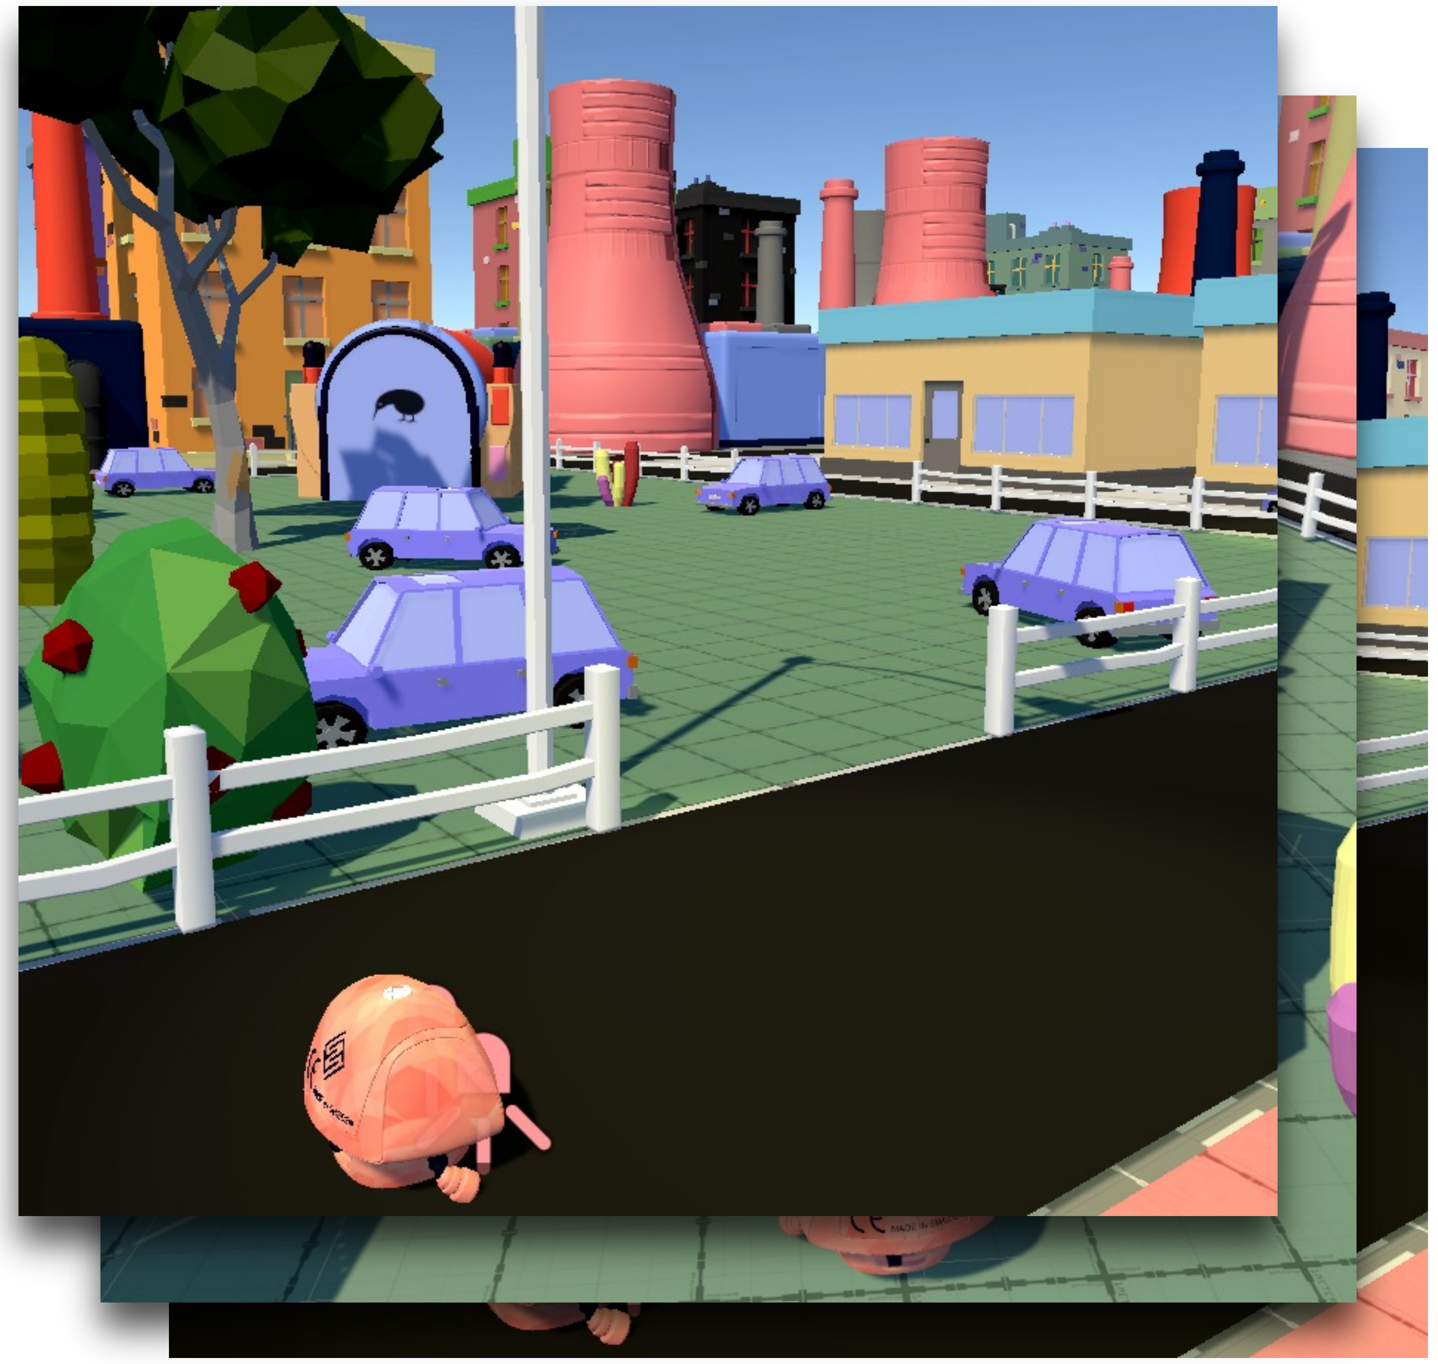

Testing Environment

[3] *Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning*, Sestini et al., 2023

Experiments - Augmentations

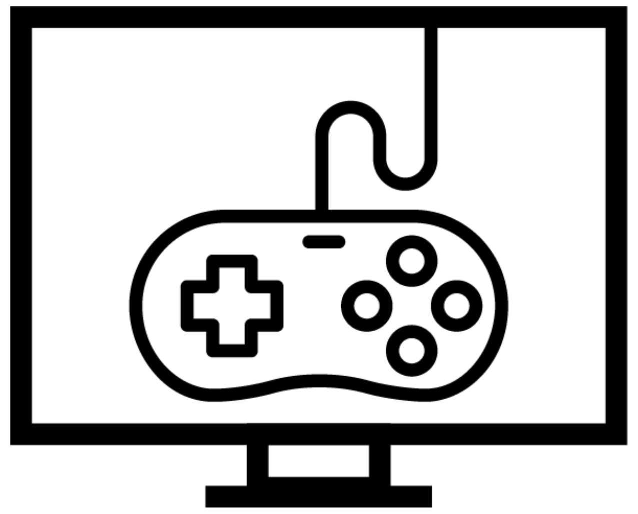

Create
Demonstrations

Experiments - Augmentations

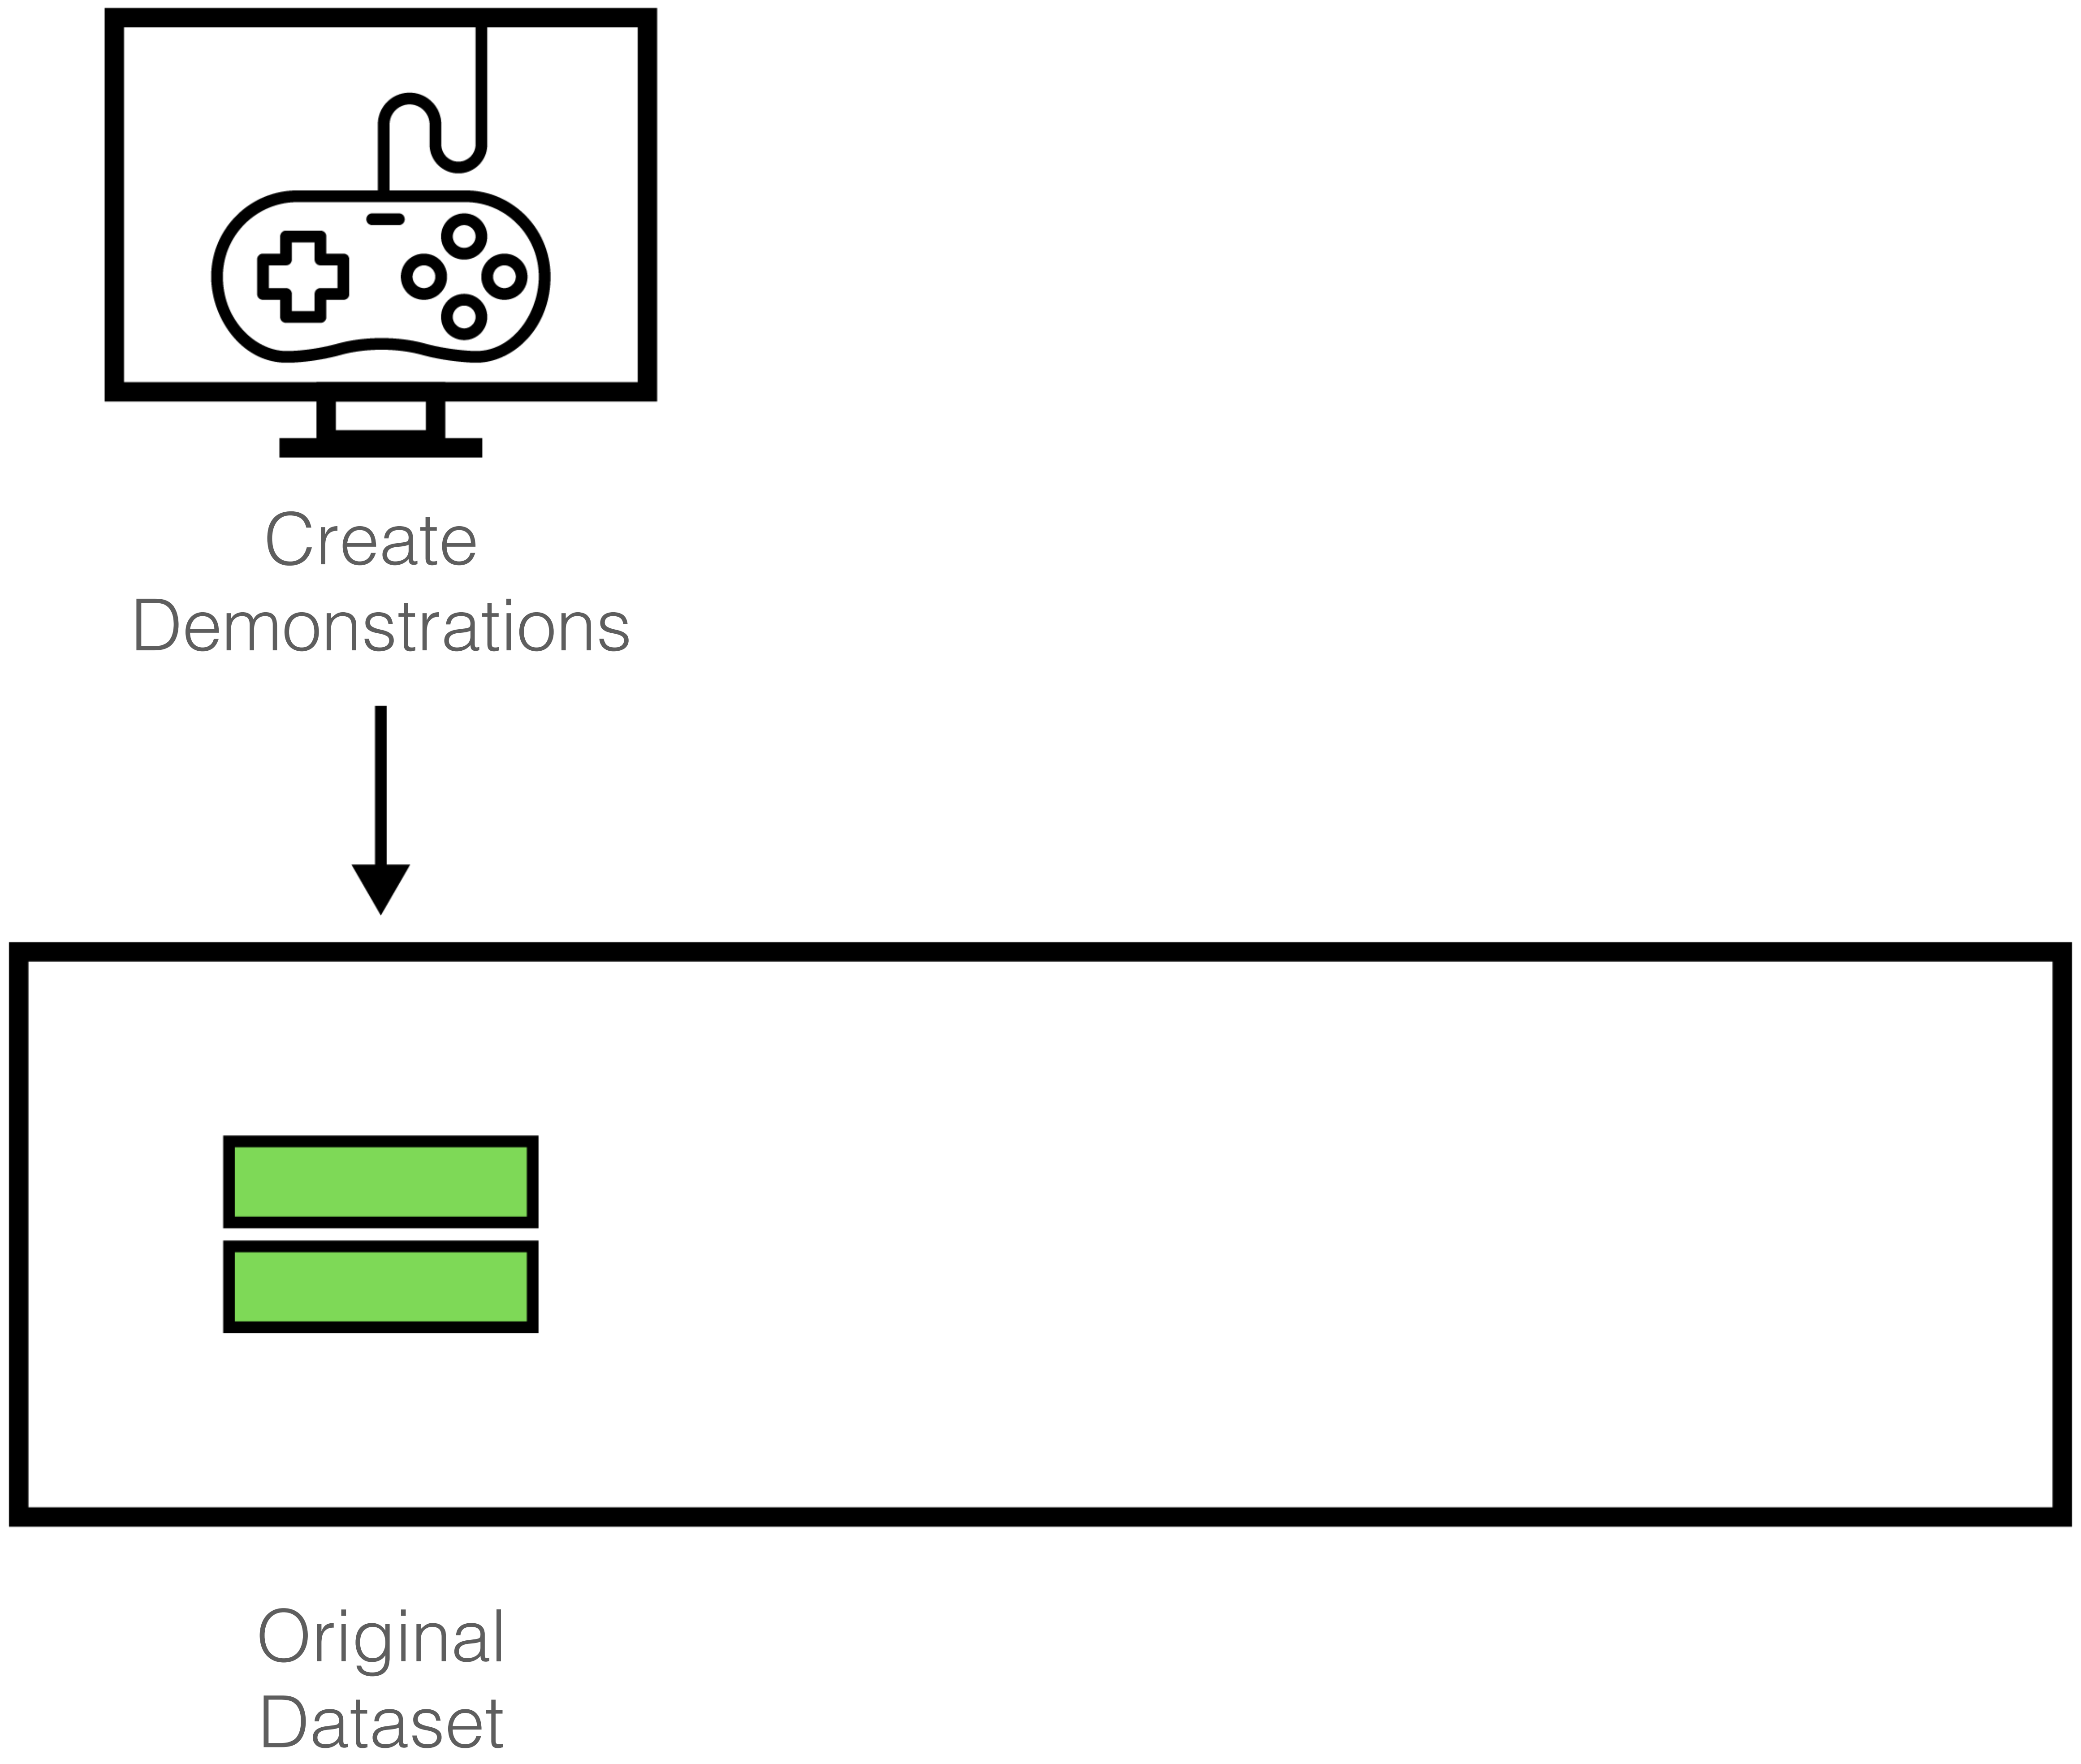

Experiments - Augmentations

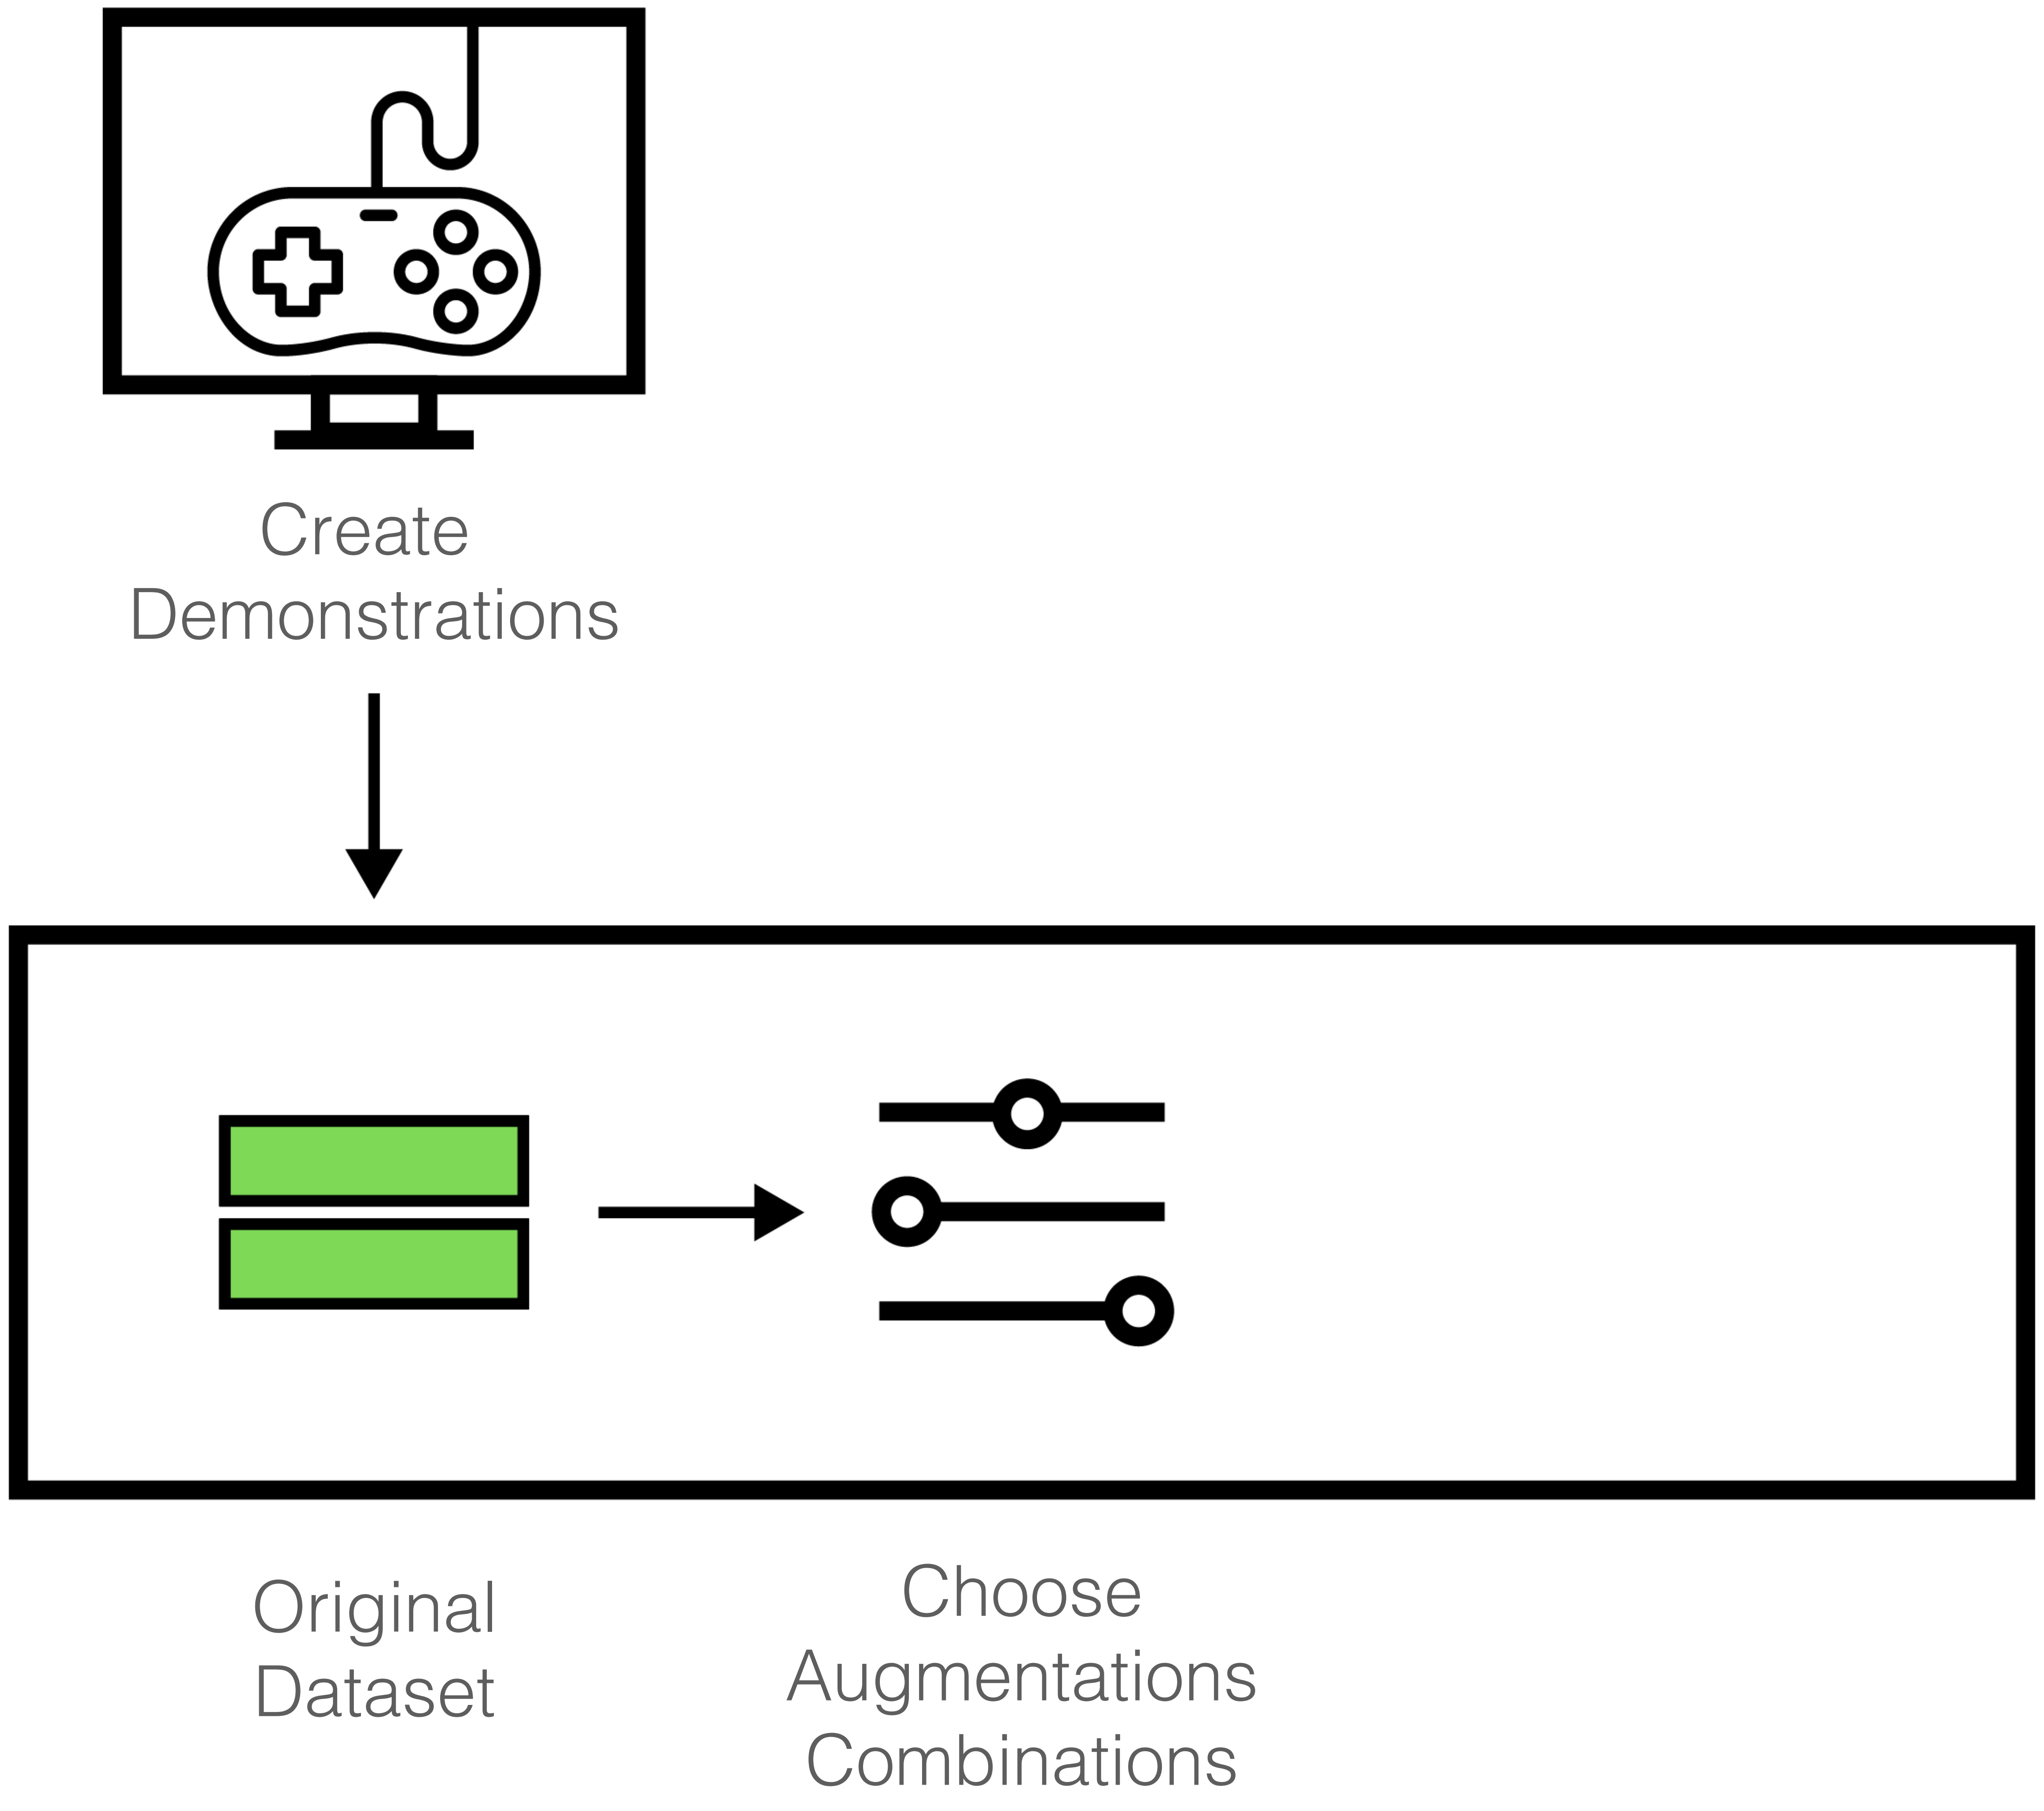

Experiments - Augmentations

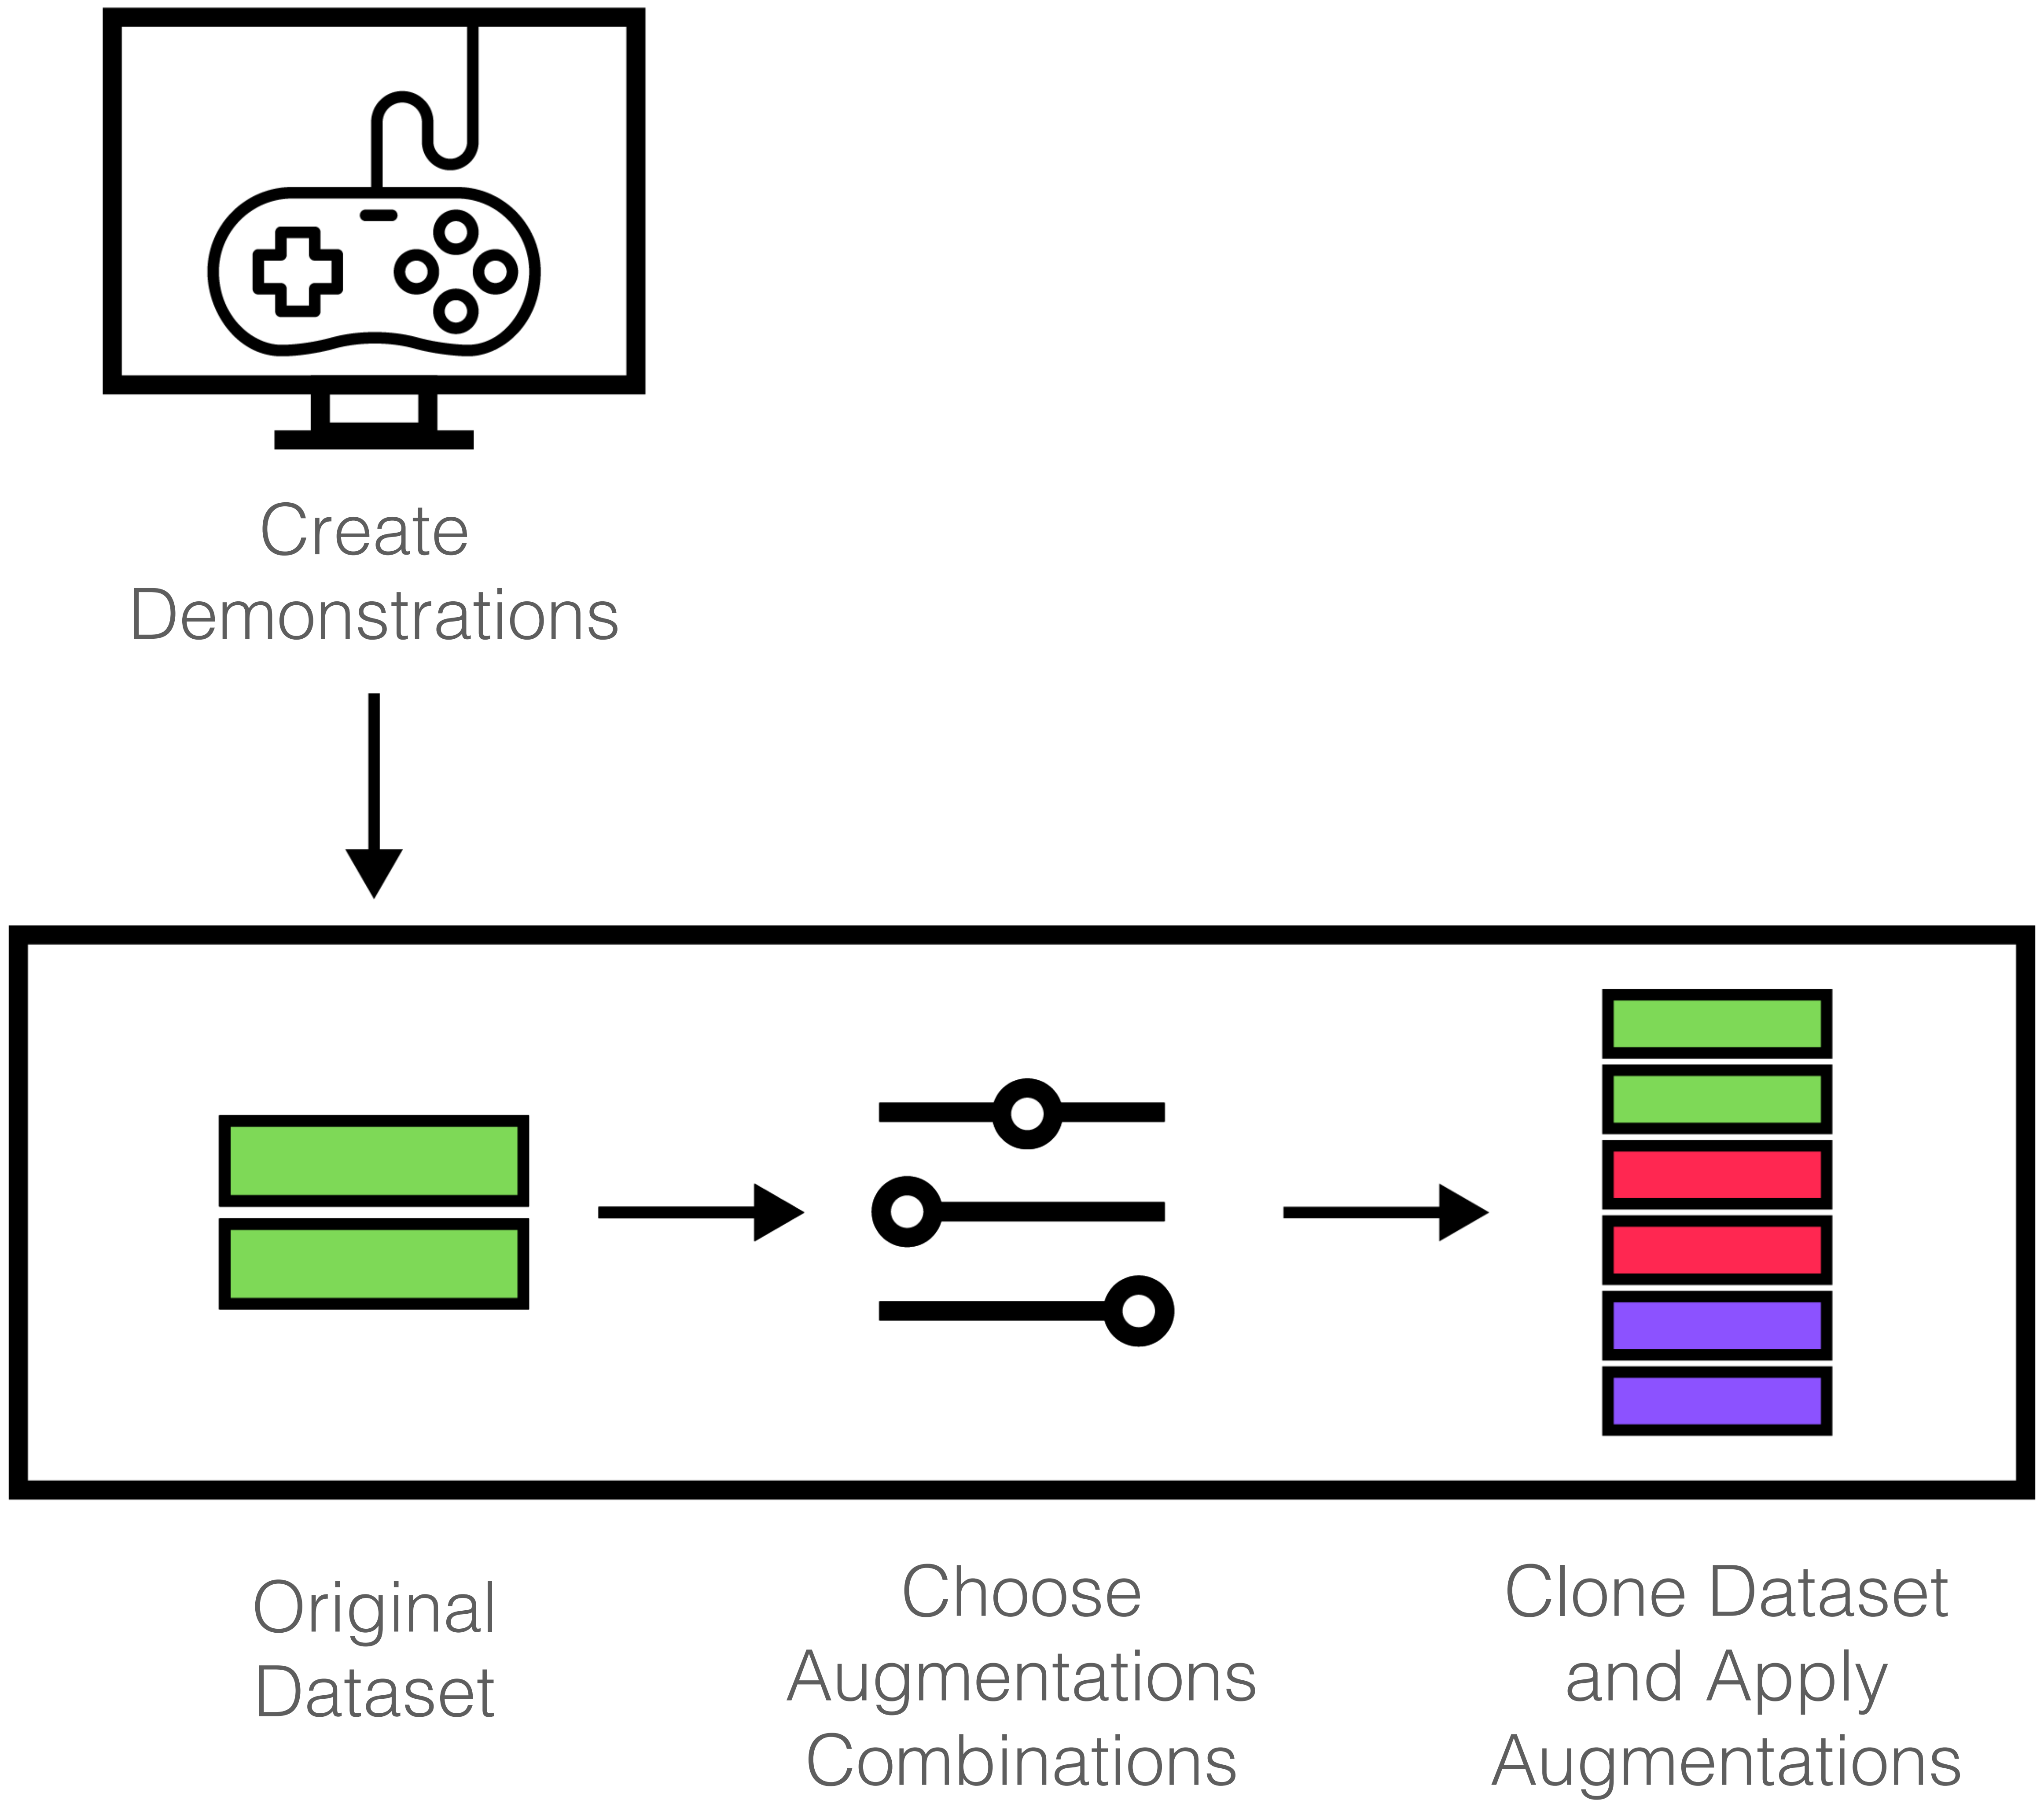

Experiments - Augmentations

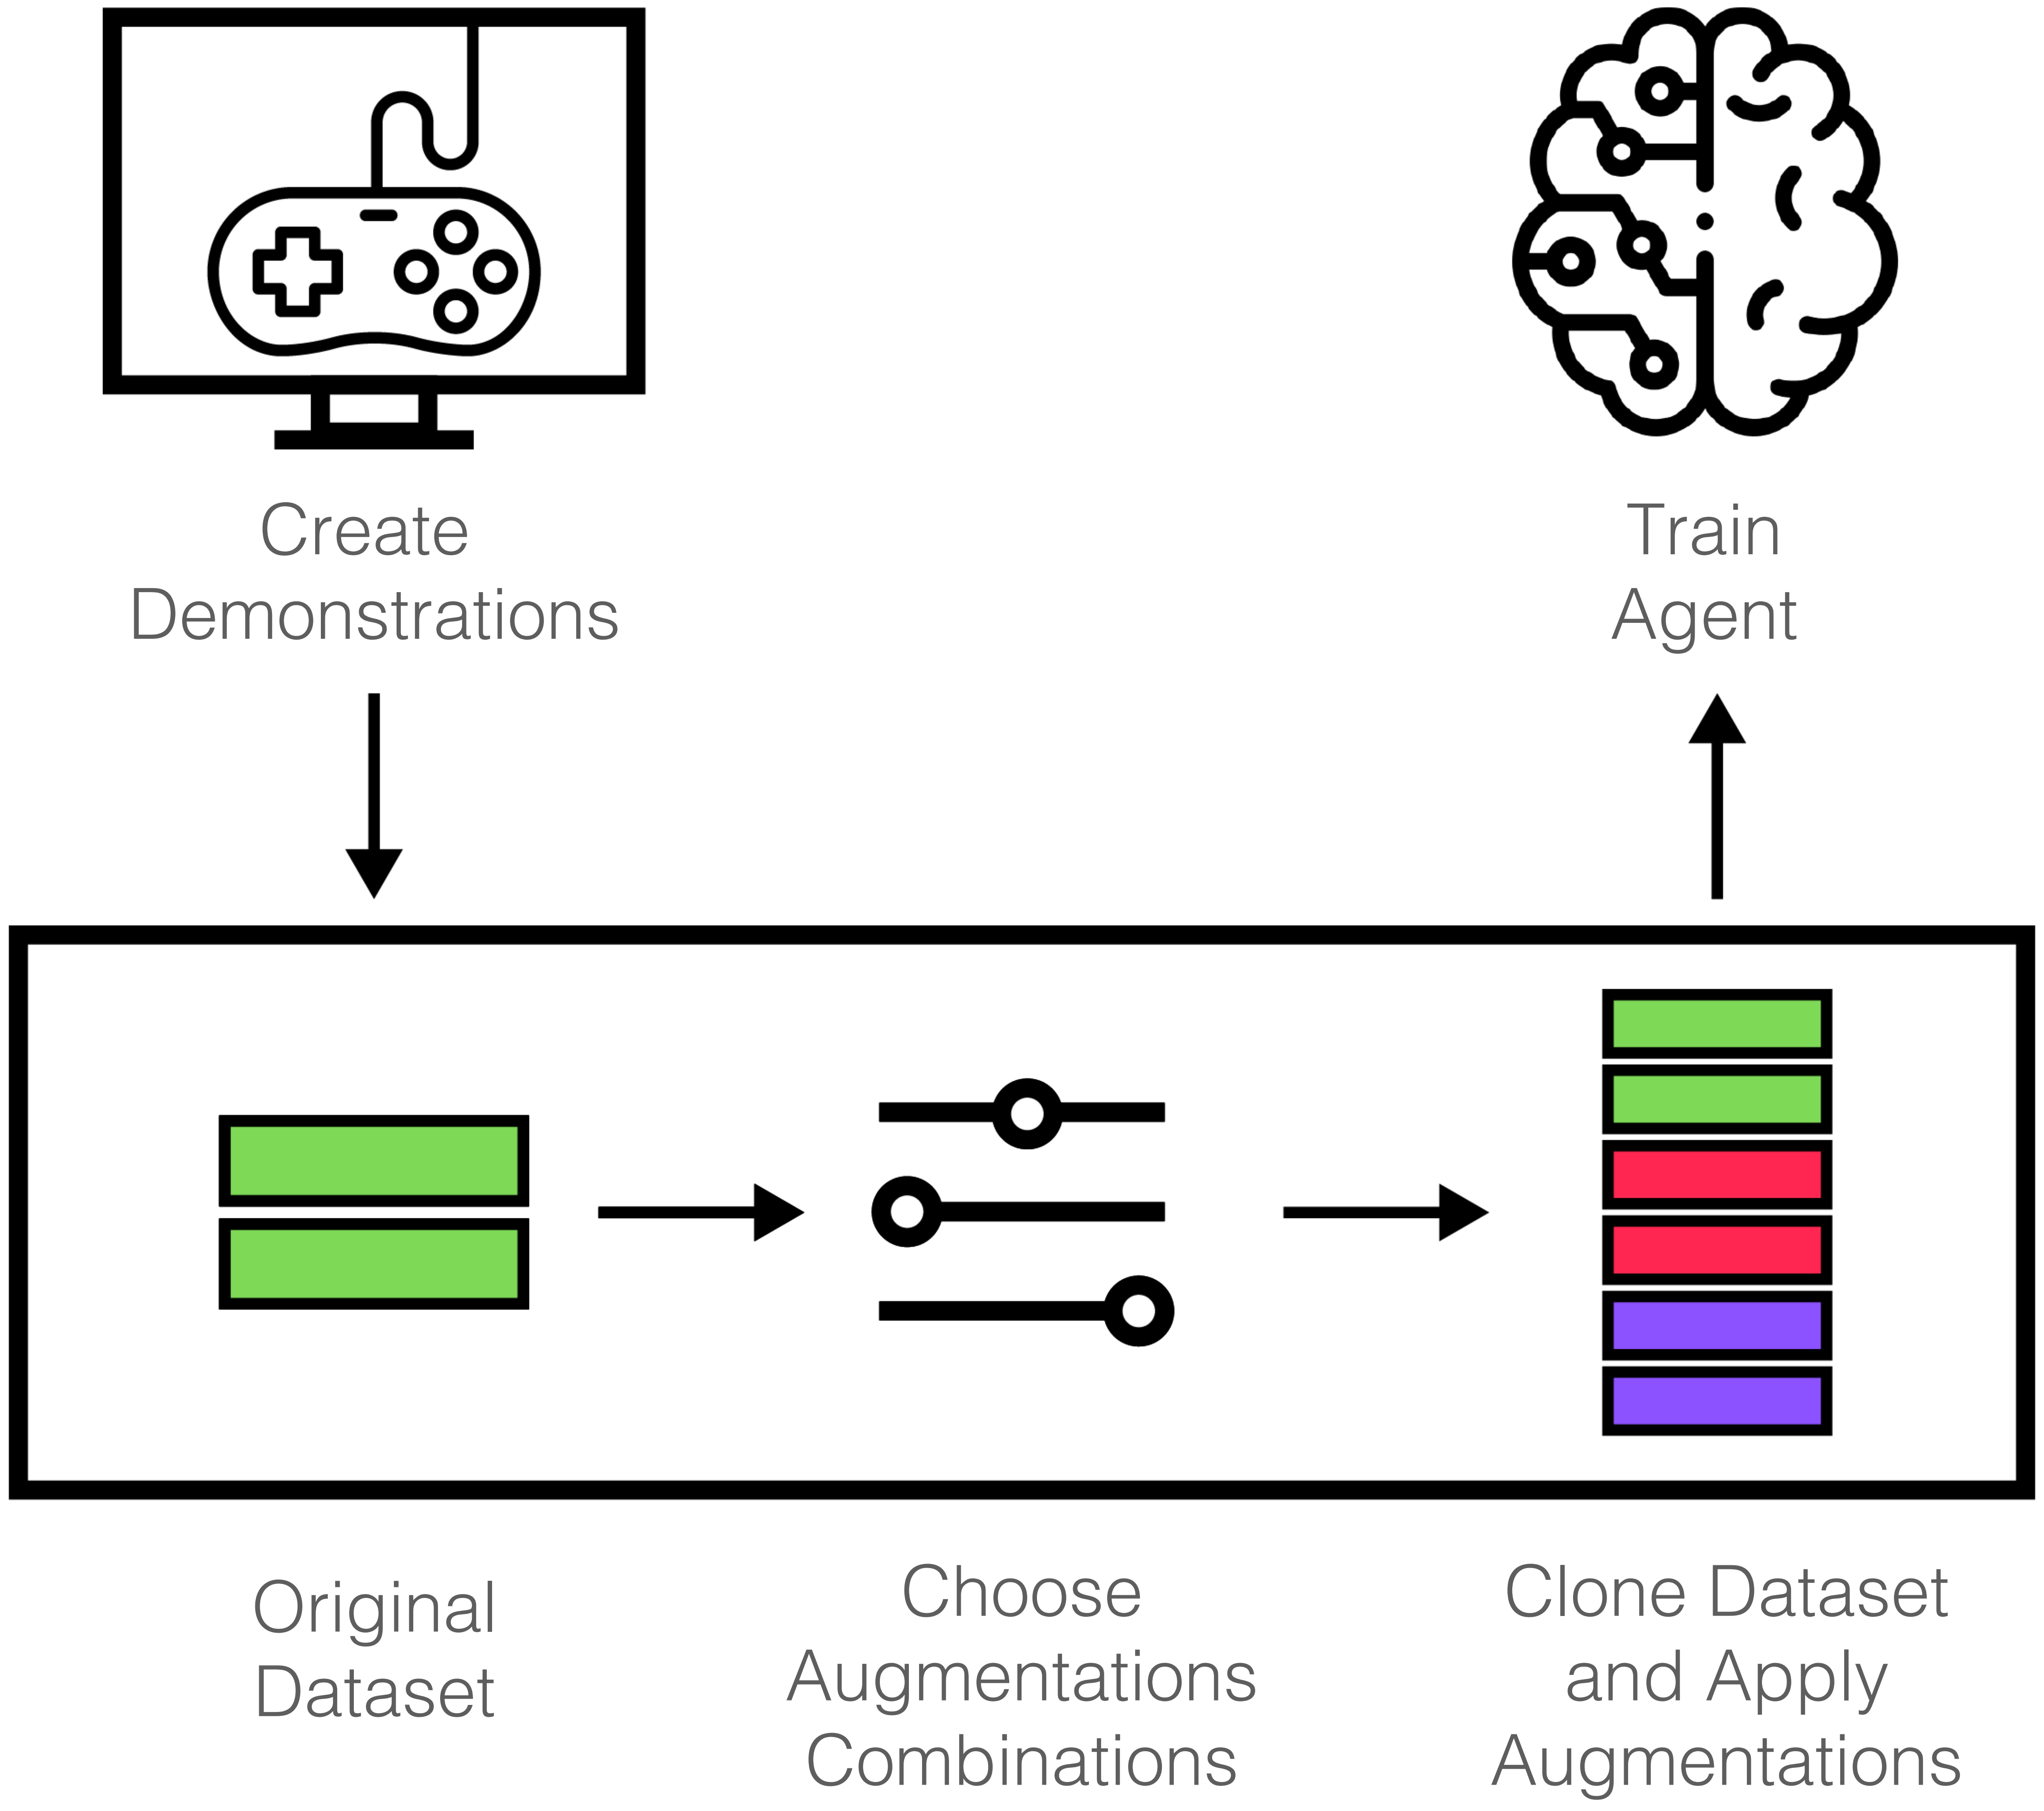

Experiments - State Space

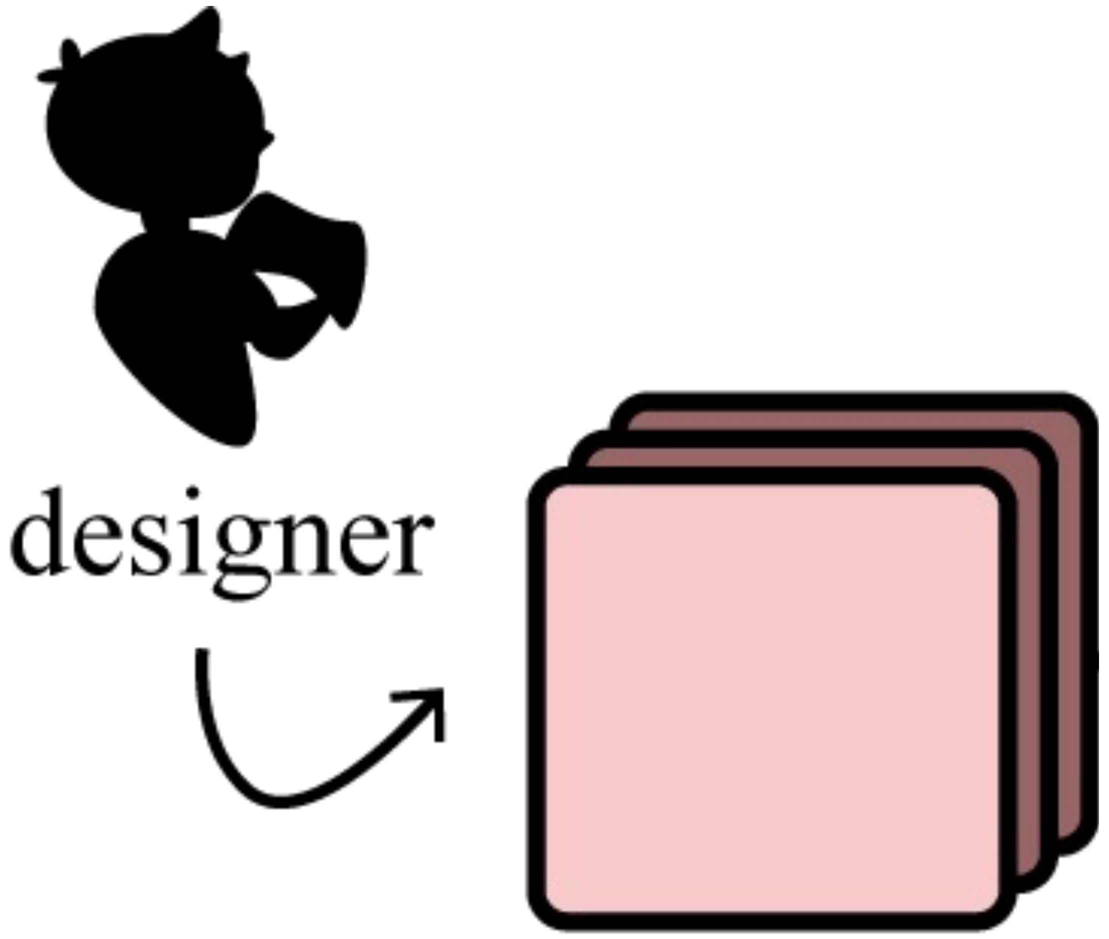

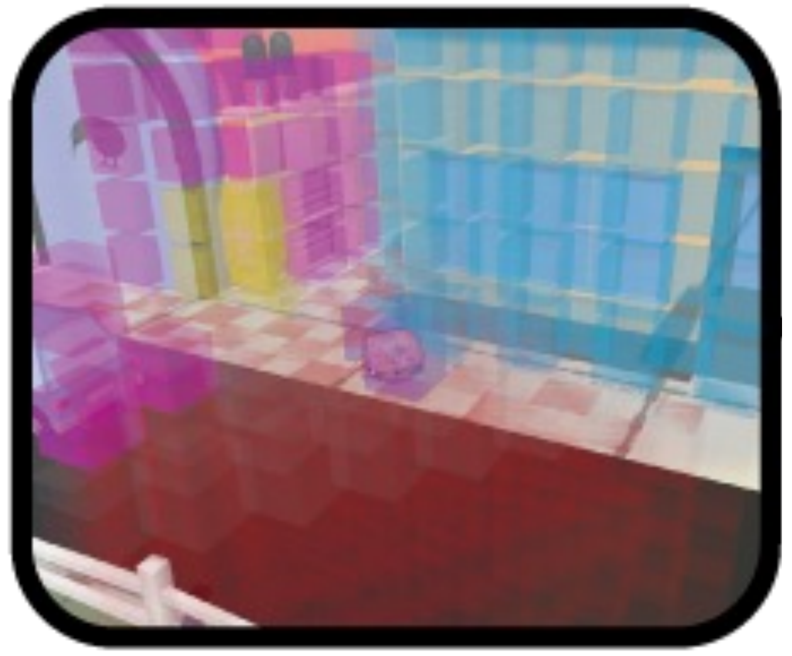

Agent Info	Entities Info	Semantic Map
agent position, current health, ammunitions, …	*objects of interest, relative and absolute positions, …*	*5x5x5 map centered in the position of the agent*

Experiments - Neural Network

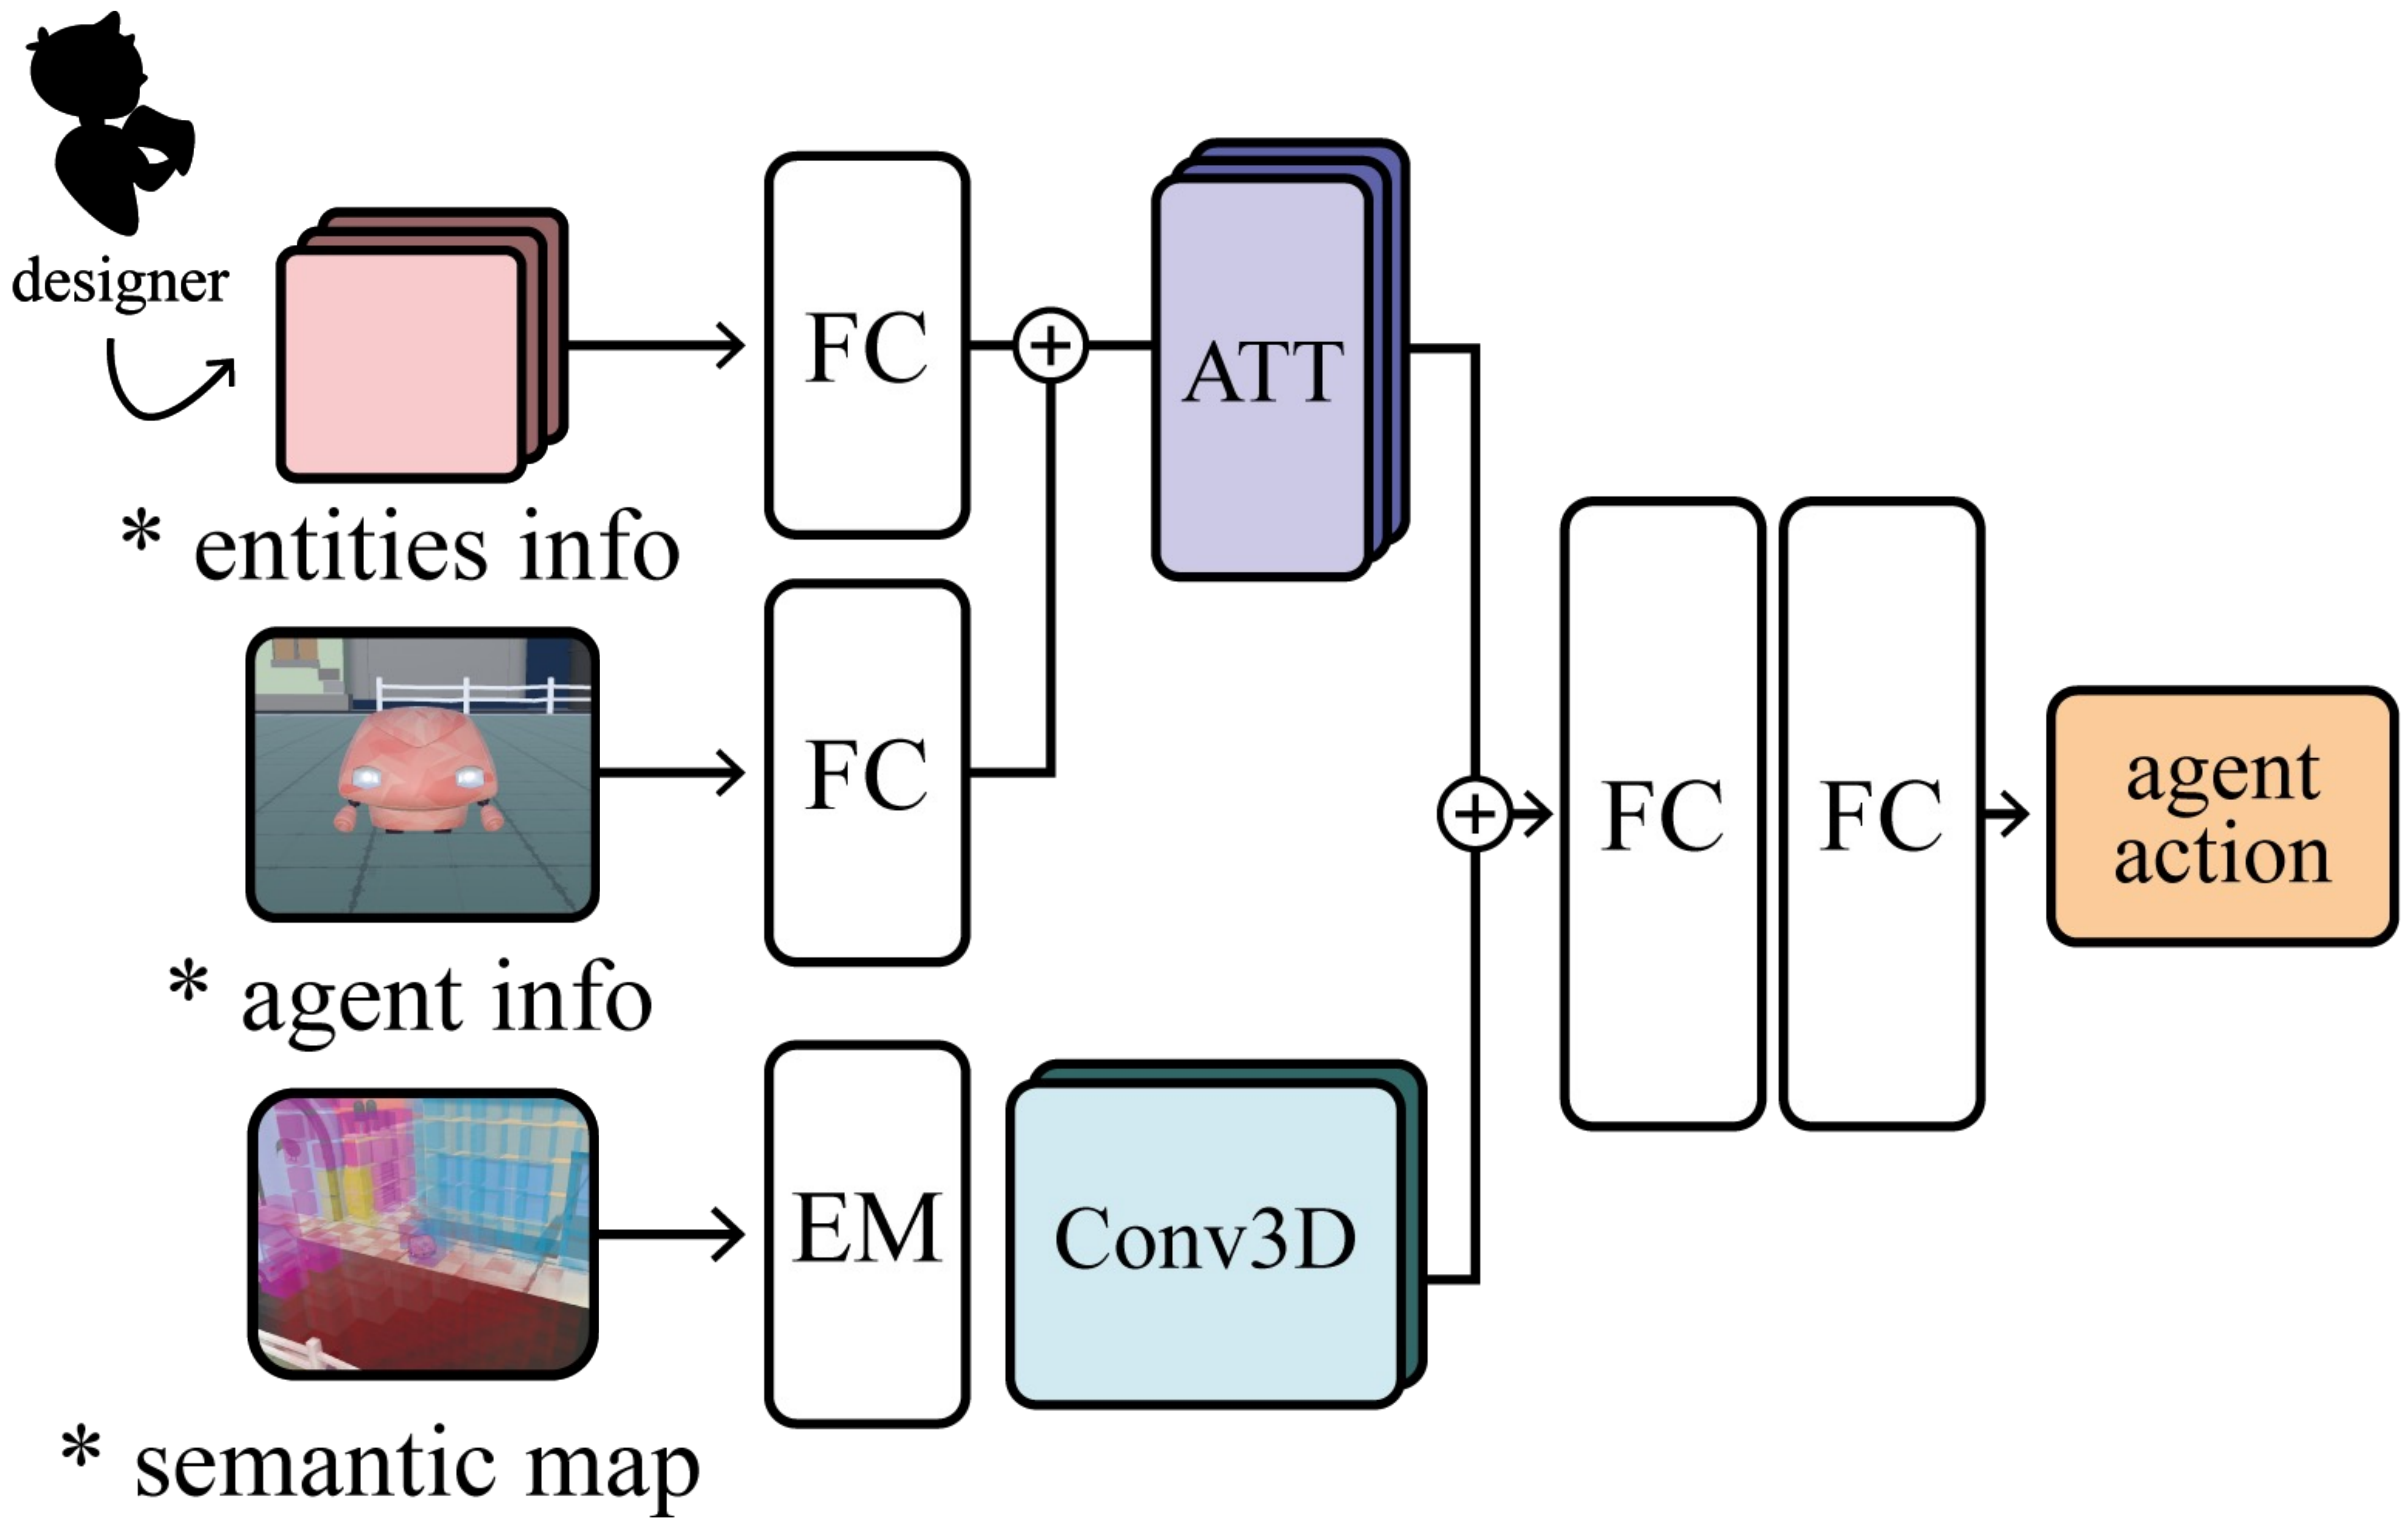

Research Questions

Research Questions

1 *Can we find at least one data augmentation combinations that* ***improve the performance*** *of the original agent, especially in the testing environments?*

Research Questions

1. *Can we find at least one data augmentation combinations that **improve the performance** of the original agent, especially in the testing environments?*

2. *What is the best combination of augmentation that has the **highest performance** on all the testing environments?*

Research Questions

1. *Can we find at least one data augmentation combinations that **improve the performance** of the original agent, especially in the testing environments?*

2. *What is the best combination of augmentation that has the **highest performance** on all the testing environments?*

3. *What is the **single most effective** augmentation? Is there a single most effective augmentation?*

1

Experiments - Quantitative Results

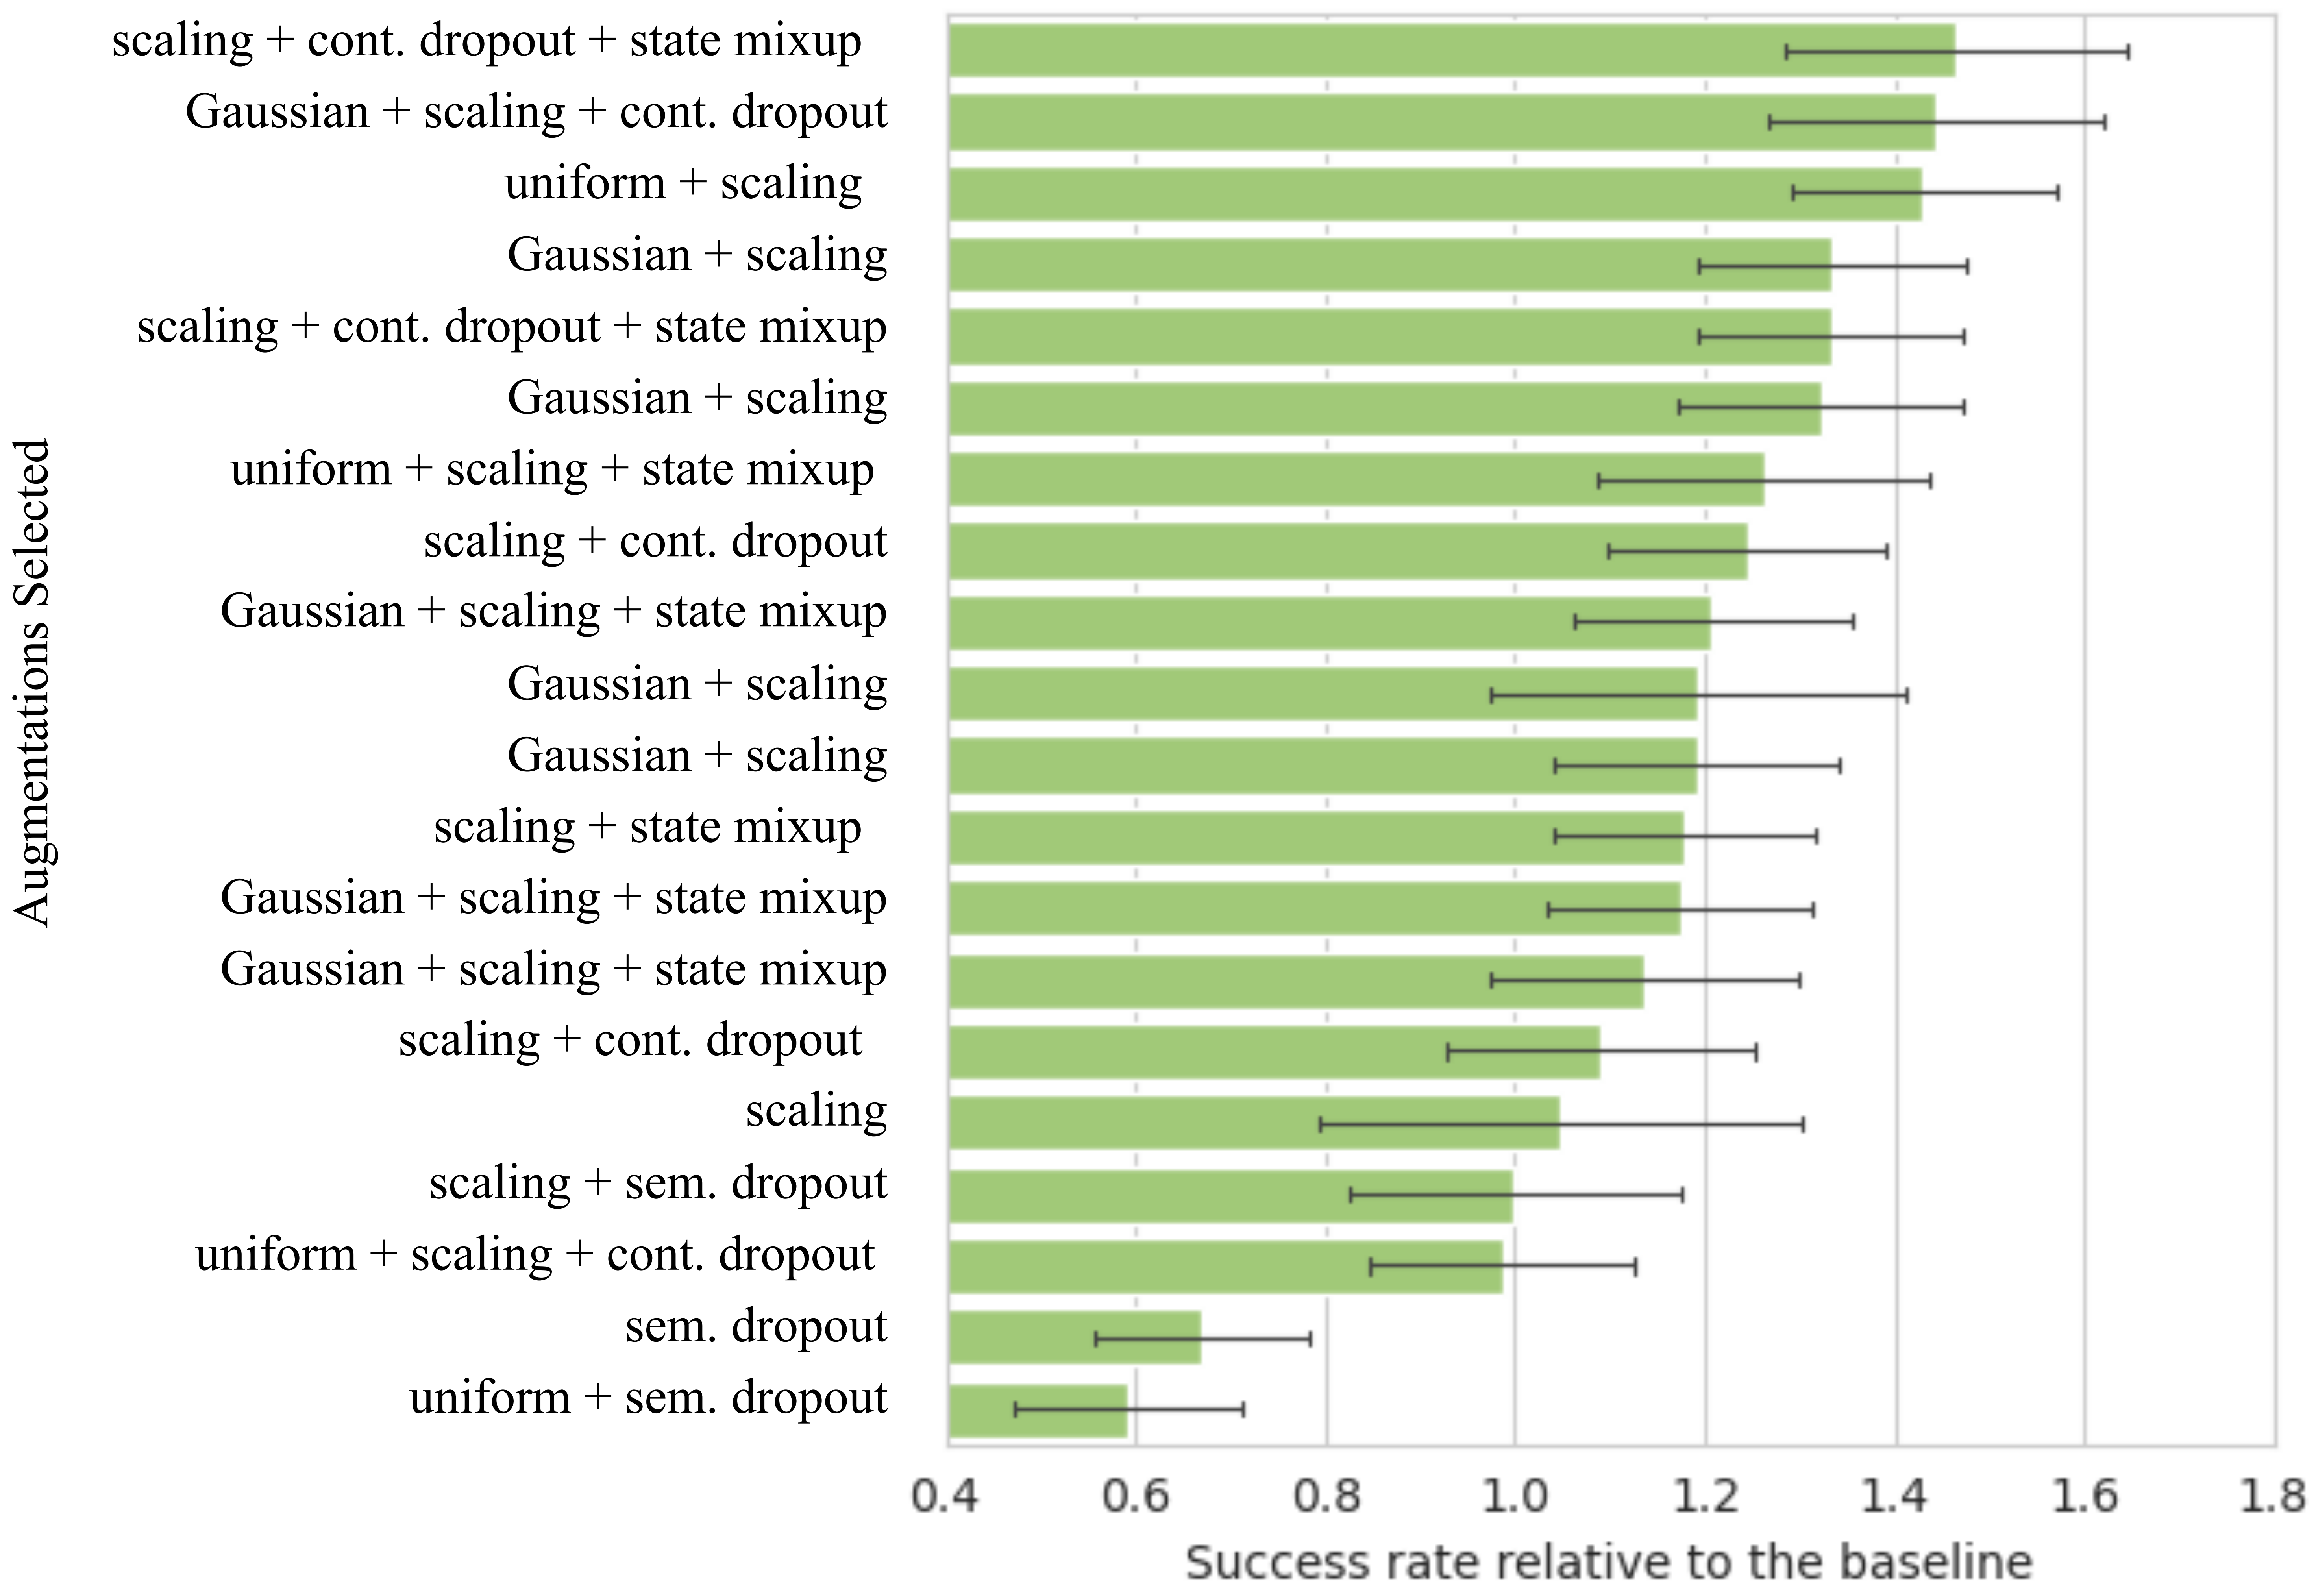

1 Experiments - Quantitative Results

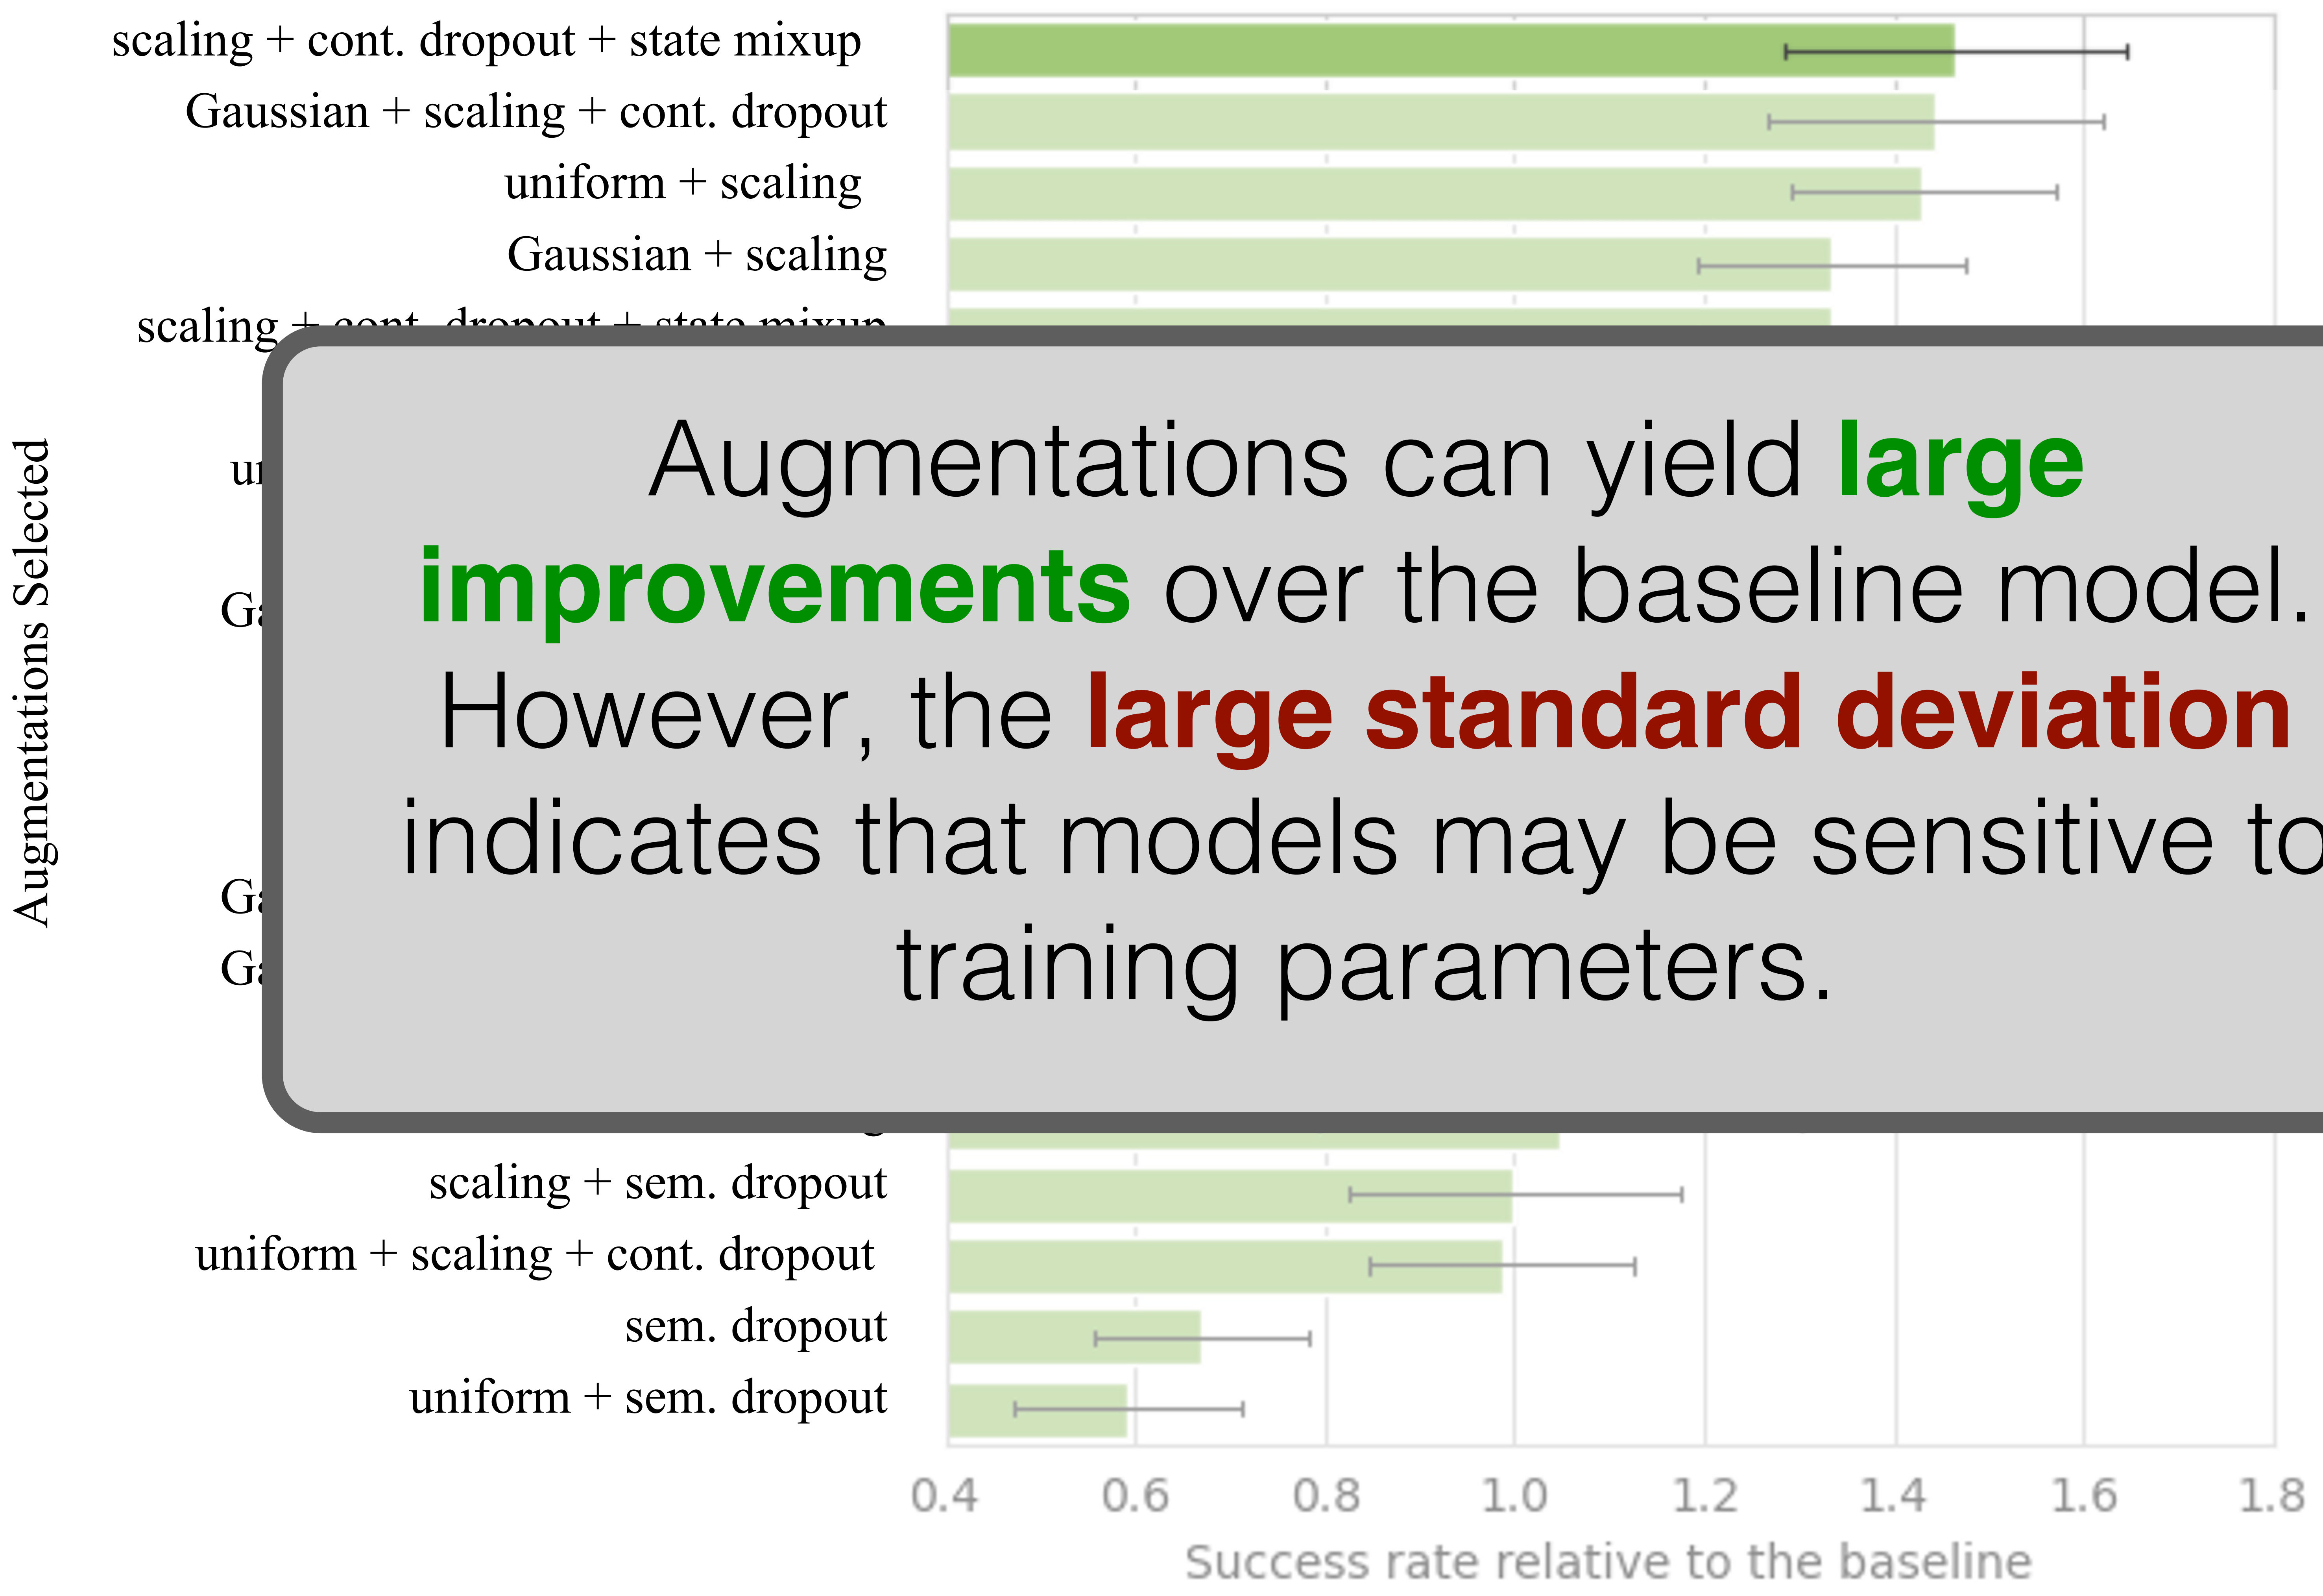

Augmentations can yield **large improvements** over the baseline model. However, the **large standard deviation** indicates that models may be sensitive to training parameters.

2 Experiments - Consistency

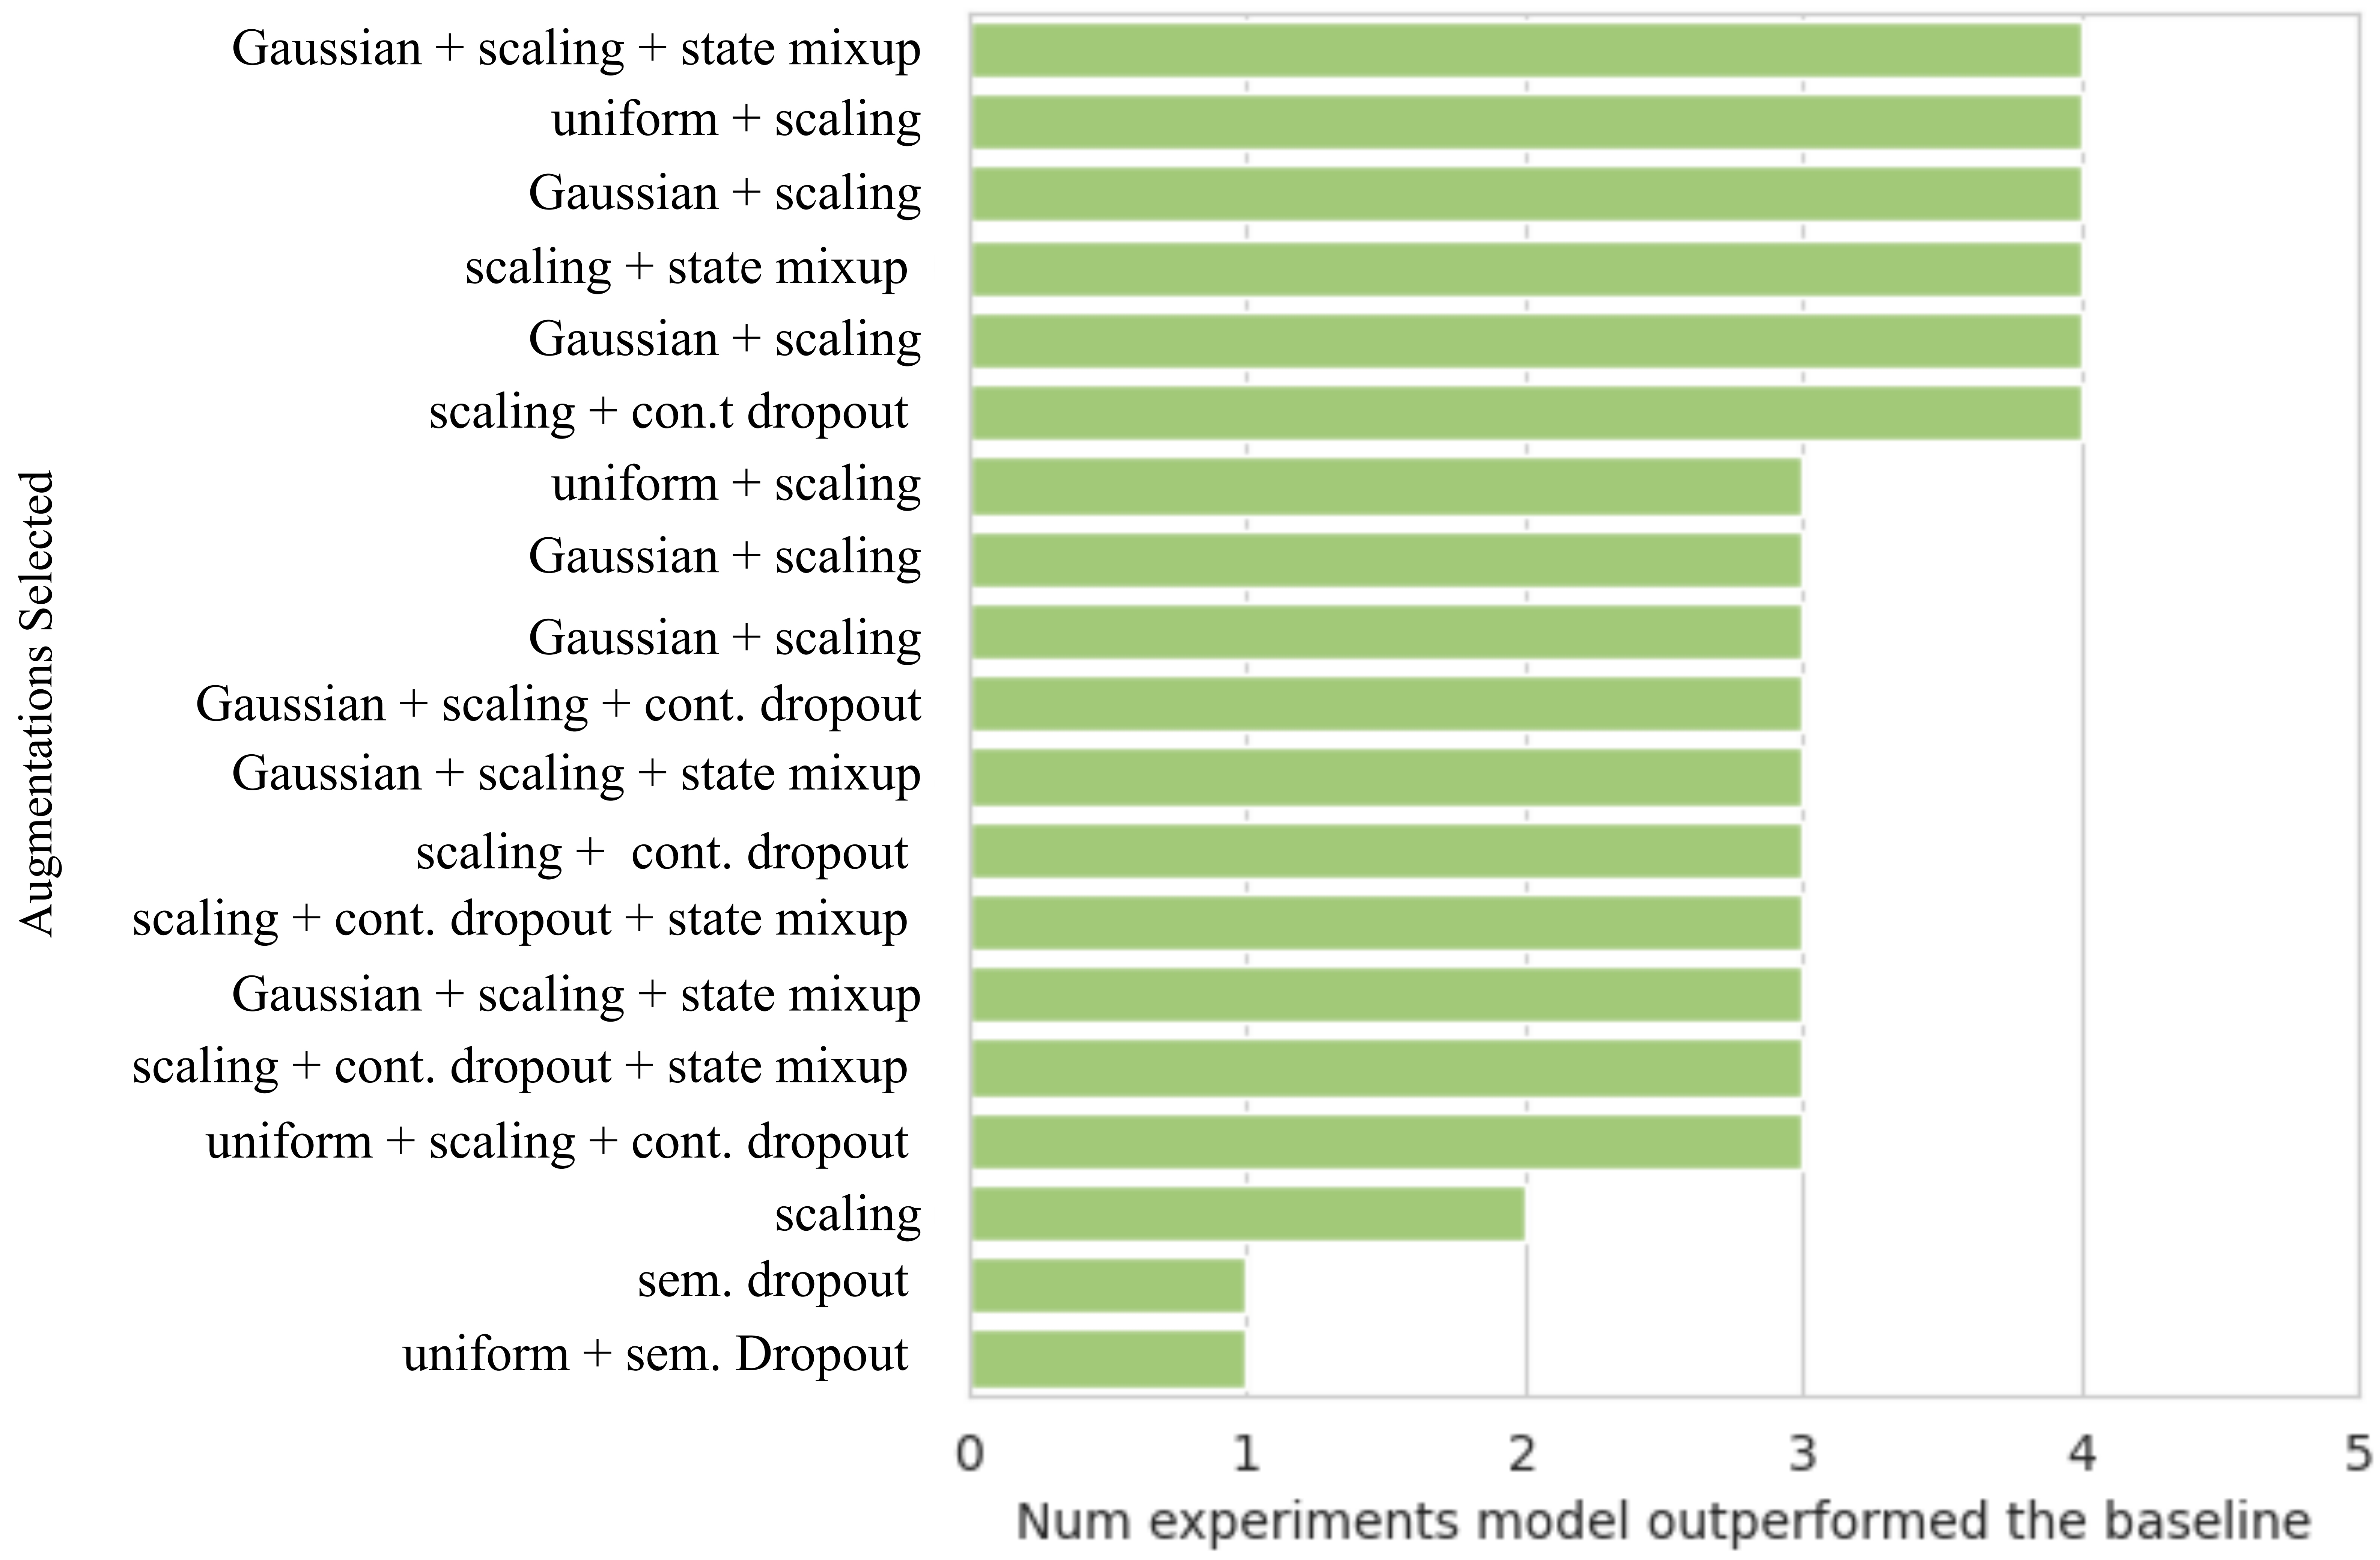

2 Experiments - Consistency

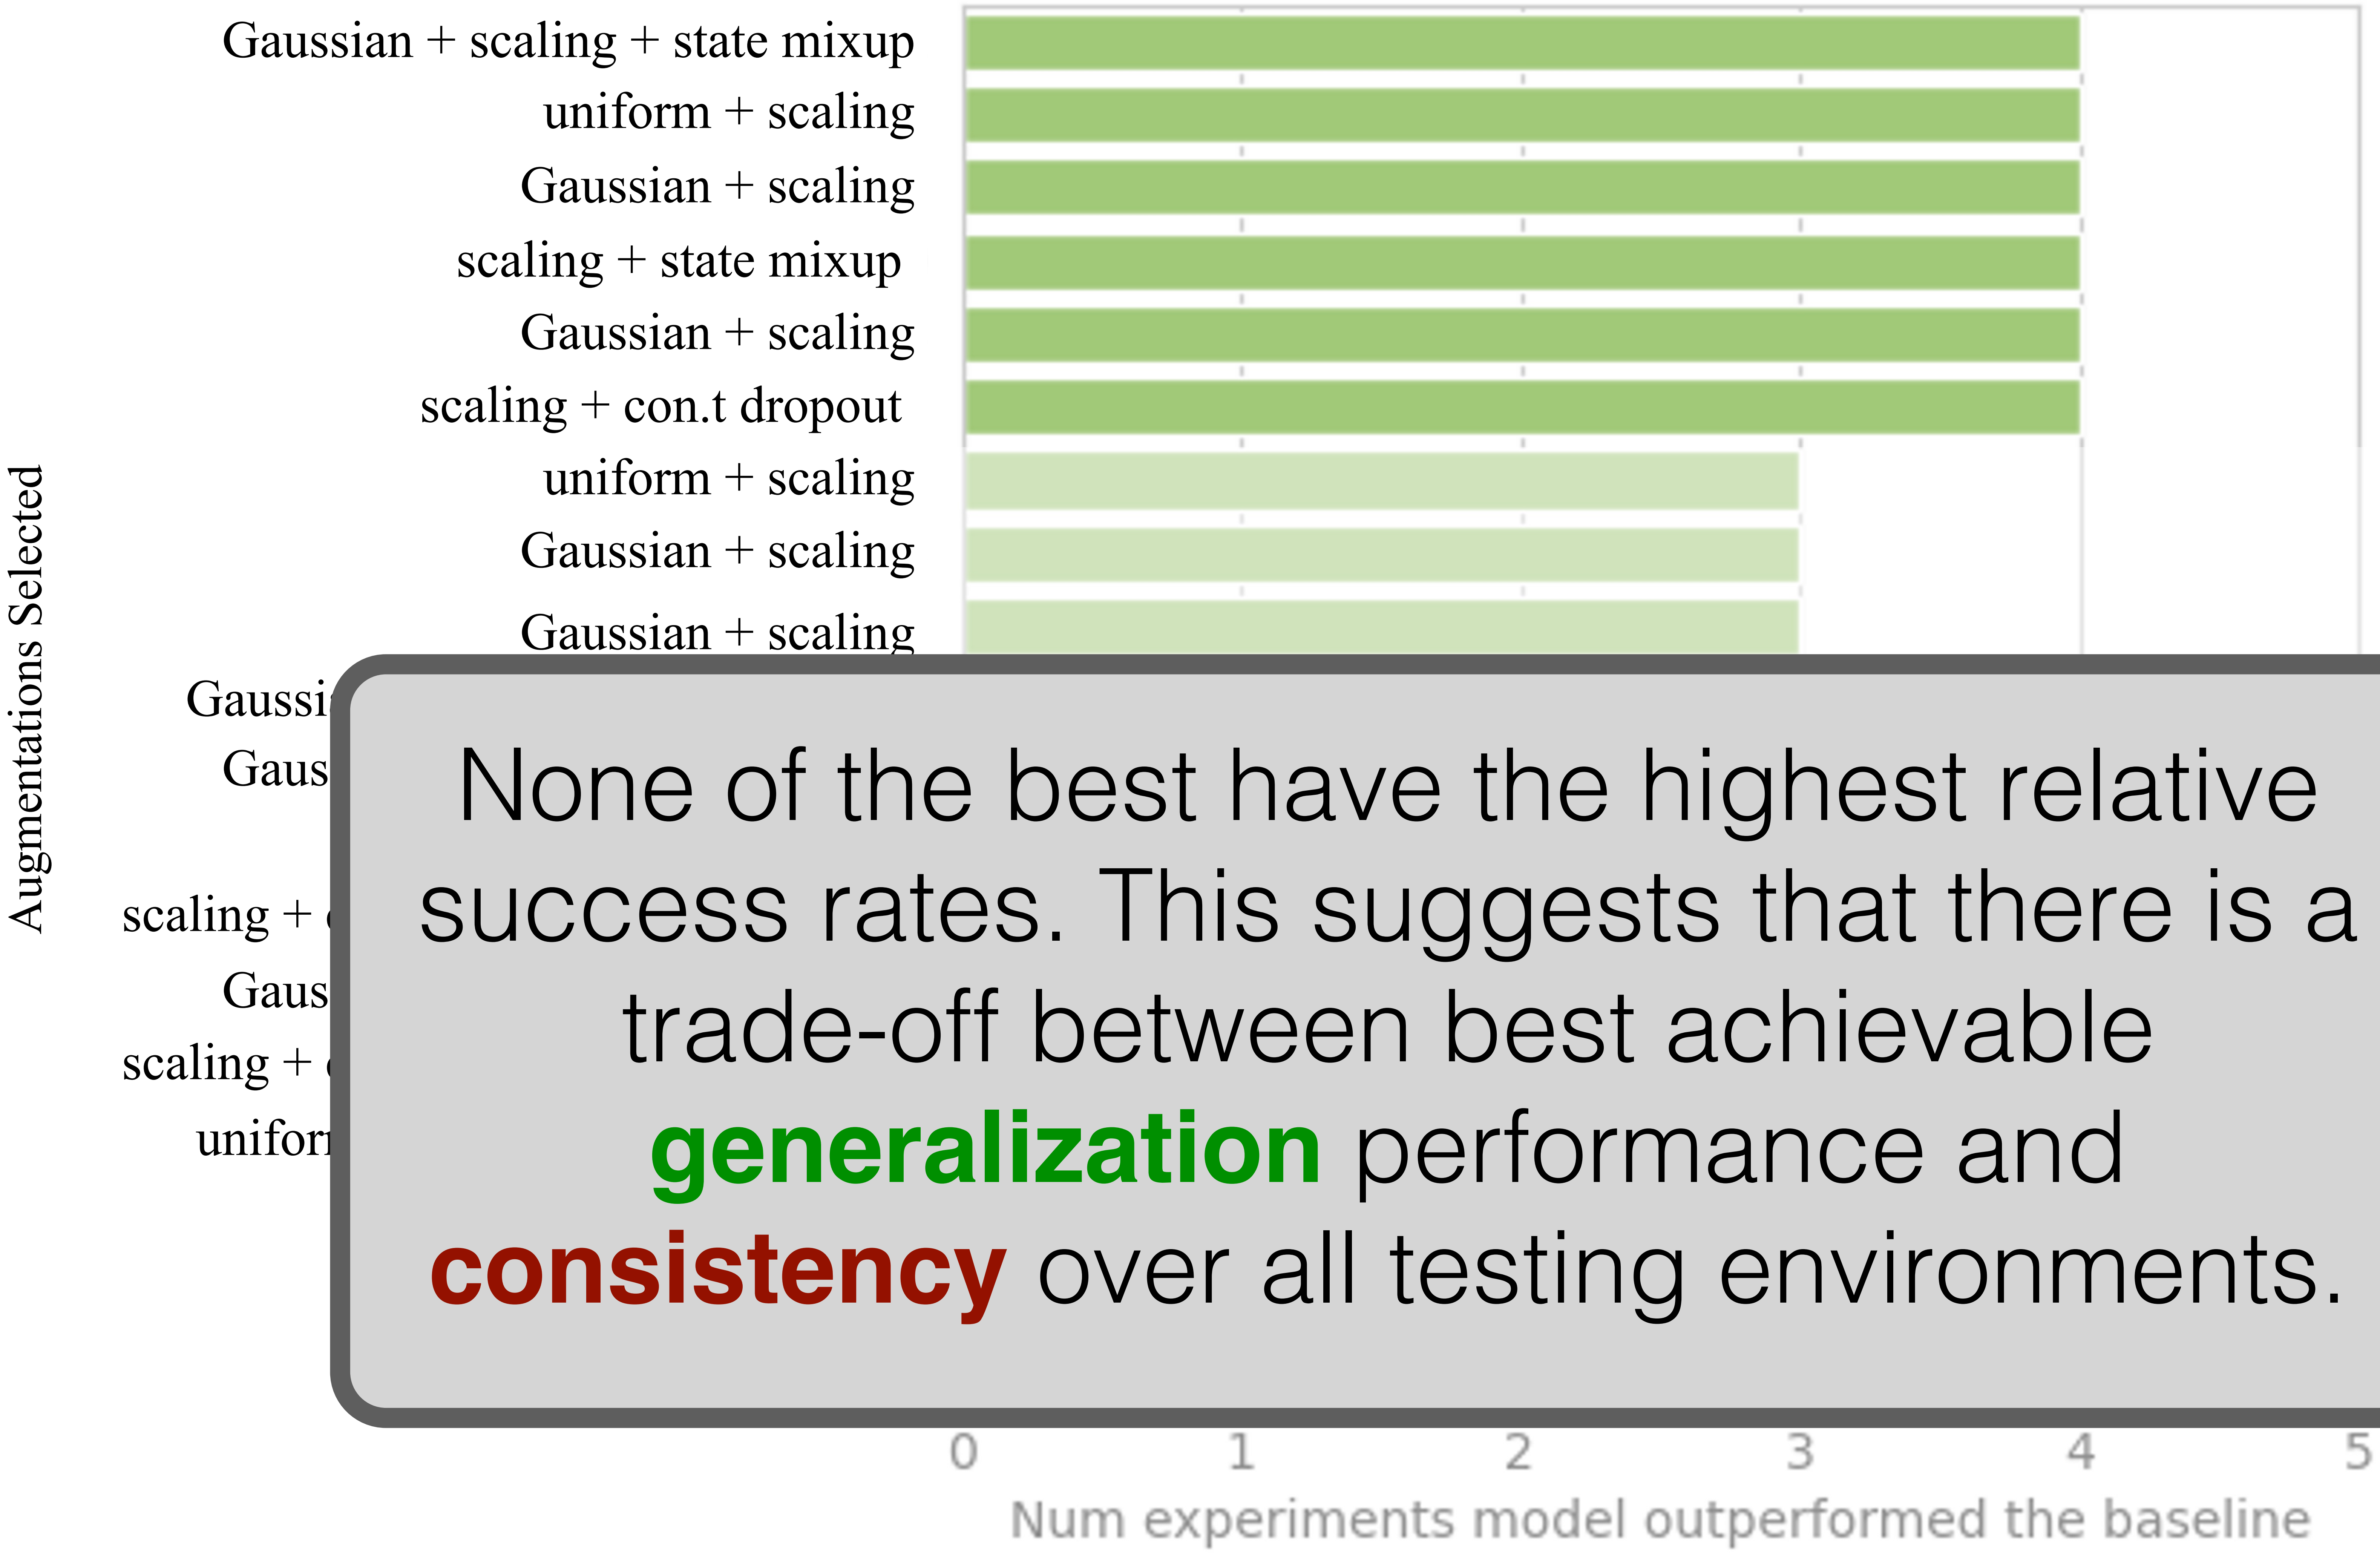

None of the best have the highest relative success rates. This suggests that there is a trade-off between best achievable **generalization** performance and **consistency** over all testing environments.

3 Experiments - Consistency

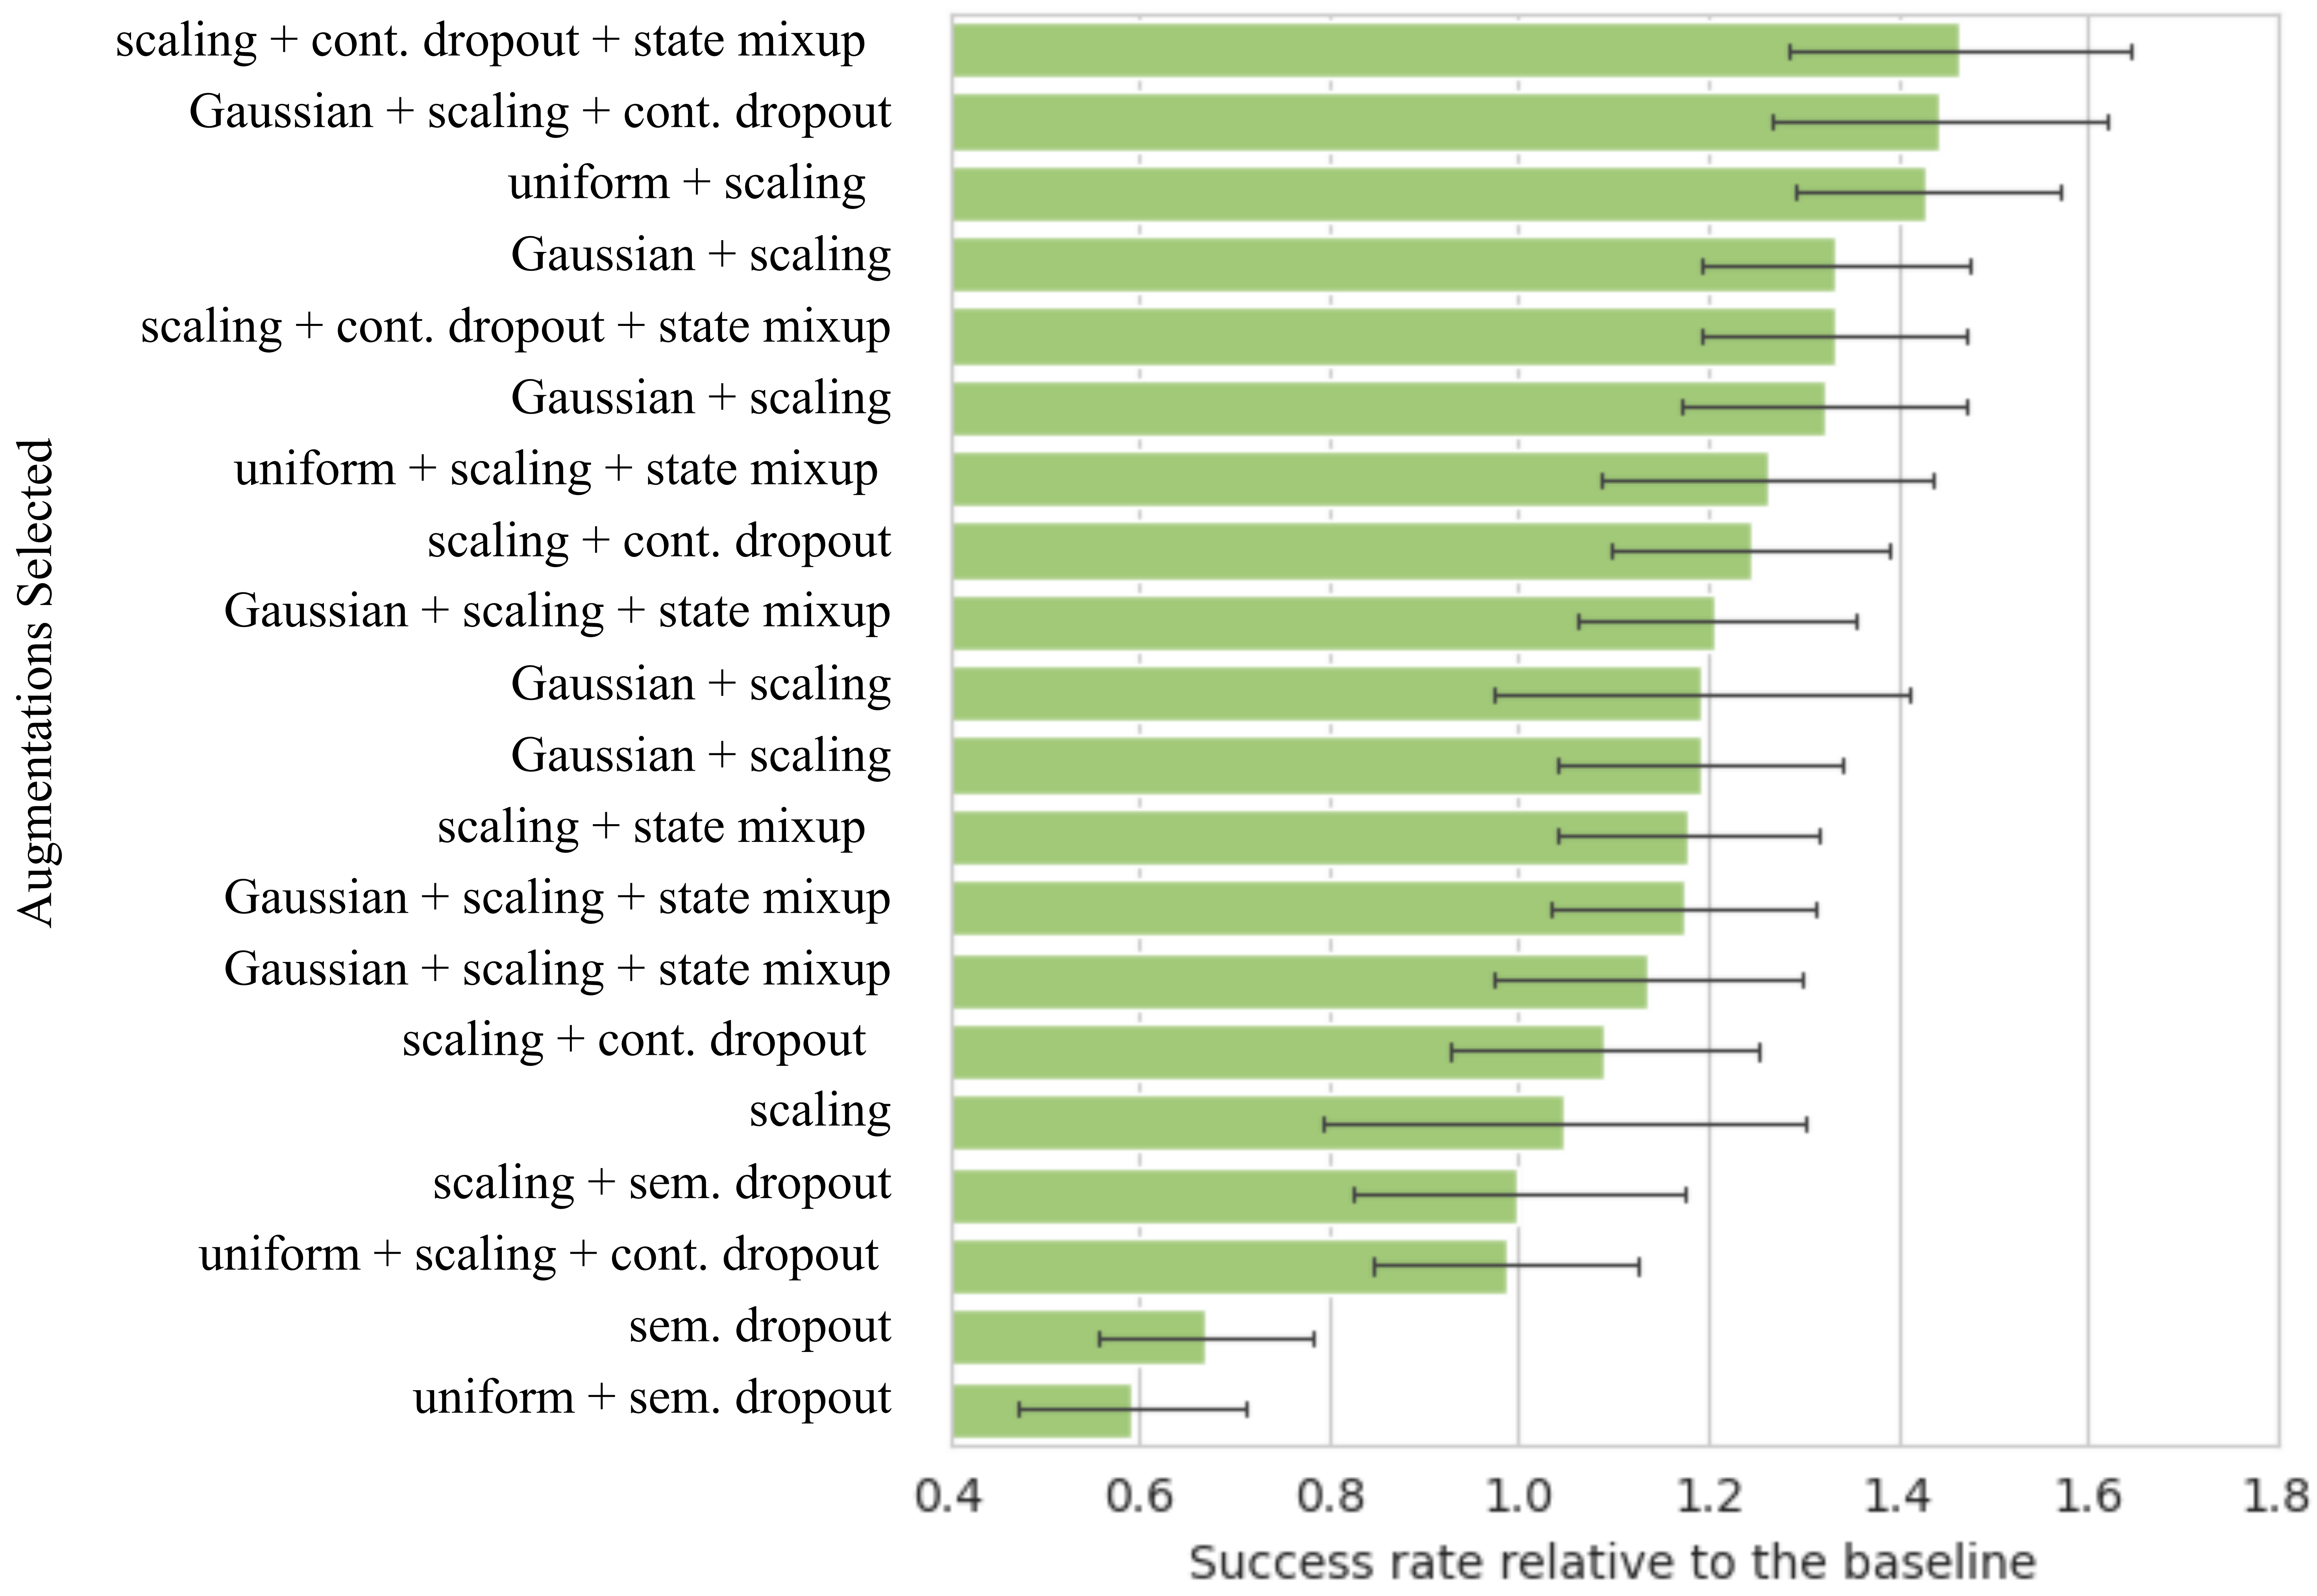

3 Experiments - Consistency

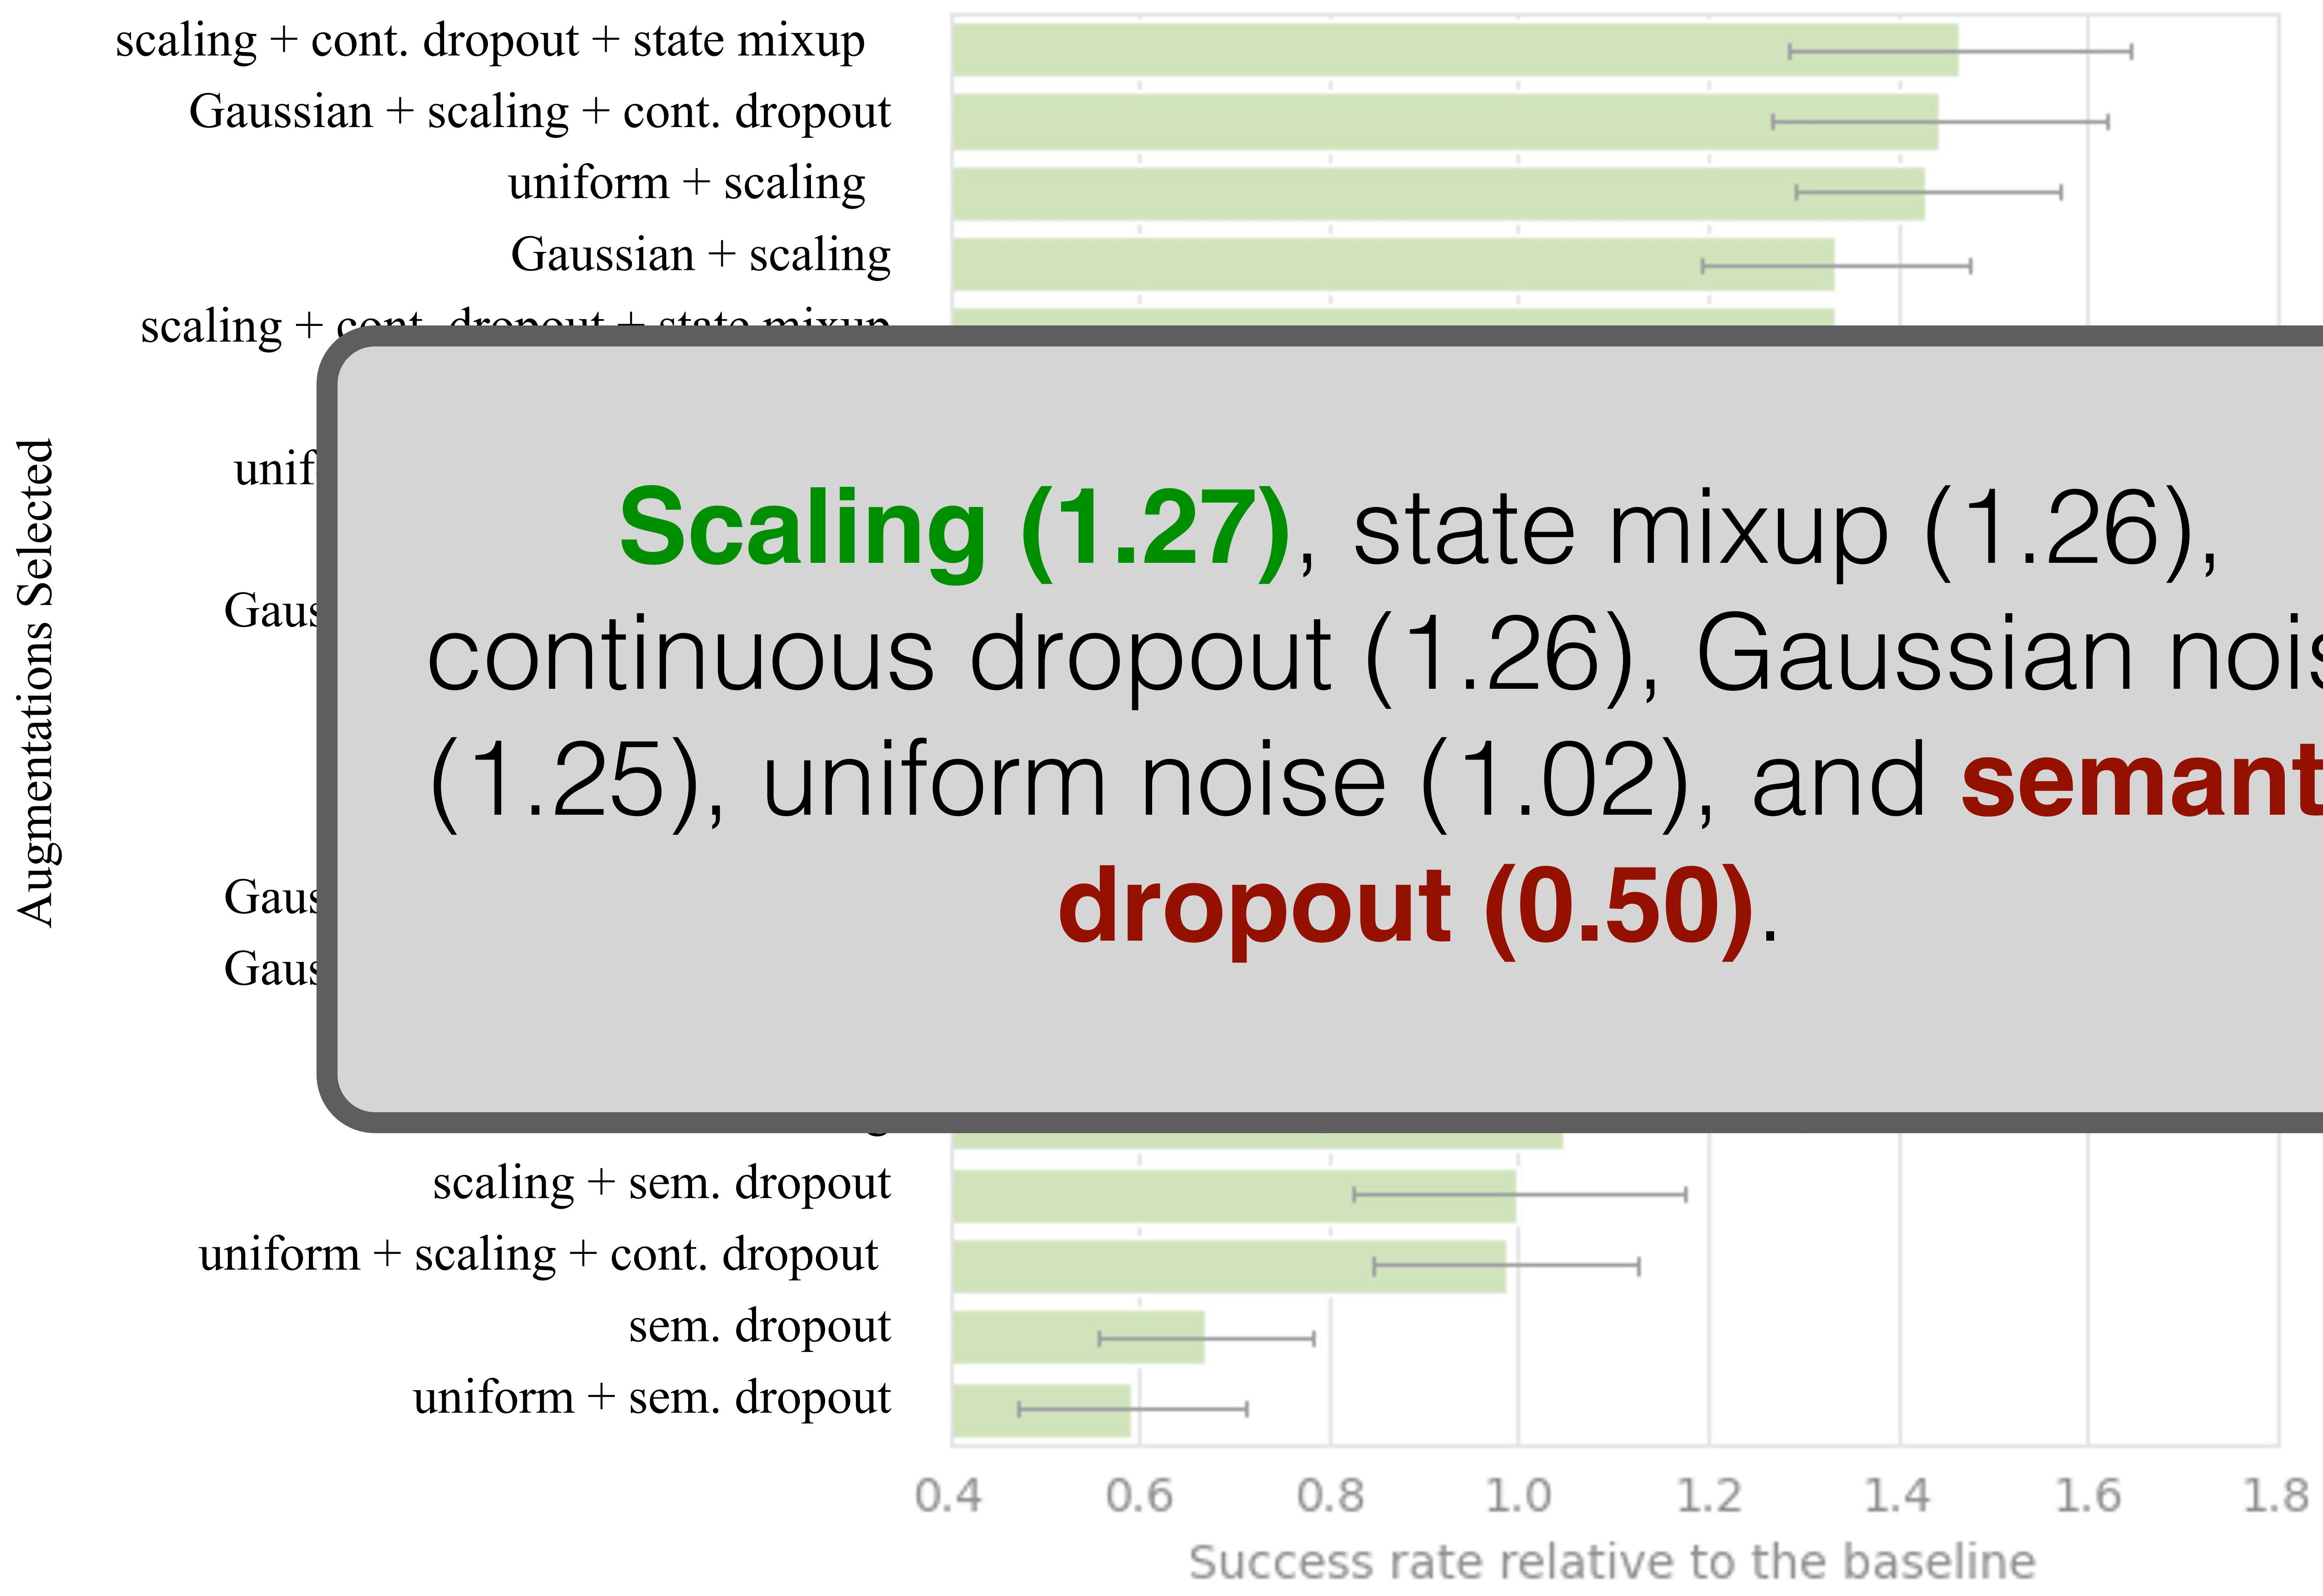

Scaling (1.27), state mixup (1.26), continuous dropout (1.26), Gaussian noise (1.25), uniform noise (1.02), and **semantic dropout (0.50)**.

Some Results - With Augmentations

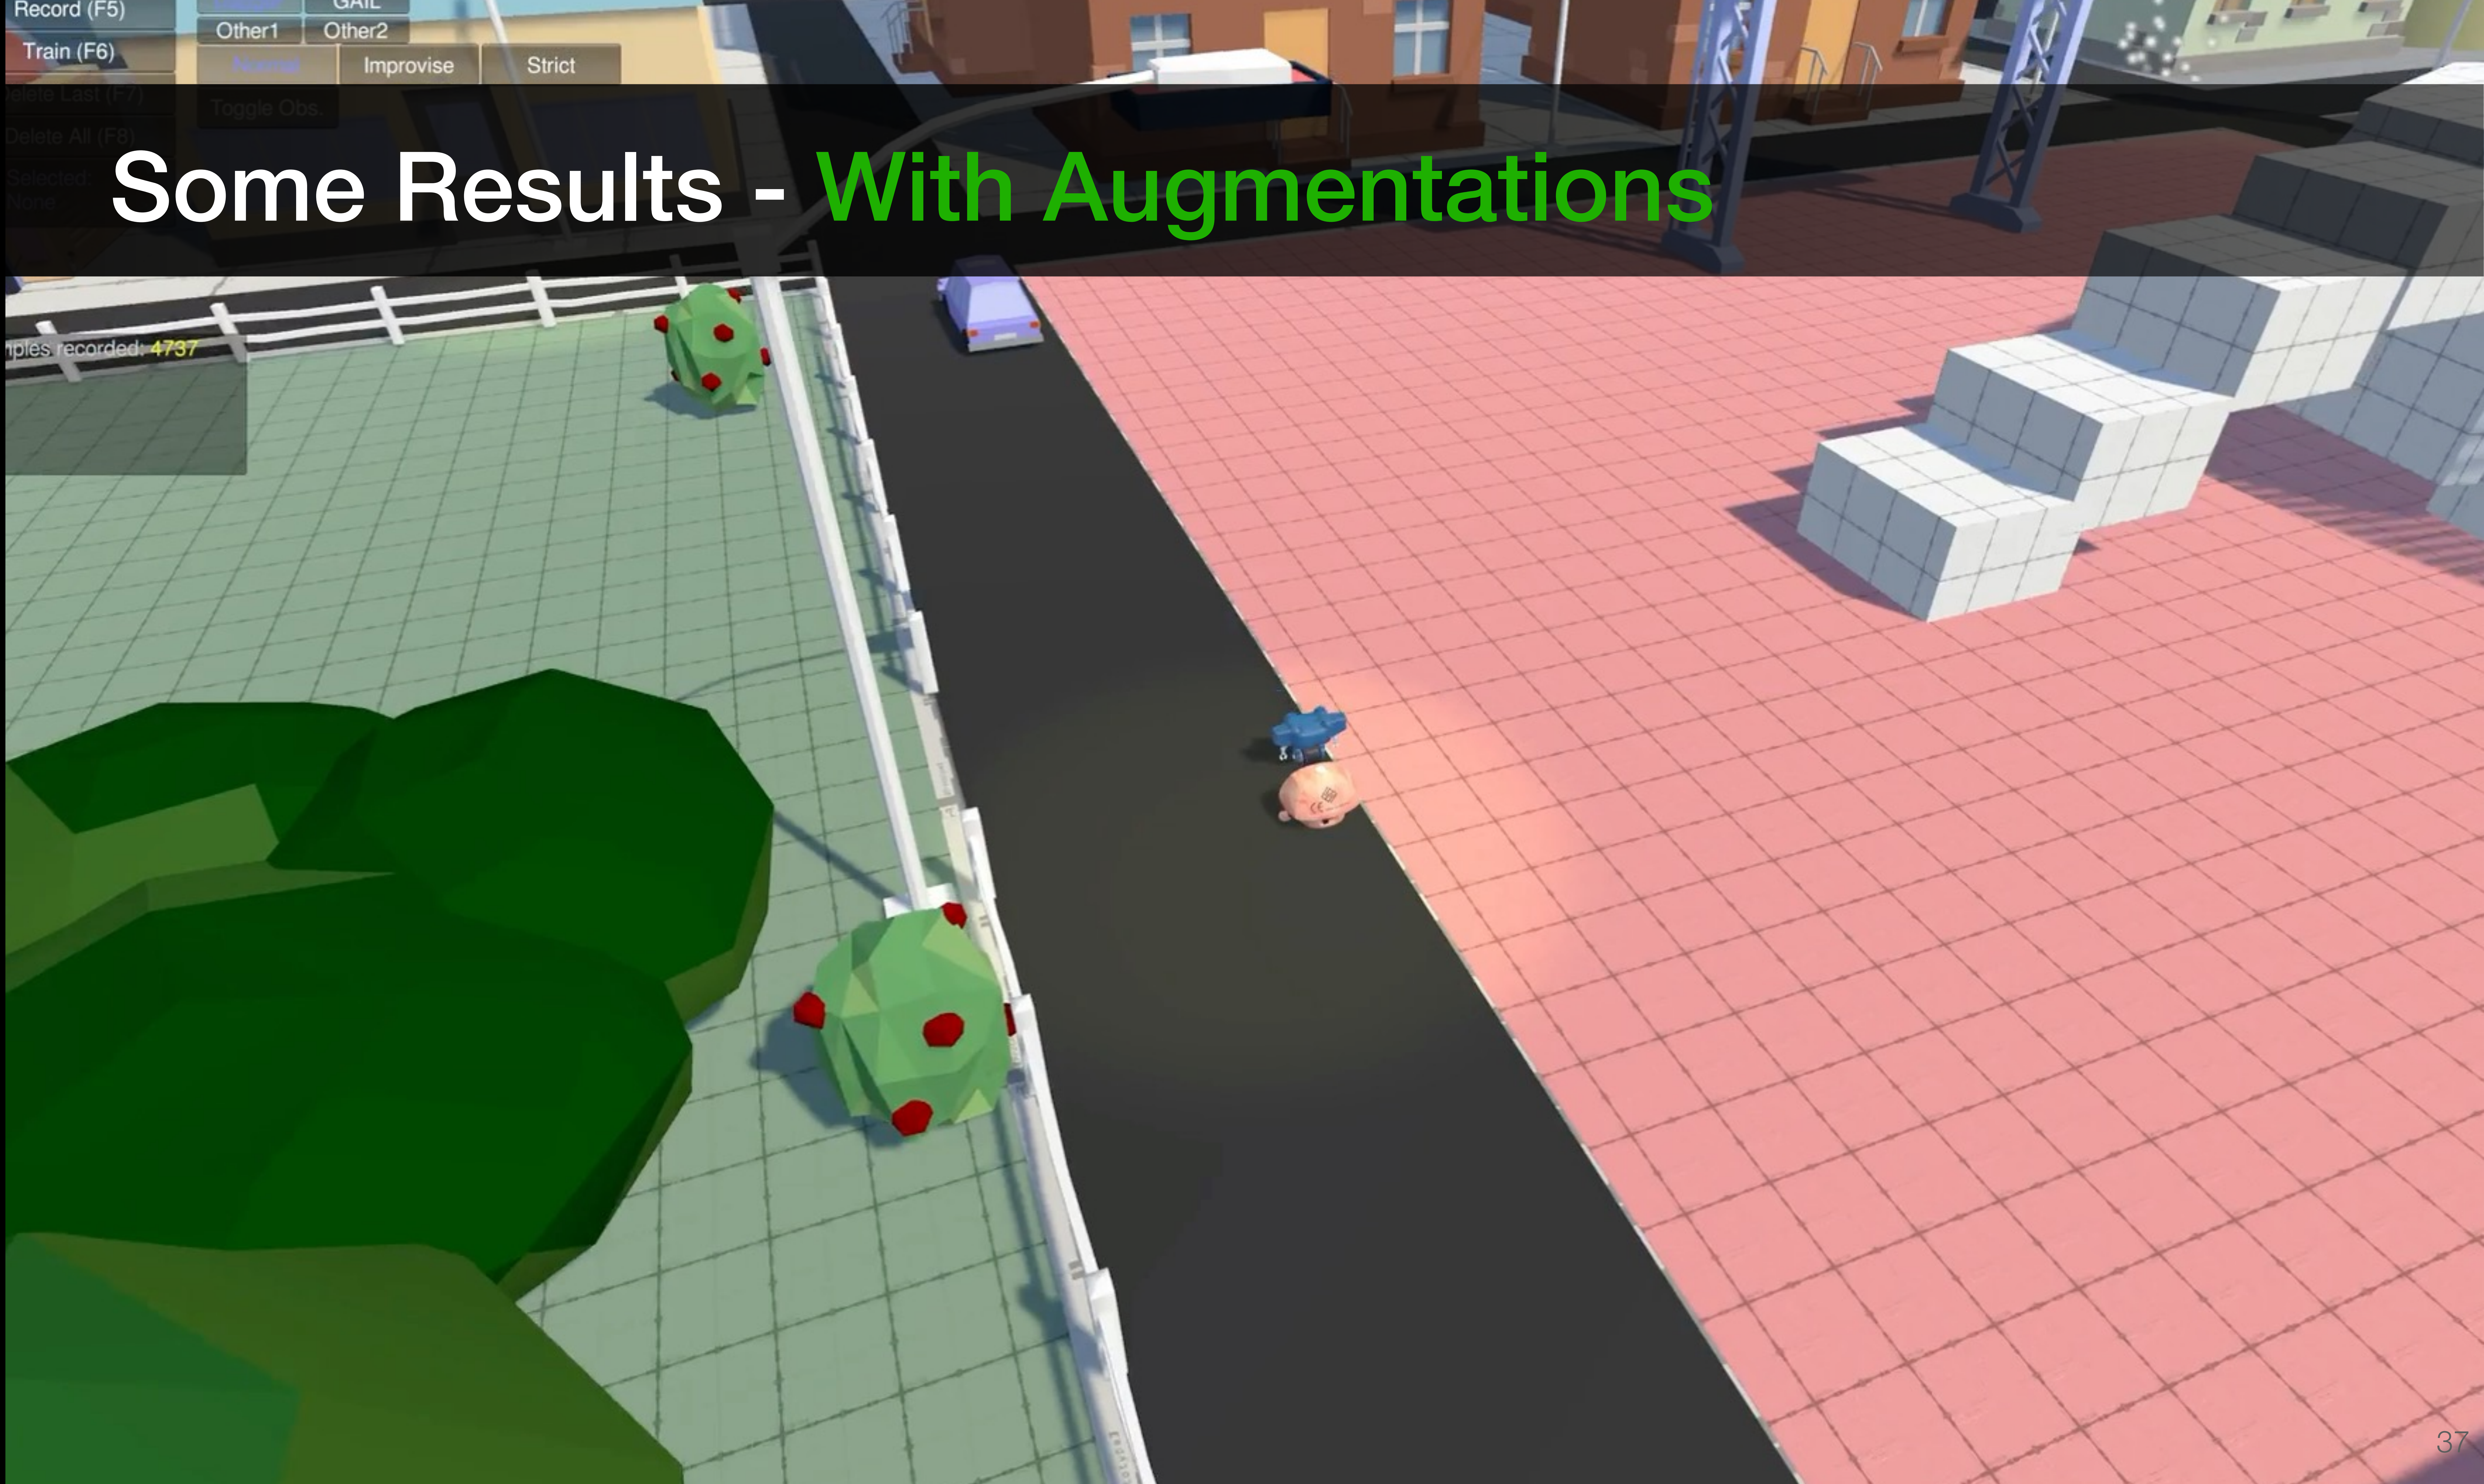

Some Results - Without Augmentations

Conclusion


GAIL
Other1
Other2
Normal
Improvise
Strict
Toggle Obs.


SIMPLE GAME-AI BEHAVIOUR

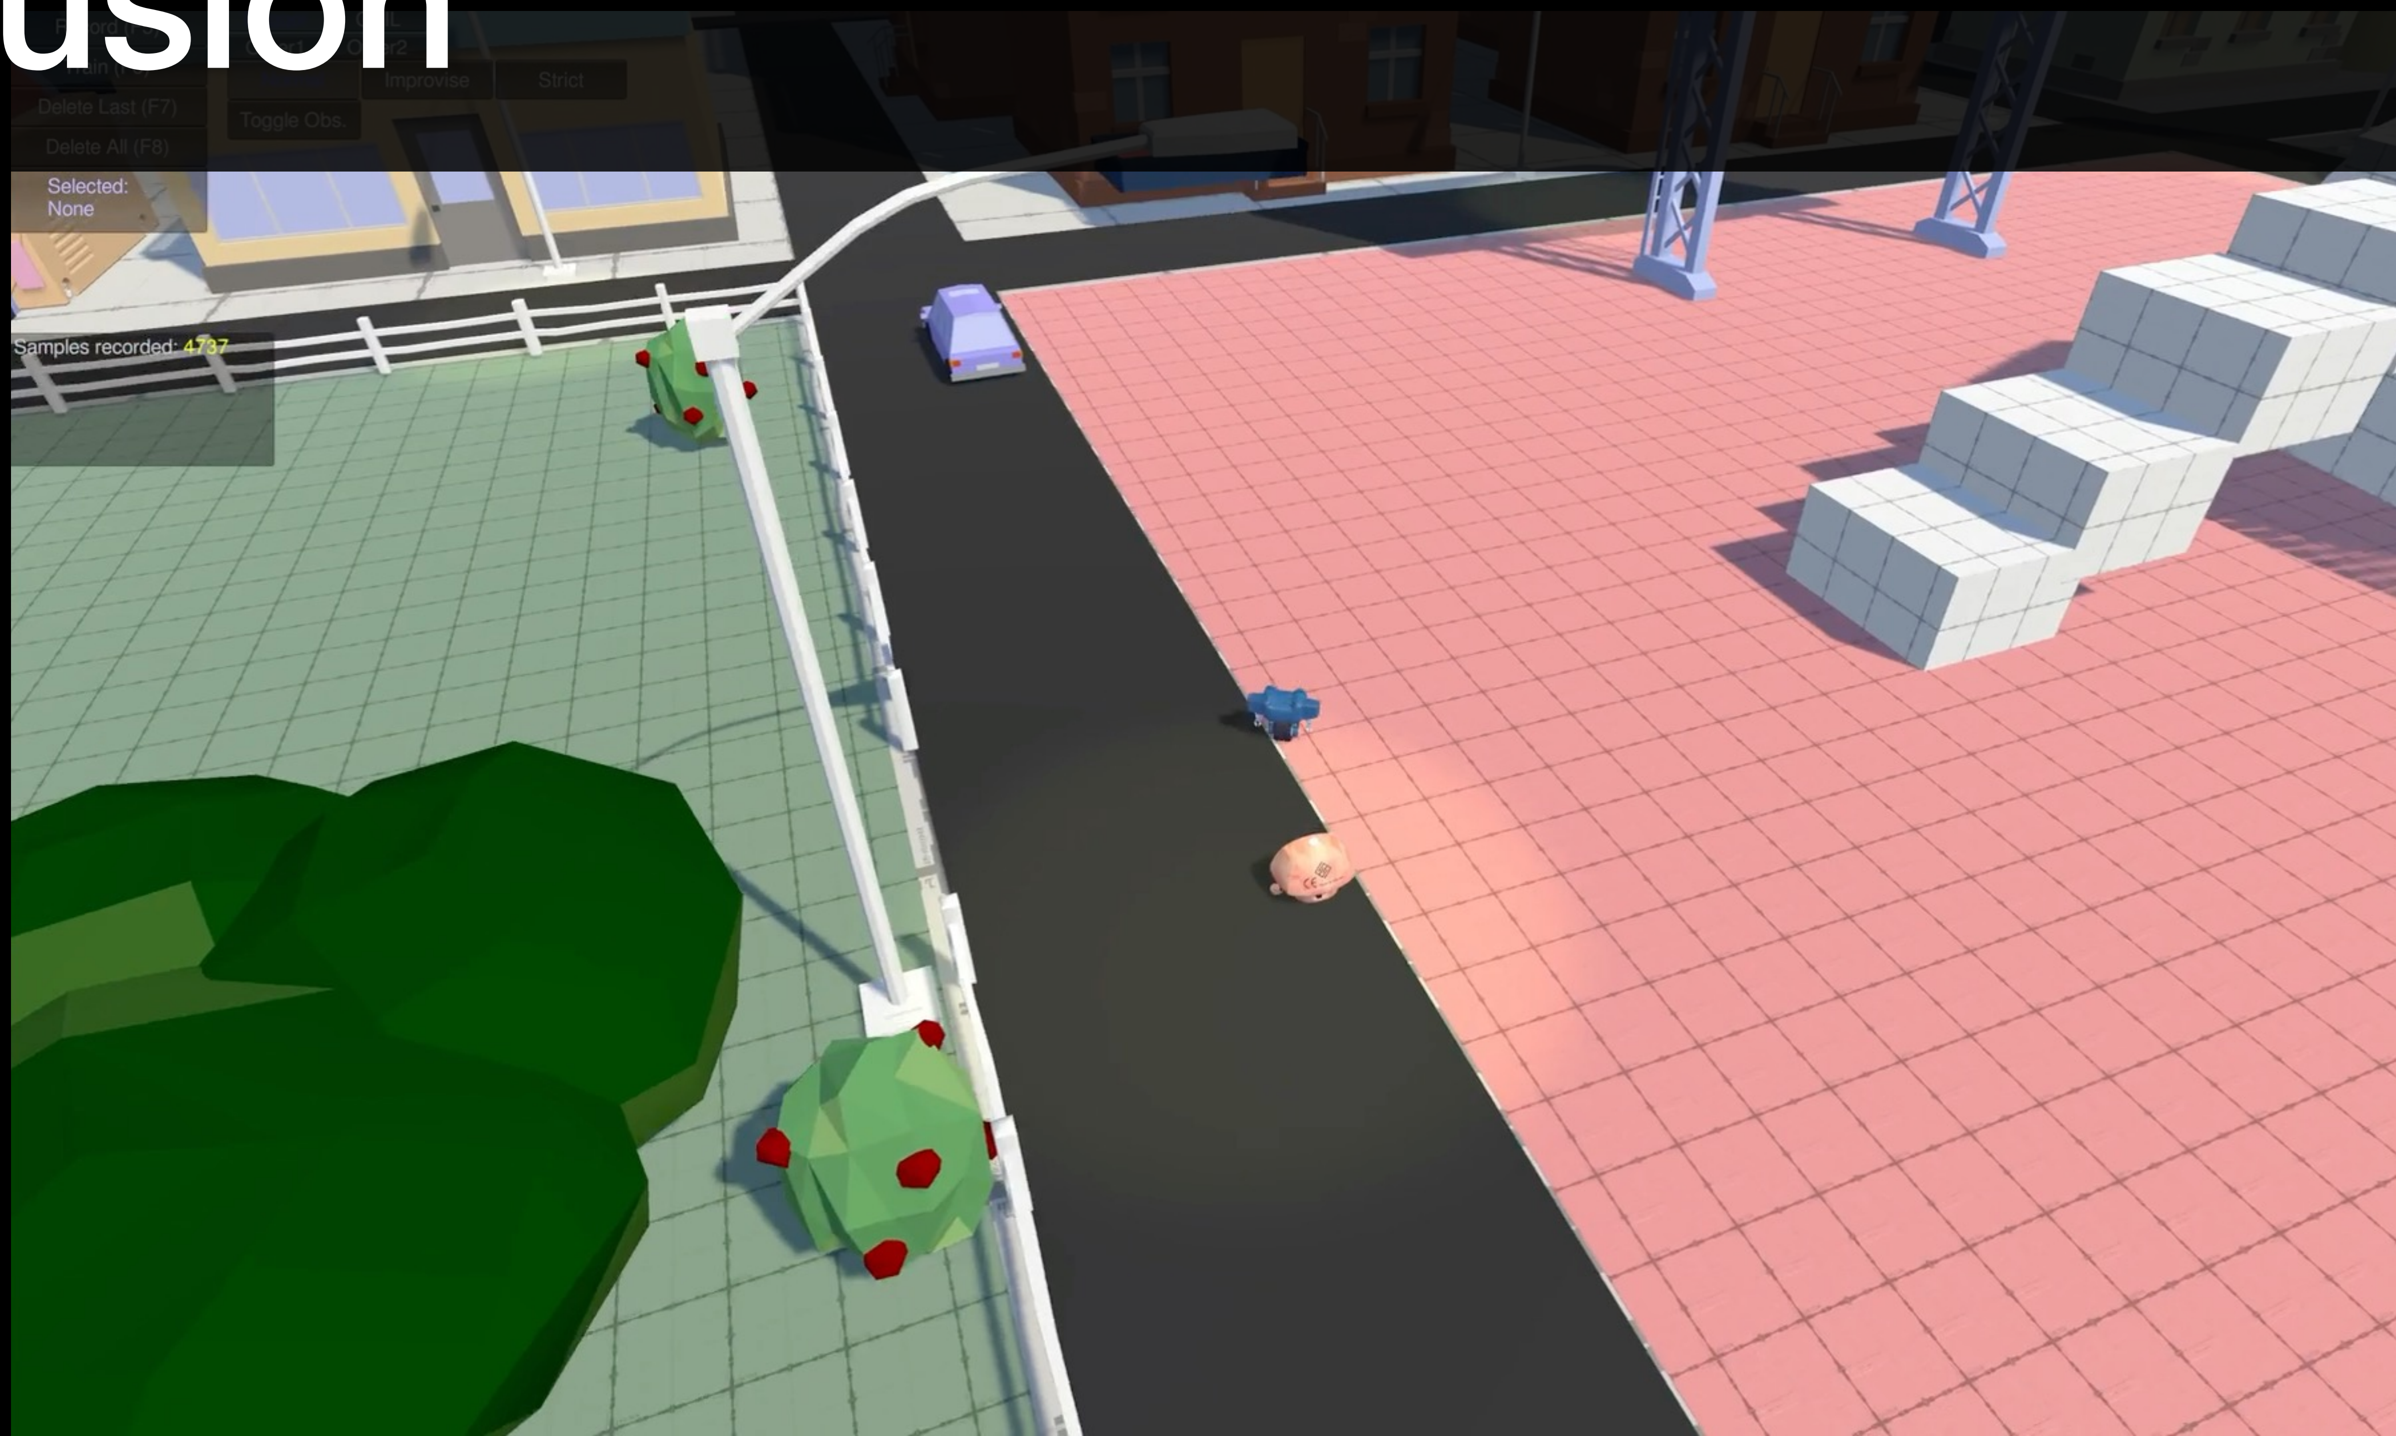

Contacts

email: asestini@ea.com

Electronic Arts: https://www.ea.com

Open Positions: ea.com/careers

SEED: seed.ea.com

We are  !

Role	Area of Interest	Location	Apply
Senior Computer Vision Researcher (Digital Humans) - SEED	Software Development	Toronto, Canada Vancouver, Canada	Apply Now
SEED Master Thesis Intern	Software Development	Stockholm, Sweden	Apply Now
Senior Research Scientist (Computer Vision) - SEED	Software Development	Toronto, Canada Vancouver, Canada	Apply Now
Senior Physics Software Engineer	Software Development	Toronto, Canada Vancouver, Canada Guildford, United Kingdom	Apply Now
Rendering Engineer - SEED	Software Development	Toronto, Canada Vancouver, Canada	Apply Now

References 📚

[1] *S4RL: Surprisingly Simple Self-Supervision for Offline Reinforcement Learning,* Sinha et al., 2021

[2] *Towards Informed Design and Validation Assistance in Computer Games Using Imitation Learning,* Sestini et al., 2023

[3] ***Improving Generalization in Game Agents with Data Augmentation in Imitation Learning***, Yadgaroff et al., 2024