

9. Supplementary

9.1. Hyperparameter Details

PPO Parameters	
Parallel simulation environments	6
Rollout length (steps per environment)	64
DDPPO sync fraction	0.6
Number of PPO Epochs	4
Mini-batches per epoch	1
Optimizer	Adam
Learning rate	2.5×10^{-4}
Epsilon (ϵ)	1.0×10^{-5}
Learning rate decay	False
PPO-clip	0.2
clip decay	False
Clip the value loss	True
Generalized advantage estimation (GAE)	True
Normalized	True
γ	0.99
τ	0.95
Value loss coefficient (c_v)	0.5
Offset regularization coefficient (c_r)	0.1146
Max gradient norm	0.2
Reward Parameters	
Success (r_{success})	10
Success distance	1.0m
slack reward (r_{slack}) scalar	-0.01

Table 6: Hyperparameters shared by all experiments.

9.2. Semantic Segmentation Predictor Details

To finetune the RedNet [16] semantic segmentation module we render 60k images (50k in train and 10k in val) from randomly sampled view points in the Matterport3D houses [7] with YCB inserted objects. RedNet is finetuned under the Cross-Entropy Loss function using the SGD optimizer with learning rates of 0.0004 for the lower layers and 0.01 for the last layer. We use the same data augmentation techniques as proposed in [16] (random scale, crop, flip and color jitter). The category level metrics are shown in Tab. 7. We observe that the segmentor overfits on the training set with performance metrics lower on the validation set for the Goal-Object categories. The same YCB objects are inserted in both the training and validation split, thus leading to similar performances on both splits. Fig. 8 shows some qualitative examples of the semantic segmentation outputs. The left column shows the RGB input frame (the depth input is not shown in this figure for simplicity). The two middle columns show the semantic segmentation ground truth and prediction. The right most column shows the GoalSeg probabilities.

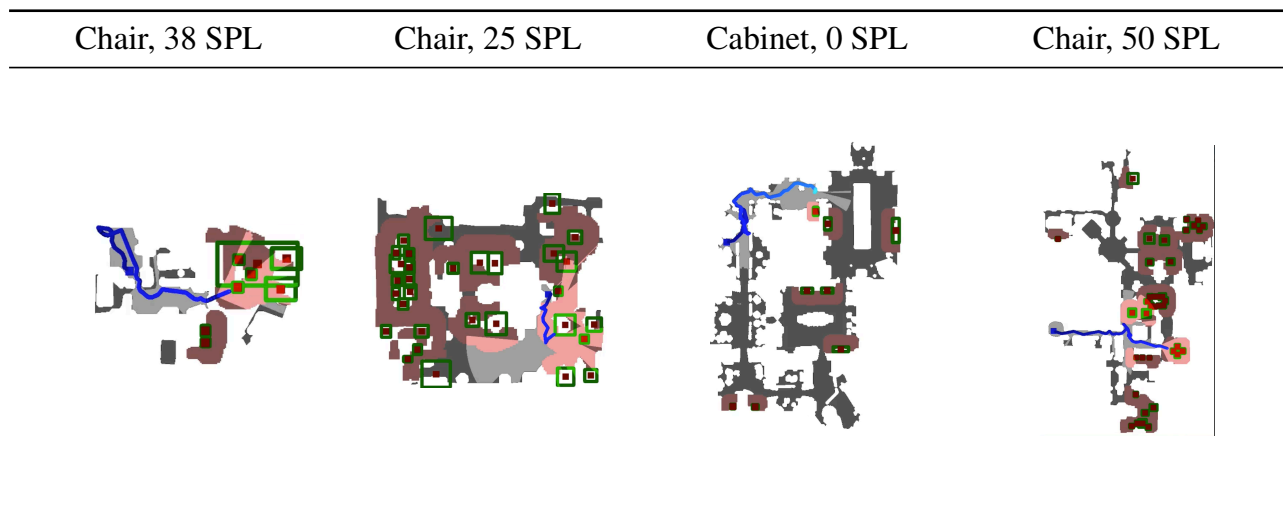


Figure 5: Example episodes for Depth + GoalSeg + THDA + fine-tune model on val split of original ObjectNav dataset.




Figure 6: Success and SPL per goal object category for Depth + GoalSeg + THDA + fine-tune model on val split of original ObjectNav dataset.

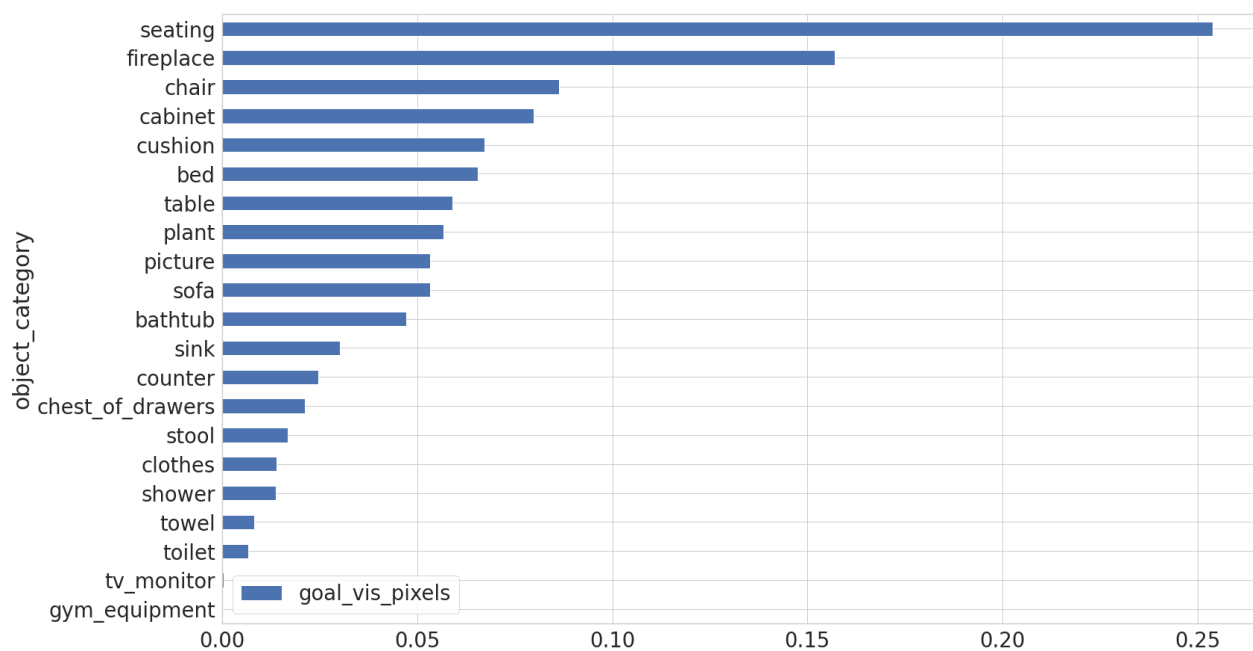


Figure 7: Goal segmentation occupation on the last frame when STOP was called per goal object category for Depth + GoalSeg + THDA + fine-tune model on val split of original ObjectNav dataset.

		Train			Val		
		IoU	Precision	Recall	IoU	Precision	Recall
Goal-Objects	background	79.97	99.66	80.19	71.45	89.36	78.09
	chair	49.91	54.41	85.77	36.13	42.45	70.83
	table	42.28	44.96	87.65	19.10	24.95	44.91
	picture	31.90	32.75	92.41	12.18	22.89	20.65
	cabinet	40.06	41.30	93.03	19.83	26.11	45.19
	cushion	49.10	51.37	91.75	35.35	42.59	67.52
	sofa	63.82	68.77	89.86	40.80	55.71	60.37
	bed	68.29	70.97	94.75	27.79	42.02	45.07
	drawers	23.77	24.05	95.40	3.75	5.81	9.52
	plant	45.14	47.60	89.73	25.60	27.75	76.77
	sink	10.80	10.87	94.74	1.12	1.24	10.22
	toilet	4.36	4.37	94.46	1.28	1.29	88.64
	stool	31.41	31.68	97.37	10.53	17.68	20.64
	towel	3.17	3.18	96.92	0.20	0.30	0.64
	tv	44.41	44.83	97.91	16.16	20.36	43.94
	shower	20.81	22.92	69.32	2.70	3.03	19.47
	bath tub	32.05	32.50	95.85	0.00	0.00	0.00
	counter	47.54	48.46	96.16	12.97	28.06	19.44
	fireplace	44.99	45.46	97.72	10.04	11.98	38.31
	gym equipment	15.75	16.02	90.12	0.20	1.29	0.24
YCB	seating	63.77	65.09	96.91	2.68	19.00	3.02
	clothes	15.09	15.12	98.69	6.38	7.28	33.98
	foodstuff	60.25	61.19	97.52	55.22	56.27	96.74
	stationery	74.67	75.68	98.24	63.85	64.97	97.38
	fruit	30.32	30.52	97.89	25.35	25.58	96.57
	plaything	62.25	63.06	97.99	61.18	61.64	98.80
	hand tool	67.37	68.55	97.52	69.93	71.02	97.85
	game equipment	90.71	91.33	99.26	83.61	86.22	96.51
	kitchenware	20.41	20.50	97.87	11.12	11.22	92.79

Table 7: Category-level performances of the Semantic Segmentation module on the Matterport3D validation split.




Figure 8: Semantic segmentation predictions. The left column shows the RGB input frame (the depth input is not show in this figure for simplicity). The two middle columns show the semantic segmentation ground truth and prediction. The right most column shows the GoalSeg probabilities.