We thank all the reviewers for their helpful feedback and positive view of our work. To address the reviewers concerns (**R3,R4,R5**), we have added a comparison to Duan et al.'s One Shot Imitation learning in Tables 3 and 2, a comparison to a non-normalized TECNet ablation, as well as an evaluation on a VizDoom navigation task in Table 1. We believe that these additions address all of the main reviewer concerns.

R3,R5 *Lacking some important references on learning using task embeddings / goal representations, etc* Thank you for the pointers to more related work; missing these was an oversight. Hausman et al. is discussed at line 95 of the paper, and we will add discussion of the additional papers, as well as the language HRL paper from R3, in the related work section. Combining CPVs with natural language task descriptions is an interesting avenue for future work.

R5 *TECNet normalizes the embedding. Is the embedding normalized in the same way in your model?* We do not normalize the CPV embeddings. To tease apart the effects of normalization, we have added a comparison to non-normalized versions of TECNet, which are labeled "TE" (task embedding) in Table 3 and 2.

R4,R5 *More Comparisons.* We have added a comparison to Duan et al.'s One shot imitation method. As the authors have not released an implementation of the method or environment, we implemented the key details of the method: the reference demonstration is encoded with a residual 1D convolution, and the LSTM policy attends over the reference trajectory. This and the TE comparison will be run on the VizDoom env for the camera ready version.

R4, R5 *Lack of motivation behind introducing new environments.* We agree that benchmark environments are ideal for the integrity of the field. To address this, we have added a navigation task from VizDoom. Unfortunately, most currently available environments are too simple to benefit from a compositional representation of tasks. The environment in Duan et al. was not made public. The environment from Sohn et al. was only released this summer (after the deadline). We are releasing our environments publicly with documentation, training code, and demonstration data.

Table 1: **ViZDoom Navigation Results.** We evaluate our method in ViZDoom where the goal is to visit waypoints in a predetermined order. The actions are "turn left," "turn right," and "go forward." The observation space consists of a first person image observation as well as the locations of the waypoints. We evaluate on trajectories that must visit 1 or 2 waypoints (skills), and also evaluate on the compositions of these trajectories. The policies were only trained on trajectories that visit up to 3 waypoints.All numbers are success rates of arriving within 1 meter of each waypoint.

Model	1 Skill	2 Skills	1+1	2+2
Naive	97	94	36.7	2
TECNet	96	95.3	48.3	0
CPV	93	90.7	**91**	**64**

Table 2: **3D Pick and Place Results.** We added the TE and TE-Pair baselines, which use a task embedding like TECNet but without normalizing to the unit ball. TE-Pair has a triplet margin loss so that embeddings of the same task should be close in feature space, which is like the TECNet margin loss. The plain TE only uses the imitation learning loss. While TE performs well at the training tasks in this environment, it does not succeed at compositions of tasks. The Duan et al. architecture fails in this environment.

Model	1 Skill	2 Skills	1,1
TECNet	82 ± 6	50 ± 2	33 ± 4
CPV	87 ± 2	$\mathbf{55} \pm 2$	$\mathbf{52} \pm 2$
TE	$\mathbf{91} \pm 2$	$\mathbf{55} \pm 5$	22 ± 2
TE-Pair	81 ± 11	51 ± 8	15 ± 3
Duan et al.	6 ± 1	0 ± 0	0 ± 0

R5 *How does CPV compare to other imitation learning algorithms such as Behavioral Cloning, Dagger, or GAIL?* CPV can be used in conjunction with any imitation learning algorithm. In our results we use behavioral cloning, and we plan to try IRL methods such as GAIL in future work.

Figure 1: First person view in VizDoom env. The agent must navigate through multiple waypoints.

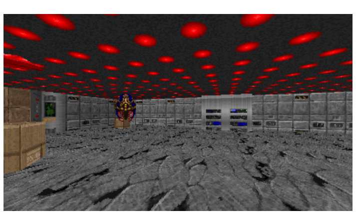

Table 3: **2D crafting results.** The TE ablation, which is like TECNet but with un-normalized embeddings performs worse than TECNet. The Duan et al. architecture performs well in this crafting environment.

Model	4 Skills	8 Skills	16 Skills	2,2	2,2,2,2	4,4
TECNet	50 ± 14	39 ± 5	8 ± 11	52 ± 15	17 ± 5	34 ± 22
CPV	65 ± 10	$\mathbf{80} \pm 3$	44 ± 9	55 ± 5	29 ± 9	$\mathbf{58} \pm 8$
CPV-Hom.	$\mathbf{84} \pm 12$	$\mathbf{82} \pm 15$	54 ± 8	$\mathbf{71} \pm 1$	29 ± 10	48 ± 14
TE	29 ± 4	21 ± 29	3 ± 2	35 ± 10	20 ± 12	15 ± 7
TE-Pair	25 ± 6	12 ± 2	0 ± 0	36 ± 9	10 ± 3	12 ± 4
Duan et al.	75	67	**80**	59	**62**	**66**