

463 **A Implementation Details**

464 **A.1 Distance function**

465 As mentioned in Section 3, we wish to learn an encoding such that our distance metric  $d$  is the  
 466 euclidean distance between the encoded states.

$$d(s, s') = \|h(s) - h(s')\|^2 \tag{5}$$

467 To learn the encoder  $h$ , we optimize a contrastive loss on encodings of the current and future states  
 468 along the same trajectory. We use the InfoNCE Loss [30],

$$\mathcal{L}_q = \log \frac{\exp(q^T W k)}{\exp\left(\sum_{i=0}^K \exp(q^T W k_i)\right)} \tag{6}$$

469 with query  $q = h(s_t^i)$  as the encoded starting state, and the keys  $k = h(s_{t+H}^i)$  as the encoded future  
 470 states along the  $K$  trajectories in the dataset  $\mathcal{D}$ .

471 **A.2 Training**

472 The training for both skill extraction and fine-tuning were done on a single NVIDIA 2080Ti GPU.  
 473 Skill extraction takes approximately 3-4 hours, and fine-tuning requires less than 10 minutes. We  
 474 build upon the SPiRL codebase at <https://github.com/clvrail/spirl>. Hyperparameters used  
 475 are listed in Table 6.

Table 6: Training Hyperparameters

Hyperparameter	Value
<b>Contrastive Distance Metric</b>	
Encoder output dim	32
Encoder Hidden Layers	128
Encoder # Hidden Layers	2
Optimizer	Adam( $\beta_1 = 0.9, \beta_2 = 0.999, \text{LR}=1\text{e-}3$ )
<b>Skill extraction</b>	
Epochs	200
Batch size	128
Optimizer	Adam( $\beta_1 = 0.9, \beta_2 = 0.999, \text{LR}=1\text{e-}3$ )
$H$ (sub-trajectory length)	10
$\beta$	$5\text{e-}4$ (Kitchen), $1\text{e-}2$ (Maze)
<b>Skill Encoder</b>	
dim- $\mathcal{Z}$ in VAE	128
hidden dim	128
# LSTM Layers	1
<b>Skill Decoder</b>	
hidden dim	128
# hidden layers	5
<b>Inverse Skill Dynamic Model</b>	
hidden dim	128
# hidden layers	5
<b>Fine-tuning</b>	
Epochs	50
Batch size	128
Optimizer	Adam( $\beta_1 = 0.9, \beta_2 = 0.999, \text{LR}=1\text{e-}3$ )

### 476 **A.3 Datasets**

477 The PointMaze and Kitchen environment datasets (both skill extraction datasets and few-shot learning  
478 datasets) are generated from an expert policy. For the AntMaze environment, the dataset was created  
479 from the D4RL dataset [29], licensed under the Creative Commons Attribution 4.0 License (CC BY).  
480 Datasets for each blocked section was created by filtering out any trajectories that passed through the  
481 blocked regions shown in Figure 3. Code for the dataset generation is included in the supplementary  
482 material.