

# PRIVILEGED INFORMATION DROPOUT IN REINFORCEMENT LEARNING

**Pierre-Alexandre Kamienny\***  
University of Oxford

**Kai Arulkumaran**  
Imperial College London

**Feryal Behbahani**  
Imperial College London

**Wendelin Böhmer**  
University of Oxford

**Shimon Whiteson**  
University of Oxford

## ABSTRACT

Using privileged information during training can improve the sample efficiency and performance of machine learning systems. This paradigm has been applied to reinforcement learning (RL), primarily in the form of distillation or auxiliary tasks, and less commonly in the form of augmenting the inputs of agents. In this work, we investigate Privileged Information Dropout (PI-Dropout) for achieving the latter which can be applied equally to value-based and policy-based RL algorithms. Within a simple partially-observed environment, we demonstrate that PI-Dropout outperforms alternatives for leveraging privileged information, including distillation and auxiliary tasks, and can successfully utilise different types of privileged information. Finally, we analyse its effect on the learned representations.

## 1 INTRODUCTION

The idea of incorporating additional information during training to accelerate the learning process has a long history, mainly in the context of supervised learning, and has been studied under different names, including learning using hints (Abu-Mostafa, 1990), hidden or privileged information (PI; Vapnik et al., 2009; Vapnik & Vashist, 2009), transfer learning (Pan & Yang, 2009), and multitask learning (Caruana, 1997). By leveraging the extra information during training as either extra inputs, extra outputs or additional loss terms, these methods add inductive biases to improve generalisation.

We focus on such a setting where privileged input features,  $\mathbf{x}^*$  are available during training, alongside the regular input features,  $\mathbf{x}$ . In particular, we consider three broad ways of leveraging PI: 1) *Distillation*: Using a “teacher” model, which receives PI as input, to regularise a separate “student” model, which only receives regular features as input (Hinton et al., 2015; Lopez-Paz et al., 2015). 2) *Auxillary tasks*: Constructing additional loss terms based on PI (Caruana & De Sa, 1997). 3) *Augmenting inputs*: Using PI as additional inputs.

Previous work on using PI in RL has focused largely on distillation (Parisotto et al., 2015; Rusu et al., 2015) and auxiliary tasks (Mirowski et al., 2016). The disadvantage of distillation is that it requires training a separate teacher model. In settings where the PI cannot be directly predicted from regular inputs, using PI instead as input would be more beneficial. For example, predicting depth from an RGB view of a scene has been shown to be a useful auxiliary task (Mirowski et al., 2016), but predicting non-overlapping views is in general not possible. However, multi-view integration—using these non-overlapping views as input—can be beneficial for tasks such as navigation.

Using PI as input in an actor-critic setting has been explored where only the critic has access to PI during training (Pinto et al., 2017). Other work has extended this principle to multi-agent systems, where information about the full state is used to train a centralised critic with decentralised actors (Lowe et al., 2017; Foerster et al., 2018; Rashid et al., 2018).

In contrast, we explore a more general way to provide PI as input to both policy and value functions through the use of PI-Dropout (Lambert et al., 2018), which extends information dropout

\*Correspondence to pierrealex1995@gmail.com

(Achille & Soatto, 2018) to leverage PI. Proposed in a supervised learning setting, PI-Dropout theoretically combines the representation learning advantages of information dropout with the ability to marginalise out  $\mathbf{x}^*$  during execution. Our contribution is to investigate the use of PI-Dropout in an RL setting and analyse its effect on trained agents. In our experimental setup we find that PI-Dropout outperforms other methods using PI, including auxiliary prediction tasks and distillation. In a key ablation, we show that PI-Dropout outperforms standard information dropout, indicating that our method does indeed benefit from using PI. Of note is parallel research from Salter et al. (2019), who also introduced a method that allows PI to be used as inputs for policies and value functions. Their method relies on aligning the attention of separate actor-critic networks, where one pair has access to PI, and is orthogonal and complementary to our work.

## 2 BACKGROUND AND METHOD

PI-Dropout (Lambert et al., 2018) is motivated by information bottleneck (IB) theory (Tishby et al., 2000). In IB, for inputs  $\mathbf{x}$  and task outputs  $\mathbf{y}$ , an “optimal” intermediate representation,  $\mathbf{z}$ , should have high mutual information (MI) with the outputs,  $I(\mathbf{z}; \mathbf{y})$ , while minimising task-irrelevant information from the inputs,  $I(\mathbf{x}; \mathbf{z})$ . The intuition is that the learned encoding  $\mathbf{z}$  will ignore uninformative patterns in  $\mathbf{x}$ , and therefore generalize well to previously unseen samples. It is possible to directly optimise for this property by using a relaxed Lagrangian formulation of the IB. This can be expressed as the minimisation of a cross-entropy term and an MI term,  $H(\mathbf{y}|\mathbf{z}) + \beta I(\mathbf{x}; \mathbf{z})$ , with hyperparameter  $\beta$ . Minimising the MI term corresponds to minimising the KL divergence between  $p(\mathbf{z}|\mathbf{x})$  and the marginal  $p(\mathbf{z})$ .

Information dropout (Achille & Soatto, 2018) uses multiplicative dropout to implement the IB Lagrangian, where the variance of the noise applied to  $\mathbf{z}$  is conditioned on  $\mathbf{x}$ . PI-Dropout instead conditions the noise on  $\mathbf{x}^*$ , allowing PI to mask out irrelevant features with noise. Concretely, in PI-Dropout,  $\mathbf{z} = h(\mathbf{x}; W) \odot \text{Lognormal}(\mathbf{0}, h^*(\mathbf{x}^*; W^*))$ , where  $h(\cdot; W)$  and  $h^*(\cdot; W^*)$  represent separate subnetworks<sup>1</sup>, parameterized by weights  $W$  and  $W^*$ , respectively. Achille & Soatto (2018) use a log-normal distribution for  $p(\mathbf{z})$ , which in PI-Dropout results in maximising the regularisation term  $\beta \log h^*(\mathbf{x}^*; W^*)$ . During execution, the PI subnetwork is ignored, and the variance is manually set to 0, marginalising out  $\mathbf{x}^*$ .

## 3 EXPERIMENTS

To compare different methods for leveraging PI under the RL setting, we designed a simple partially-observed environment (depicted in Figure 1), where  $\mathbf{x}$  is an observation of the environment, and  $\mathbf{x}^*$  is a privileged feature. We initially experimented with A3C (Mnih et al., 2016), PPO (Schulman et al., 2017) and DRQN<sup>2</sup> (Hausknecht & Stone, 2015), and use the latter as our baseline as it outperformed the other methods.



Figure 1: **Left**) Agent’s egocentric observation. **Middle**) Full state/default form of PI, with the goal highlighted in green for illustration purposes. **Right**) Alternative PI, consisting of the entirety of the current room, with the correct corridor (blue) given as a sub-goal.

### 3.1 ENVIRONMENT

Our environment is a gridworld consisting of three  $6 \times 6$  rooms linked together. The agent (red square), which is initialised in a random position in the left or centre rooms, must navigate to the

<sup>1</sup>In practice we restrict the variance  $\in (0, 1)$  using a sigmoid function after  $h^*(\mathbf{x}^*; W^*)$ .

<sup>2</sup>In addition to the original settings, we used Polyak updates for the target network and additional “burn-in” of the recurrent layer’s hidden state (Kapturowski et al., 2018) when training the network.

goal within 100 timesteps, receiving a reward of  $-0.1$  for each action and  $+10$  for reaching the goal. By default, we utilise the agent’s observation, a  $5 \times 5$  egocentric view, as  $\mathbf{x}$ , and the full state (FS) of the environment as  $\mathbf{x}^*$ . We also experiment with an alternative form of  $\mathbf{x}^*$ , in which the correct corridor to traverse—a sub-goal (SG)—is coloured blue.

### 3.2 AGENTS

The standard architecture for our DRQN-based agents consists of two convolutional layers, a GRU (Cho et al., 2014), and two fully-connected layers, with ReLU nonlinearities. In our PI-Dropout (PI-D) agent, we use two separate convolutional sub-networks for  $\mathbf{x}$  and  $\mathbf{x}^*$ , with multiplicative dropout applied to the output of the latter before the features of the subnetworks are concatenated (Figure 2).

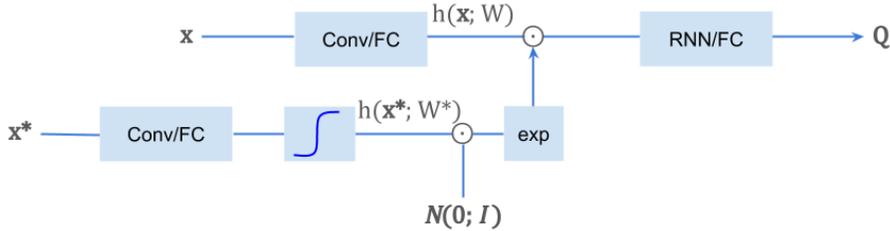


Figure 2: The architecture for a  $Q$ -learning-based PI-D agent. The regular features  $\mathbf{x}$  and privileged features  $\mathbf{x}^*$  are first processed by separate layers. During training, the intermediate representation of  $\mathbf{x}$  is multiplied with log-normal noise conditioned on  $\mathbf{x}^*$ , while during testing, no noise is added.

We also implemented several baselines:

- *DRQN*: Standard agent (no PI).
- *Auxiliary task (AUX)*: An agent with deconvolutional layers on top of the GRU, which is additionally trained to reconstruct  $\mathbf{x}^*$  at every timestep.
- *Distillation (DIS)*: A student agent that only receives  $\mathbf{x}$  as input, but has an additional loss to mimic the outputs of a pre-trained teacher network that receives  $\mathbf{x}^*$  as input.
- *Naive Dropout (ND)*: An agent with the PI-D architecture, where the multiplicative dropout noise is manually annealed from 0 to 1 over 3000 training episodes.
- *Information Dropout (I-D)*: An agent that uses information dropout (no PI).

All models are trained for  $10^4$  episodes, using  $\epsilon$ -greedy exploration. Test performance is calculated as the reward of the greedy policy, averaged over episodes starting from all possible initial positions (in the left and centre rooms). In the results we suffix all methods with  $[\mathbf{x}, \mathbf{x}^*]$ . For example, PI-D[5x5,FS] indicates a PI-Dropout agent that uses the egocentric view as  $\mathbf{x}$  and full state as  $\mathbf{x}^*$ .

### 3.3 RESULTS

The results for all the PI-based algorithms, using their best hyperparameters, are shown in the Figure 3a. The performance of the DRQN agent with full state view as input (DRQN[FS, $\emptyset$ ]) is plotted as an *oracle* upper-bound. We also show the DRQN[5x5, $\emptyset$ ] as the simplest baseline. We see that our PI-Dropout agent outperforms all methods, getting closest to the performance of the oracle agent. Under detailed inspection, we observed that only our PI-Dropout agent and the oracle agent were able to consistently solve the task when initialised from any position in the environment. Other baselines fail at solving the task from the leftmost room which is the farthest from the goal, and hence show sub-optimal performance. The Naive Dropout agent (ND[5x5,FS]) in particular has poor performance across different initial positions and random seeds, highlighting the benefit of the principled PI-Dropout approach.

Figure 3b shows that PI-Dropout performs equally well, and even converges faster, with subgoals provided as PI (PI-D[5x5,SG]), demonstrating its generality. We also include an I-D[5x5,5x5] agent, which corresponds to standard information dropout (no PI). Given that its performance is the same as

the DRQN[5x5, $\emptyset$ ] agent, this indicates that our agents benefit from PI, and not simply regularisation from dropout. In Supplementary Figure 4, we show the sensitivity of PI-Dropout to the value of  $\beta$ . With FS as PI, PI-Dropout is quite sensitive, but less so with SG as PI.

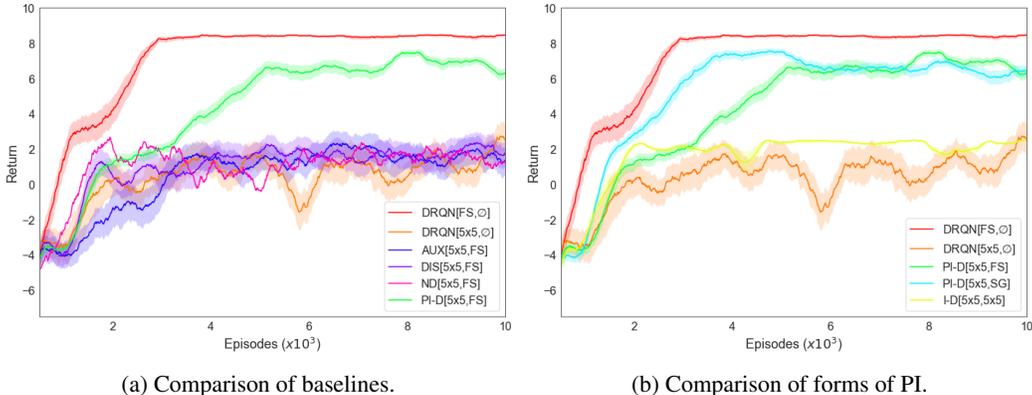


Figure 3: Average test performance of agents  $\pm 1$  standard error, calculated over 5 random seeds.  $[\cdot, \cdot]$  indicates  $\mathbf{x}$  and  $\mathbf{x}^*$ , respectively. Our method, with either full state PI (PI-D[5x5,FS]; green curve) or sub-goal PI (PI-D[5x5,SG]; cyan curve), surpasses the baselines and achieves the closest performance to that of the *oracle* (DRQN[FS, $\emptyset$ ]; red curve).

### 3.4 ANALYSIS

We set out to understand the effect of PI-Dropout on our agent’s learned representations, with the hypothesis that the information of the agent’s position might be more easily decodable from the PI-Dropout agent, in comparison to baselines. To test this, we generated a dataset of GRU activations (which implicitly form a belief state in partially-observable environments) by collecting these from multiple rollouts for each agent from all initial positions. We then trained linear classifiers to predict the position of the agent based on these activations. Confusion matrices for the results are shown in Supplementary Figure 5.

The classifier for the auxiliary task agent (AUX[5x5, $\emptyset$ ]) had the highest accuracy of 94.8%, which is perhaps expected given that it has the explicit auxiliary task of reconstructing the full state—and hence the agent’s location. Surprisingly, the classifier for the PI-Dropout agent (PI-D[5x5,FS]) did not have a significantly higher accuracy than that of the baseline DRQN (DRQN[5x5, $\emptyset$ ]) classifier (88.8% vs. at 90.1%). However, that solving the task successfully may not necessitate linearly decodable representations of the agent’s position, and further analysis is needed to better understand the role of PI-Dropout in our trained agents. We are also interested in understanding how PI-Dropout might improve exploration, and how it compares with or can be combined with other methods that use noise to aid exploration (Fortunato et al., 2017; Plappert et al., 2017).

## 4 DISCUSSION

In this work, we investigated the use of PI-Dropout in the context of reinforcement learning, which enables augmenting the inputs of any RL algorithms with privileged information. In a simple partially-observable environment, we demonstrated improved performance using PI-Dropout, in comparison with other methods that use PI, including distillation and auxiliary prediction tasks. We further showed the generality of PI-Dropout by utilising two different types of PI—the full state of the environment and sub-goals. An ablation against standard information dropout confirms that the use of PI is indeed beneficial and responsible for improved performance. An area for future research is to apply PI-Dropout in more challenging domains and to better understand the contributions of PI-Dropout in learning better representations or improving exploration.

## REFERENCES

- Yaser S Abu-Mostafa. Learning from hints in neural networks. *Journal of complexity*, 6(2):192–198, 1990.
- Alessandro Achille and Stefano Soatto. Information dropout: Learning optimal representations through noisy computation. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2897–2905, 2018.
- Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- Rich Caruana and Virginia R De Sa. Promoting poor features to supervisors: Some inputs work better as outputs. In *Advances in Neural Information Processing Systems*, pp. 389–395, 1997.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, 2014.
- Jakob N Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, et al. Noisy networks for exploration. *arXiv preprint arXiv:1706.10295*, 2017.
- Matthew Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. In *2015 AAAI Fall Symposium Series*, 2015.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Steven Kapturowski, Georg Ostrovski, John Quan, Remi Munos, and Will Dabney. Recurrent experience replay in distributed reinforcement learning. In *International Conference on Learning Representations*, 2018.
- John Lambert, Ozan Sener, and Silvio Savarese. Deep learning under privileged information using heteroscedastic dropout. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8886–8895, 2018.
- David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. *arXiv preprint arXiv:1511.03643*, 2015.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in neural information processing systems*, pp. 6379–6390, 2017.
- Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, et al. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*, 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*, pp. 1928–1937, 2016.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. Actor-mimic: Deep multitask and transfer reinforcement learning. *arXiv preprint arXiv:1511.06342*, 2015.
- Lerrel Pinto, Marcin Andrychowicz, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. Asymmetric actor critic for image-based robot learning. *arXiv preprint arXiv:1710.06542*, 2017.

Matthias Plappert, Rein Houthoofd, Prafulla Dhariwal, Szymon Sidor, Richard Y Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*, 2017.

Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 4295–4304, 2018.

Andrei A Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. Policy distillation. *arXiv preprint arXiv:1511.06295*, 2015.

Sasha Salter, Dushyant Rao, Markus Wulfmeier, Raia Hadsell, and Ingmar Posner. Attention privileged reinforcement learning for domain transfer. *arXiv preprint arXiv:1911.08363*, 2019.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural networks*, 22(5-6):544–557, 2009.

Vladimir Vapnik, Akshay Vashist, and Natalya Pavlovitch. Learning using hidden information (learning with teacher). In *2009 International Joint Conference on Neural Networks*, pp. 3188–3195. IEEE, 2009.

SUPPLEMENTARY

RESULTS

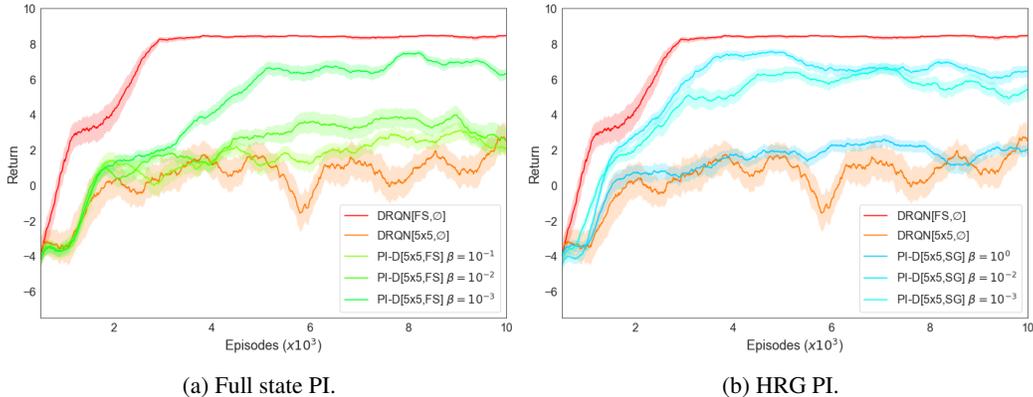


Figure 4: Average test performance of agents  $\pm 1$  standard error, calculated over 5 random seeds.  $[\cdot, \cdot]$  indicates  $\mathbf{x}$  and  $\mathbf{x}^*$ , respectively.

ANALYSIS

Supplementary Figure 5) shows confusion matrices for the linear classifiers trained on belief states from three different agents: DRQN[5x5,∅], AUX[5x5,∅], and PI-D[5x5,FS]. For the PI-Dropout agent, we include results from both its operation during test-time (does not receive  $\mathbf{x}^*$  as input) and train-time (receives  $\mathbf{x}^*$  as input). Intriguingly, during standard train-time operation the PI-Dropout agent has the highest uncertainty over the agent’s location, as compared to all other settings, but it always correctly identifies the correct room, unlike in all other settings.

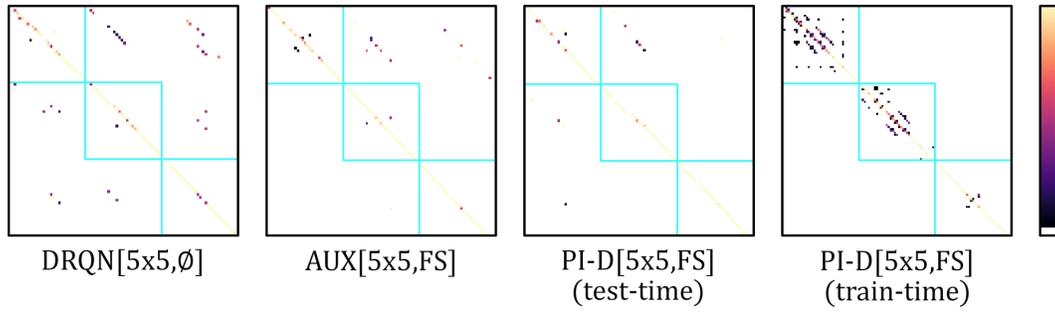


Figure 5: Real location vs. predicted location, from small (black) to large (yellow) probability mass; white indicates zero mass (to highlight low probability predictions). Room boundaries are shown in cyan. All agents without PI make errors outside of the correct room, but the PI-Dropout agent given  $x^*$  only makes localisation errors within the correct room.