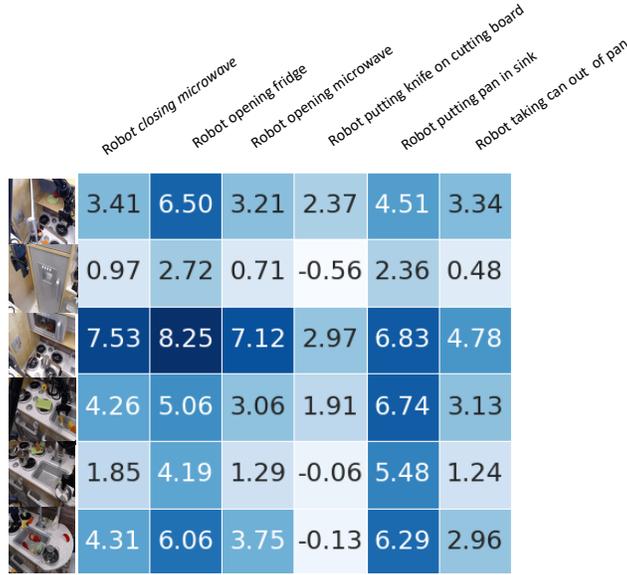


488 **Supplemental Material**

489 **A Confusion Matrix Analysis**

490 To test out the alignment of the VLM with other domains we also perform our confusion matrix  
 491 analysis on Berkeley Bridge Dataset [Ebert et al., 2021] and 2 human video datasets — one from a  
 492 stock gif repository and another a dataset collected from the kitchen in our laboratory. The Bridge  
 493 dataset videos consist of a robot teleoperated to interact with toy versions of everyday kitchen objects  
 494 such as forks, knives, pans, etc. The stock photo set consists of random videos from disparate  
 495 domains. The lab dataset is collected by the same actor and with the same camera angles, lighting,  
 496 etc., with the goal of this small dataset to study whether removing confounders from the gifs like  
 lighting and camera angles makes the task of video retrieval easier or not.



**Figure 9: Domain Alignment of the Berkeley Bridge Dataset.** We study how well videos taken from the Berkeley Bridge Dataset [Ebert et al., 2021] are correctly identified by the S3D [Xie et al., 2018] pretrained on HowTo100M [Miech et al., 2019]. Each entry  $(i, j)$  corresponds to the similarity score between the  $i^{th}$  video and the  $j^{th}$  text label. We find that the S3D model often gets distracted by background objects and mislabels the videos in this context. For example, for the last video of *robot taking can out of pan*, the model sees the sink and labels the video with a high score for the text label *robot putting pan in sink*.

497

498 As shown in Figure 9, the alignment of the Berkeley Bridge Dataset is poor. This is potentially due to  
 499 the large amount of clutter and distractor objects in the background of the videos. This often leads  
 500 to confusion during inference in model as it often correctly identifies the background objects and  
 501 fails to attend to the task object. For example, for the last video of *robot taking can out of pan*, the  
 502 model sees the sink and labels the video with a high score for the text label *robot putting pan in*  
 503 *sink*. The sink is a prominent distractor object resulting in high scores across multiple videos (see  
 504 vertical column corresponding to the label *Robot putting pan in sink*). Another mode of failure is  
 505 typically in the confusion of object identities. This is especially prominent in the first 3 videos where  
 506 the microwave and the fridge are repeatedly mixed up.

507 We subsequently study whether videos that are visually similar to the domain of HowTo100M [Miech  
 508 et al., 2019] are correctly labeled by S3D. We find that videos are generally well aligned with all the  
 509 videos we tested to be correctly labeled as seen in Figure 11. The model transfers to these video  
 510 domains demonstrating appreciable generalizability including correctly labeling videos of inanimate  
 511 objects like the video of a waterfall. The inference capabilities also generalize to a variety of camera  
 512 views with the videos from a first-person camera. For example, *Human turning knob* and *Human*  
 513 *opening microwave* are correctly classified in addition to the third-person video of *Human opening*  
 514 *cabinet*.

	Human opening cabinet	Human opening microwave	Human turning knob	waterfall
	1.71	-2.04	1.45	0.40
	4.02	6.85	1.34	2.19
	1.54	2.35	3.63	0.12
	0.66	-0.09	2.13	3.57

**Figure 10: Domain Alignment of the Stock Gif Repository Dataset.** We study how well videos taken from online Stock Gif Repositories are correctly identified by the S3D [Xie et al., 2018] pretrained on HowTo100M [Miech et al., 2019]. Each entry  $(i, j)$  corresponds to the similarity score between the  $i^{th}$  video and the  $j^{th}$  text label. We find that the model aligns well with videos that are visually similar to the domain on which it was trained, i.e., HowTo100M [Miech et al., 2019] indicating appreciable generalization capabilities. Each of the motions are correctly identified along with the objects, including videos of inanimate objects like the waterfall in the last video.

515 Finally, we study whether videos sourced from the same domain, i.e., videos without any confounders  
 516 such as camera view, actors, etc., can be correctly labeled by the pretrained VLM. To this end, we  
 517 use the dataset collected in our lab with the same actor, camera view, lighting etc., and perform the  
 518 confusion matrix analysis. We find that this is generally successful as the model can generally identify  
 519 objects well and confuses one type of door for another. It does not, however, possess a common  
 520 sense understanding of the world as it cannot tell that turning a light on is achieved by the motion of  
 521 flipping a switch.

## 522 B Implementation Details

### 523 B.1 Policy Learning

524 We use PPO [Schulman et al., 2017] to train our agents. However, to ensure slow yet stable  
 525 convergence, we set the batch size for updates equal to the number of parallel environments times the  
 526 number of steps per environment. This effectively ensures that the entire replay buffer is loaded at  
 527 once to perform updates. We also set the number of epochs to 1. This ensures that the policy does  
 528 stray from the trust region [Schulman et al., 2017] resulting in a slower, yet more stable convergence.  
 529 This of course means that the training can be sped up by playing around with these parameters, but  
 530 we find these settings to work in all the Metaworld experiments as well as Kitchen experiments.

531 We also find that increasing the weight for the entropy term in the loss for PPO provides significant  
 532 improvements. We hypothesize this to be the case since the learned reward function contains two  
 533 kinds of maxima — one which corresponds to the actual task defined by the descriptor and the other  
 534 a spurious maximum corresponding to an adversarial video which, to the model, "looks" like the  
 535 task descriptor but does not actually solve the task defined by the descriptor. This is expected as  
 536 the VLM is never finetuned on environment-specific data and thus it is unreasonable to expect that  
 537 the RoboCLIP reward will perfectly align with the downstream task reward. We find that setting a  
 538 high coefficient for entropy term in the loss function of PPO [Schulman et al., 2017] improves the  
 539 exploration bonus provided to the policy so that it explores multiple maxima in the RoboCLIP reward  
 540 allowing the agent to solve the task by eventually optimizing the RoboCLIP maxima corresponding  
 541 to task success.

542 The policy is implemented using StableBaselines3 [Raffin et al., 2021] with the following parameters:

	Human opening door	Human opening microwave	Human opening slide door	Human turning light on
Human opening door	3.43	0.67	4.59	0.58
Human opening microwave	3.57	6.88	4.78	0.31
Human opening slide door	6.89	5.15	7.07	3.92
Human turning light on	4.64	2.85	5.33	2.39

**Figure 11: Domain Alignment of the Lab Dataset.** We study how well videos collected from a first-person view in our lab kitchen are correctly identified by the S3D [Xie et al., 2018] pretrained on HowTo100M [Miech et al., 2019]. Each entry  $(i, j)$  corresponds to the similarity score between the  $i^{th}$  video and the  $j^{th}$  text label. We find that the model aligns well with videos that are visually similar to the domain on which it was trained, i.e., HowTo100M [Miech et al., 2019] indicating appreciable generalization capabilities. While the diagonal matrix is not diagonal-heavy, the model is generally able to identify door-like objects and is able to correctly identify a microwave. However, it is unable to correctly identify that a light switch turns a light on and lacks commonsense reasoning capabilities that to turn a light on, one needs to flip a switch.

```

543     • learning_rate=0.0003
544     • n_steps=128
545     • batch_size=2048
546     • n_epochs=1
547     • gamma=0.99
548     • gae_lambda=0.95
549     • clip_range=0.2
550     • clip_range_vf=None
551     • normalize_advantage=True
552     • entropy_coef=0.5
553     • value_function_coef=0.5
554     • max_grad_norm=0.5
555     • use_sde=False
556     • sde_sample_freq=-1
557     • target_kl_divergence=None
558     • stats_window_size=100
559     • seed=None
560     • device='auto'
561     • _init_setup_model=True

```

562 The policy network consists of a fully connected neural network with 2 layers and 64 hidden units per  
563 layer. The feature extractor for the  $Q$  function and the policy are shared. The remaining parameters  
564 are listed below:

```

565     • activation_function=tanh
566     • ortho_init=True
567     • use_sde=False
568     • log_std_init=0.0
569     • full_std=True
570     • use_expln=False
571     • squash_output=False

```

- 572 • features\_extractor\_class=Flatten
- 573 • features\_extractor\_kwargs=None
- 574 • share\_features\_extractor=True
- 575 • normalize\_images=True
- 576 • optimizer\_class=Adam
- 577 • optimizer\_kwargs=None

## 578 **B.2 Reward Generation**

579 Here, we describe the implementation details for generating the RoboCLIP conditioned on various  
580 demonstration types. We find that in Metaworld, normalizing the video of the demonstration is not  
581 needed. Typically, images are stored as 3D arrays with width, height, and channels, and with a  
582 per-pixel value of 0 – 255. S3D [Xie et al., 2018] pretrained on HowTo100M [Miech et al., 2019]  
583 was trained on per pixel values normalized to 0 – 1. Experimentally, we found that not normalizing  
584 the images for the Metaworld experiments yielded higher zero-shot rewards. We also found this to  
585 occur when training on Kitchen videos. Not normalizing yielded benefits in terms of higher zero-shot  
586 rewards. We also shorten the episode videos from 128 timesteps per episode to a 32 frame video  
587 using uniform downsampling. This ensures that we conform to the S3D pretraining methodology.