

Are We There Yet? Learning to Localize in Embodied Instruction Following

Shane Storks^{1*}, Qiaozi Gao², Govind Thattai², Gokhan Tur²

¹ University of Michigan ² Amazon Alexa AI
sstorks@umich.edu, {qzgao, thattg, gokhatur}@amazon.com

Abstract

Embodied instruction following is a challenging problem requiring an agent to infer a sequence of primitive actions to achieve a goal environment state from complex language and visual inputs. Action Learning From Realistic Environments and Directives (ALFRED) is a recently proposed benchmark for this problem consisting of step-by-step natural language instructions to achieve subgoals which compose to an ultimate high-level goal. Key challenges for this task include localizing target locations and navigating to them through visual inputs, and grounding language instructions to visual appearance of objects. To address these challenges, in this study, we augment the agent’s field of view during navigation subgoals with multiple viewing angles, and train the agent to predict its relative spatial relation to the target location at each timestep. We also improve language grounding by introducing a pre-trained object detection module to the model pipeline. Empirical studies show that our approach exceeds the baseline model performance.

Introduction

As in-home robots become a reality, there are still fundamental problems along the intersection of vision, language, and robotics to be solved. A particularly challenging problem is embodied instruction following, where an agent must complete a task by following a human teacher’s language instructions. The language instructions may require the agent to navigate through a space, and manipulate objects in the space. Thus, navigation skills and language grounding are essential capabilities for such agents, and this research area has received significant attention in recent years.

While much progress has been made toward this in graph-structured, discrete navigation environments (Anderson et al. 2018; Zhu et al. 2020; Majumdar et al. 2020), the problem remains challenging in continuous navigation environments. Action Learning From Realistic Environments and Directives (ALFRED) is a recently introduced benchmark dataset for embodied task learning, providing human-written natural language instructions for an agent to perform tasks in a virtual environment with continuous navigation

*Work performed during an internship with Amazon Alexa AI. Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

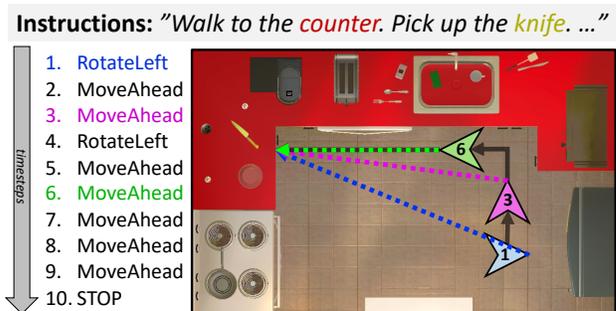


Figure 1: To improve an embodied agent’s navigation capability, we use enhanced inputs from the environment to track the direction of the goal location from the agent’s perspective at every timestep (colored dotted lines). This signal guides the agent during navigation.

(Shridhar et al. 2020). We focus our work on this challenging benchmark.

To address key weaknesses of existing approaches for ALFRED, we implement several enhancements in this work. First, we train the model on a step-by-step granularity of language instructions. Next, we augment the inputs for navigation by adding panoramic visual observations at each timestep of navigation, and object segmentation masks of all objects in the environment. We use these inputs to train an object detection module that can be used at inference time. Lastly, we propose a novel, transformer-based localization module which we train to predict the agent’s orientation angle with respect to the goal location. This module is then used at inference time to guide the agent during navigation, as shown in Figure 1. From these modifications, we show some improvements over a strong baseline approach.

Related Work

Language, vision, and robotics. Within the intersection of language, vision, and robotics, there are several threads of work toward embodied instruction following, navigation, and localization within a virtual environment. As such problems rely on the availability of virtual environments and large-scale data, many of these threads have spawned from the release of benchmarks and challenge datasets. Chang

et al. (2017), Kolve et al. (2017), Yan et al. (2018), Savva et al. (2019), and Xia et al. (2019) propose various virtual environments for embodied AI research, each with their own considerations and features. Das et al. (2018) framed the traditional task of QA in a virtual environment, posing questions that required an agent to explore an environment to find the answer. Qi et al. (2020) proposes a remote object grounding challenge, where an embodied agent must find a particular object in a virtual environment. Xia et al. (2020) frames robotic motion planning in a virtual environment with movable objects as a reinforcement learning problem. Anderson et al. (2018) proposes the vision-and-language navigation (VLN) task, where an agent must use human-written language instructions to reach a goal location in a photorealistic virtual environment, and Thomason et al. (2019) expands this task to a dialog-based formulation.

Vision-and language navigation. A particularly substantial body of work toward navigation in embodied instruction following exists in the VLN thread. Fried et al. (2018) introduced a pragmatic reasoning module to the problem that learned to re-tell candidate navigation routes in natural language, and use this output to re-rank the routes. Jain et al. (2019) introduced a new graph-based metric for navigation fidelity, and used this as a new training objective to improve the fidelity of agents to navigation instructions. Ma et al. (2019a) introduced another auxiliary training objective for the agent to estimate its progress toward the goal location, and Ma et al. (2019b) expanded upon this to give the agent the ability to backtrack during navigation. Hu et al. (2019) improve language grounding in navigation agents by using a pre-trained object detection system to augment visual inputs with localized representations of objects in the scene. Zhu et al. (2020) combined several previously proposed and novel self-supervised auxiliary training objectives for VLN to achieve state-of-the-art performance on the task. Majumdar et al. (2020) pre-trained a transformer on a large-scale language grounding task to further improve performance.

ALFRED Benchmark and Baseline

In this work, we focus on the ALFRED benchmark. ALFRED is a recent embodied task learning benchmark where an agent must complete household tasks to achieve a specific goal state in a virtual environment powered by AI2 Thor (Kolve et al. 2017). First, we introduce the problem and the existing baseline approach formally.

Problem Formulation

An instance of ALFRED consists of a high-level *goal* G which consists of a sequence of N *subgoals* $g_i \in G$. Each subgoal may be a *navigation* subgoal or a *manipulation* subgoal, e.g., to pick up, put down, heat, or cool an object. Navigation subgoals consist of primitive navigation actions, e.g., to turn, move forward, or change the vertical heading angle, while manipulation subgoals consist mostly of primitive manipulation actions, e.g., to pick up, slice, or toggle an object. A typical ALFRED instance consists of alternating navigation and manipulation subgoals, and all primitive actions are

performed sequentially.

These subgoals, when completed in order, achieve a goal state S^* in the virtual environment. The internal representation of goals and subgoals are not provided for inference. Instead, the agent needs to reason from the associated human annotated natural language instructions. From language instructions L_G and $L_i, i \in \{1, 2, \dots, N\}$ describing the goal and subgoals respectively, as well as visual observations v_t at each timestep t , the agent must predict an action a_t and (if applicable) a binary mask m_t over v_t which highlights an object to interact with. After an episode length of T timesteps, an inferred sequence of predicted actions a_1, a_2, \dots, a_T and masks m_1, m_2, \dots, m_T is considered successful if and only if the final state S_T is equal to S^* .

Evaluation ALFRED employs three primary modes of evaluation, which we call *action-by-action*, *subgoal-by-subgoal*, and *goal-by-goal*. Action-by-action evaluation is used while training the model to select the best model instance during hyperparameter search. It aims to score only the action sequence predicted by the agent by comparing it to the ground truth action sequence. To facilitate this evaluation, at each timestep, we simply supply the model with the ground truth action sequence up until the timestep, and compare the model’s output action with the ground truth action. This can give us some ideas on how well the model can choose the next primitive action based on language instructions and the ground truth episode history.

To evaluate the composition of both an agent’s action and object interaction mask prediction in achieving particular subgoals, we use subgoal-by-subgoal evaluation. Here, we run the agent and underlying model through the ground truth episode up until a particular subgoal, then allow the agent to attempt to complete the subgoal. The episode will continue until the agent successfully completes the subgoal, predicts stop, or reaches a time limit. The agent will be considered successful if all subgoal conditions are satisfied, e.g., if the instructions direct the agent to pick up an object, that the object is now held by the agent.

To measure how well the agent’s sequential predictions for all subgoals achieve the overall goal in each ALFRED task, we use goal-by-goal evaluation, the primary mode used for ranking on the ALFRED leaderboard.¹ It places the agent into the virtual environment under the initial conditions, and then requires the agent to execute all steps to achieve all subgoals toward a goal. The episode ends when the agent predicts a stop action, reaches some limit of timesteps, or encounters some maximum number of API errors when interacting with the AI2 Thor environment. The agent is then evaluated based on whether the goal state is true in the environment. An agent is considered successful if all goal conditions are satisfied, but we can also score the agent based on the proportion of goal conditions it achieved. For example, if a goal is given as “Rinse off a mug and place it in the coffee maker;”² the goal conditions may be that there exists a mug in the coffee maker, and that mug is clean.

¹<https://leaderboard.allenai.org/alfred/submissions/public>

²Example from Shridhar et al. (2020).

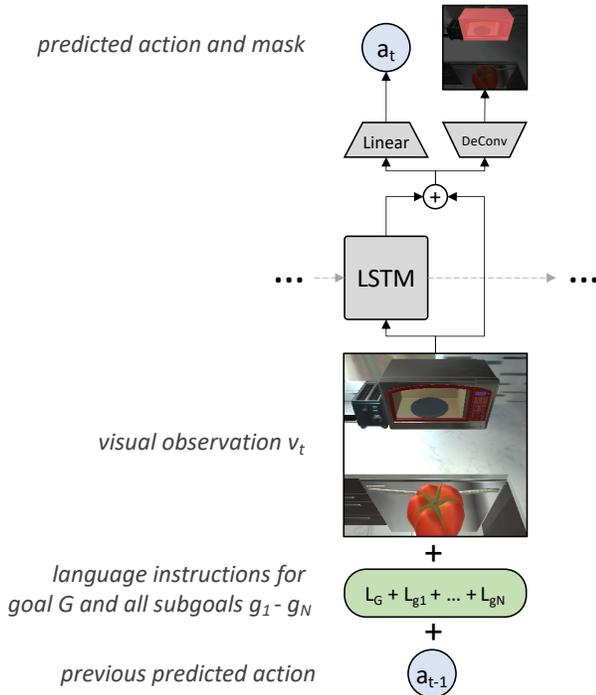


Figure 2: At each timestep, the SEQ2SEQ model takes several encoded inputs: the visual observation v_t , all language instructions $L_G + L_1 + \dots + L_N$ (reweighted by an attention mechanism at every timestep), and the previous predicted action a_{t-1} . These encoded inputs are passed into an LSTM, which generates a history-dependent representation. The original inputs and this representation can be passed to a linear layer to predict a primitive action to take and (if a manipulation action is predicted) to a deconvolutional layer to predict a binary mask over the target object.

Baseline Approach

The baseline SEQ2SEQ approach is shown in Figure 2. First, we generate encoded representations of all language instructions for the example, the current visual observation image, and the previously predicted action. Language is encoded in a bidirectional LSTM (Hochreiter and Schmidhuber 1997), images are encoded in a frozen pre-trained ResNet (He et al. 2016) model, and an embedding is learned for actions. At each timestep, these inputs are passed to a decoder LSTM. The LSTM’s hidden state along with these task inputs are used to predict a primitive action to take (through a linear layer), and a binary mask to be used if a manipulation action is predicted (through a deconvolutional layer). More detailed information on the implementation of the baseline model is provided by Shridhar et al. (2020).

Key Limitations We identify several key opportunities for improvement of the baseline approach. First, the **required sequence of actions is typically quite long** before the goal state is achieved, requiring the completion of several subgoals, which each consist of several primitive actions. Such long-distance dependency is very challenging for current

neural networks.

Second, the agent’s **performance in navigation is a bottleneck to the overall performance**. Nearly every other subgoal in ALFRED task instances requires the agent to navigate to a new location; if the agent failed to arrive at the correct location, it will not find the objects it must manipulate afterward. Using the pre-trained weights of the baseline agent,³ we found that navigation subgoals achieved a subgoal-by-subgoal success rate of 61.9%, relatively low compared to many manipulation subgoal types, such as to cool or toggle an object, which both achieved over 90%. As navigation occurs before every manipulation subgoal and thus is a prerequisite to all other subgoal types, if we can improve the navigation performance, then we expect to improve the overall performance of the model.

We suspect that the navigation performance is so low for two reasons. First, because the baseline model is trained by imitation learning, there are no opportunities to explore during training. During navigation, a task that relies heavily on exploration, this becomes a challenge. The agent is trained only on expert demonstrations where the agent knows exactly where to go, even if the initial location is far away or out of view of the destination. However, if the agent encounters these conditions at inference time, it is unclear whether it could have learned any exploratory behavior to be able to find the destination.

Another possible factor is that the agent is not given any explicit training toward language grounding during navigation subgoals. While the agent is trained to generate binary masks over the target objects of manipulation actions during manipulation subgoals, there is no such supervision available during navigation. Human navigation relies on first identifying landmarks, then finding routes between them (Siegel and White 1975). Navigation instructions in ALFRED often point out landmarks along the way, for example, “Turn around and walk towards the *bed*, then hang a left and walk up to the *wall*, turn left again and walk over to the right side of the wooden *desk*.” Therefore, it is advantageous for the agent to identify them, especially if the landmark refers to the destination.

Proposed Improvements

Considering the key limitations of the baseline model, we propose three main improvements. First, in order to lighten the load on the LSTM used to decode the sequence of actions, we train the model to execute only one subgoal at a time rather than the entire sequence of subgoals. Second, to enable better visual understanding and language grounding, we augmented the task inputs with additional visual observations for panoramic view angles at each timestep, and applied an object detection module to all of the agent’s visual observations. Lastly, to guide navigation, we use these augmented inputs to predict an additional input to the model at every timestep: the angle between the agent’s view and the goal location. This input is predicted using a new *localizer* module before being passed into the LSTM.

³<https://github.com/askforalfred/alfred>

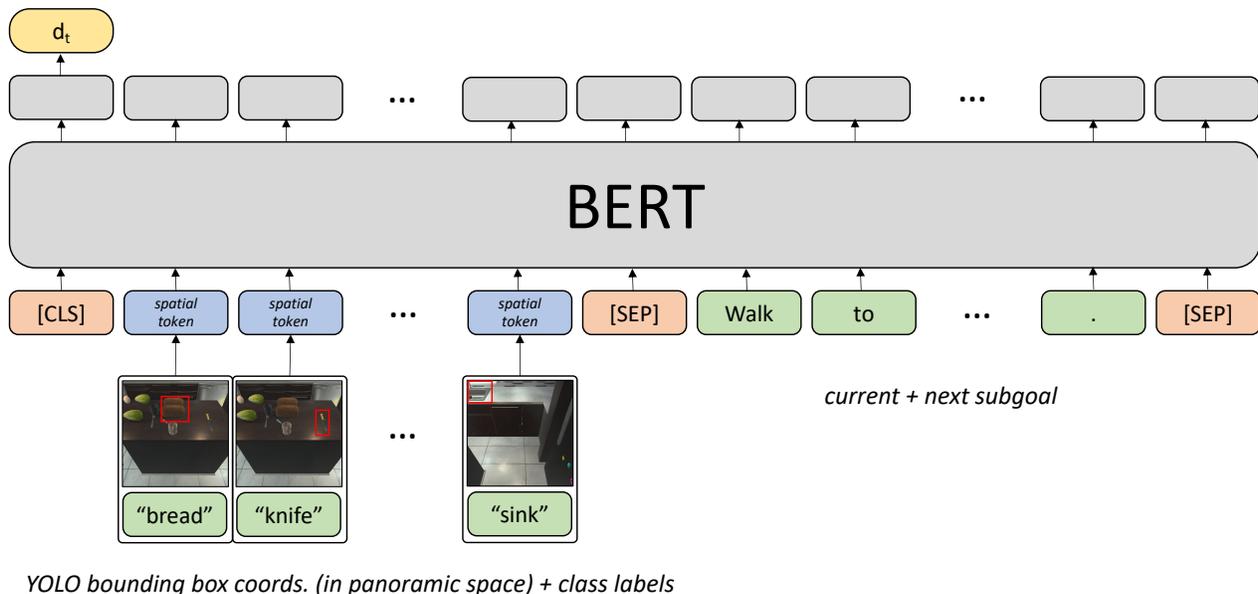


Figure 3: For a navigation subgoal g_k , the multimodal BERT-based localizer takes text and spatial tokens as input. The text input consists of language instructions L_k and L_{k+1} for the current and next subgoal, while the spatial input is comprised of the coordinates (projected into panoramic space) and BERT-embedded object class labels of panoramic bounding boxes at timestep t . The pre-trained BERT model generates a spatial representation of the model’s percepts and language instructions, and we input this representation for the special [CLS] token into a linear layer to predict d_t , the angle from the agent’s orientation to the goal location.

Granular Training

The LSTM-based SEQ2SEQ baseline for ALFRED is trained to predict a sequence of primitive actions from a long sequence of text (consisting of language descriptions of the overall goal and all subgoals), visual observations of the agent, and previously executed actions. However, most subgoals are disentangled from one another with minimal cross-referencing. As such, we reduced the strain on the LSTM-based model by breaking the execution down into one subgoal at a time: for each instance of inference, the model will receive the language instructions for one subgoal, and output the actions and interaction masks required to complete the subgoal. As the language instructions are delimited by subgoal at both training and test time, we can simply iterate through them to facilitate this.

Augmented Navigation

We augment the input data for ALFRED in two ways. First, for navigation subgoals, we add panoramic image observations to the visual input at each timestep. Second, we train an object detection model to identify objects in all visual input images.

Panoramic Visual Observations As the baseline agent is trained by imitation learning, it does not learn any exploratory behavior. During navigation, this becomes problematic. The agent can only see directly in front of itself, so if the referred landmarks or destination are not in view, it may cause difficulties linking language to visual perception.

To mitigate this, for all navigation subgoals in the training data, we use AI2 Thor (Kolve et al. 2017) to generate the agent’s visual observations at eight rotation angles for each timestep. These observations can be used to approximate a panoramic view for the agent at each timestep during navigation, which has proven beneficial in the related vision-and-language navigation (VLN) task (Anderson et al. 2018; Fried et al. 2018). In training, this allows the agent to effectively “look around” before taking a step. During inference, the additional observations can be generated by forcing the agent to take eight consecutive rotation actions at each timestep.

Object Detection Model Training While the baseline agent is trained to explicitly identify objects from language instructions during object manipulation subgoals, there is no such training during navigation subgoals. This inhibits the agent’s ability to ground navigation instructions to its surroundings.

To enable better language grounding, especially during navigation, we introduced YOLO V4 (Bochkovskiy, Wang, and Liao 2020), an object detection model, to the pipeline. The model was trained on robot view images, with bounding boxes supplied by the AI2 Thor. Once trained, this object detection model enables the agent to identify 203 categories of objects in view during inference.

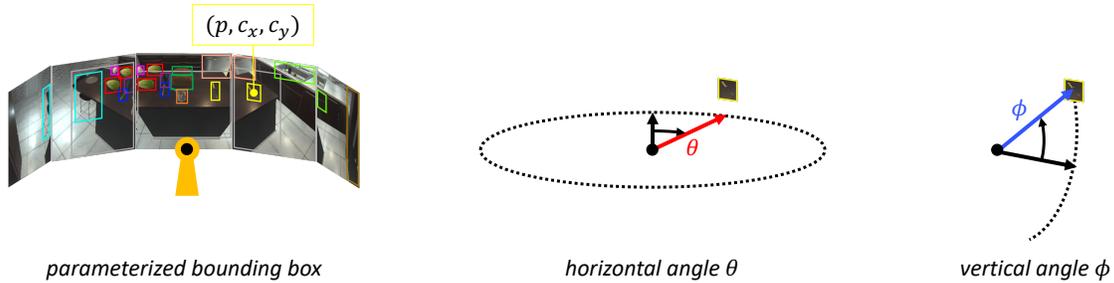


Figure 4: Spatial representation of the bounding box for a *knife*. At each timestep, the agent collects visual inputs for eight panoramic view angles, and applies the trained object detection model to them. Each resulting bounding box can be parameterized by the ordinal view angle index p of its image and the centroid coordinates (c_x, c_y) of the bounding box within the image. We project the parametric representation for each bounding box into panoramic space by using it to calculate the horizontal angle θ and vertical angle ϕ .

Localizer Module

Lastly, we combine these new inputs to introduce a new module to the pipeline: the *localizer*. At each timestep during navigation subgoals, this module predicts a spatial vector d_t representing the angle between the agent’s orientation and direction to the goal location. d_t is then appended to the SEQ2SEQ model’s inputs at each timestep to guide navigation toward the goal location. During non-navigation subgoals, a zero vector is passed instead.

Inputs The agent must predict the angle toward the goal location using task inputs that are available at inference time. These include only the language instructions and any visual observations resulting from taken actions.

In order to characterize the goal location of a navigation subgoal and subsequently predict its direction, it is critical to consider the target objects of both the navigation subgoal and the following object manipulation subgoal. Consider two consecutive instructions: “Walk to the counter” and “Pick up the knife”. The instruction for the navigation subgoal refers to a large object, *the counter*, as the destination. If we only look at this instruction, anywhere around the counter could be a correct location. However, from the next instruction, we know that the agent should stop at a location that it can reach for *the knife*.

As such, given the language instructions L_k and L_{k+1} for consecutive navigation and object manipulation subgoals g_k and g_{k+1} , along with the panoramic visual observations and bounding boxes for all visible objects as described earlier, we aim to predict the direction from the agent’s view to the goal location. As ALFRED provides the ground truth sequence of agent movements and the final goal location, it is straightforward to calculate the ground truth for d_t . To successfully perform this task, the localizer must learn how the position and size of bounding boxes correspond to the direction and distance of objects around the agent, and how these properties may be affected by object type⁴ and context.⁵

⁴For example, at the same distance, the expected sizes of a knife

Implementation The full architecture of the localizer is shown in Figure 3. It is implemented using a multimodal variant of BERT (Devlin et al. 2019) which encodes the panoramic bounding boxes available at each timestep along with the current and next subgoal language instructions to produce a spatial representation of the agent’s observations at each timestep. We perform a regression on the model output to estimate the current angle to the goal location. This model incorporates techniques proposed by Li et al. (2020), who input image region features into BERT alongside language features for impressive results on various multimodal tasks, and Miyazawa et al. (2020), who encode spatial information as input to a similar multimodal BERT model.

Bounding box coordinates in panoramic space. As shown in Figure 4, the augmented navigation inputs give us labeled bounding boxes for images from eight panoramic view angles. In order to uniquely represent the position of each bounding box with respect to the agent, we project it into three-dimensional, panoramic space through two polar coordinates. In this space, we represent a bounding box by its horizontal angle θ and vertical angle ϕ . θ will be calculated relative to the agent’s body orientation, while ϕ will be relative to the center of the agent’s vertical range of head motion.

To calculate these angles, we first parameterize each bounding box by three values: the ordinal index p of the panoramic view angle for the image the bounding box belongs in, and the centroid coordinates (c_x, c_y) of the bounding box within its image normalized by the image width and height. Note that two consecutive panoramic views have an angle of 45 degrees between them. Given these parameters and the horizontal field of view F_x , we can calculate the horizontal angle θ of a particular bounding box (in degrees) by

and of a dining table in the agent’s field of view are different.

⁵For example, if two objects of the same type are visible at once, only the context given by the language instructions can disambiguate the correct destination, e.g., “The knife on the left.”

Model	Action-by-Action F ₁ (%)		Navigation Subgoal Success Rate (%)		Goal Condition Success Rate (%)	
	Val. Seen	Val. Unseen	Val. Seen	Val. Unseen	Val. Seen	Val. Unseen
baseline	84.5	75.6	31.0	27.5	1.6	0.0
step-by-step	91.6	85.3	30.0	26.5	1.3	0.0
oracle	<u>93.9</u>	86.9	<u>67.8</u>	<u>35.4</u>	<u>2.8</u>	0.0
localizer	93.8	88.7	25.4	28.8	1.4	0.0

Table 1: Results for the compared models. Best non-oracle result in bold, and best overall result underlined.

$$\theta = \arctan \left[2(c_x - 0.5) \tan \frac{F_x}{2} \right] + 45p \quad (1)$$

Additionally given the vertical field of view F_y and the agent’s current vertical heading angle δ , we can calculate the vertical ϕ of a particular bounding box (in degrees) by

$$\phi = \arctan \left[2(0.5 - c_y) \tan \frac{F_y}{2} \right] + \delta \quad (2)$$

Input and output details. Once we have calculated θ and ϕ , we encode the spatial representation as a 5-dimensional vector of $\sin \theta$, $\cos \theta$, $\sin \phi$, and the width w and height h of the bounding box, respectively. By incorporating sines and cosines, we account for the circular nature of angles and restrict the range of the values. This encoding is repeated to be comparable size to a BERT embedding, then added to the BERT embedding for the bounding box class label, e.g., *knife*, to create a spatial token ready for input to BERT. As shown in Figure 3, the full input to the localizer is a concatenation of the special [CLS] token, the spatial tokens, the [SEP] token, the BERT embedding for the current and next subgoal language instructions L_k and L_{k+1} , and another [SEP] token.

After being processed by BERT into a spatial representation for the agent’s inputs, we use a linear layer to predict d_t , the sine and cosine of the angle toward the goal location. Specifically, we use the generated representation for the special [CLS] token. Again, sine and cosine are used to account for the circular nature of angles.

Empirical Results

Due to prohibitive training time, we conduct experiments on examples from one task type *stack and place*, where each instance consists of both navigation actions and manipulation actions on multiple objects. We compare the performance of the following approaches:

- SEQ2SEQ baseline (*baseline*)
- SEQ2SEQ baseline + granular training (*step-by-step*)
- SEQ2SEQ baseline + granular training + oracle spatial tracking (*oracle*)
- SEQ2SEQ baseline + granular training + localizer spatial tracking (*localizer*)

For the remainder of this section, we introduce the oracle approach, then present the results.

Oracle spatial tracking. To gauge the effectiveness of this spatial tracking approach and compare our model to a perfect model, we introduce an oracle spatial tracking model. At each timestep during navigation subgoals, rather than predict the angle toward the goal location d_t , we feed the ground truth value of this into the model directly. This gives us an upper bound on the performance of this spatial tracking approach.

Metrics. At inference time, we evaluate the model with metrics in the three increasingly strict modes of evaluation for our models derived from Shridhar et al. (2020): action-by-action, subgoal-by-subgoal, and goal-by-goal. To judge the agent’s prediction of primitive actions, we use the *action-by-action F-measure* of the predicted sequences of primitive actions for entire goal trajectories compared to the ground truth. To judge navigation performance, we also calculate the *navigation subgoal success rate*, i.e., the percentage of navigation subgoals the agent successfully completes. As we trained our models on a relatively small subset of the data, no models achieve a viable goal completion success rate in this evaluation. Consequently, to judge the ability of the agent to achieve the overall goal of ALFRED tasks, we only report the *goal condition success rate*, the percentage of all goal conditions achieved by the model.

All metrics are calculated for ALFRED’s two validation sets: validation *seen* and *unseen*. The seen dataset uses virtual rooms from AI2 Thor which were seen in training, while the unseen uses rooms that were not seen in training.

Results interpretation. The evaluation results for the compared models are listed in Table 1. We see that introducing granular training in the step-by-step model gives us a significant improvement in the prediction of individual actions, but achieves comparable results for other metrics. While the oracle spatial tracking approach offers drastic improvements in navigation subgoal and goal performance, the non-oracle BERT-based localizer approach does not come close to this upper bound. Nonetheless, it provides an overall net improvement in action-by-action F-measure of 9.3% and

13.1% over the baseline for seen and unseen rooms, respectively, and provides a slight improvement in navigation subgoal success rate over the baseline. This may indicate that in the task of predicting the spatial direction of targets, the transformer-based model still has much room for improvement.

Conclusion and Future Work

In this work, we investigated several methods to improve the ALFRED baseline SEQ2SEQ model, including subgoal-by-subgoal granular training, augmenting navigation inputs with panoramic visual observation images and a full coverage of object segmentation masks, and combining these new inputs to propose a novel BERT-based spatial tracking module.

While granular training considerably improves precision of predicted actions, we show through an oracle approach that if the agent is given the angle toward the goal location as input at every timestep, we can achieve astounding performance improvements in navigation and goal completion. This suggests that improving navigation performance can indeed be a key to solving this problem. Our fair, learning-based spatial tracking approach makes some progress toward this grand goal of closing the gap to oracle performance, but we expect there is much more work to be done. Beyond this, future work may include extending BERT or similar transformer-based models to handle the entire task end-to-end, including action and object interaction mask prediction.

References

- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and van den Hengel, A. 2018. Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Bochkovskiy, A.; Wang, C.-Y.; and Liao, H.-Y. M. 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv preprint arXiv:2004.10934*.
- Chang, A.; Dai, A.; Funkhouser, T.; Halber, M.; Niessner, M.; Savva, M.; Song, S.; Zeng, A.; and Zhang, Y. 2017. Matterport3D: Learning from RGB-D Data in Indoor Environments. *International Conference on 3D Vision (3DV)*.
- Das, A.; Datta, S.; Gkioxari, G.; Lee, S.; Parikh, D.; and Batra, D. 2018. Embodied Question Answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics. doi:10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- Fried, D.; Hu, R.; Cirik, V.; Rohrbach, A.; Andreas, J.; Morency, L.-P.; Berg-Kirkpatrick, T.; Saenko, K.; Klein, D.; and Darrell, T. 2018. Speaker-Follower Models for Vision-and-Language Navigation. *arXiv:1806.02724 [cs]*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation* 9(8): 1735–1780. ISSN 0899-7667. doi:10.1162/neco.1997.9.8.1735.
- Hu, R.; Fried, D.; Rohrbach, A.; Klein, D.; Darrell, T.; and Saenko, K. 2019. Are You Looking? Grounding to Multiple Modalities in Vision-and-Language Navigation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6551–6557. Florence, Italy: Association for Computational Linguistics. doi:10.18653/v1/P19-1655. URL <https://www.aclweb.org/anthology/P19-1655>.
- Jain, V.; Magalhaes, G.; Ku, A.; Vaswani, A.; Ie, E.; and Baldrige, J. 2019. Stay on the Path: Instruction Fidelity in Vision-and-Language Navigation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 1862–1872. Florence, Italy: Association for Computational Linguistics. doi:10.18653/v1/P19-1181. URL <https://www.aclweb.org/anthology/P19-1181>.
- Kolve, E.; Mottaghi, R.; Han, W.; VanderBilt, E.; Weihs, L.; Herrasti, A.; Gordon, D.; Zhu, Y.; Gupta, A.; and Farhadi, A. 2017. AI2-THOR: An Interactive 3D Environment for Visual AI. *arXiv*.
- Li, X.; Yin, X.; Li, C.; Hu, X.; Zhang, P.; Zhang, L.; Wang, L.; Hu, H.; Dong, L.; Wei, F.; Choi, Y.; and Gao, J. 2020. Oscar: Object-Semantics Aligned Pre-training for Vision-Language Tasks. *arXiv preprint arXiv:2004.06165*.
- Ma, C.-Y.; Lu, J.; Wu, Z.; AlRegib, G.; Kira, Z.; Socher, R.; and Xiong, C. 2019a. Self-Monitoring Navigation Agent via Auxiliary Progress Estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*. URL <https://arxiv.org/abs/1901.03035>.
- Ma, C.-Y.; Wu, Z.; AlRegib, G.; Xiong, C.; and Kira, Z. 2019b. The Regretful Agent: Heuristic-Aided Navigation through Progress Estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. URL <https://arxiv.org/abs/1903.01602>.
- Majumdar, A.; Shrivastava, A.; Lee, S.; Anderson, P.; Parikh, D.; and Batra, D. 2020. Improving Vision-and-Language Navigation with Image-Text Pairs from the Web. *arXiv:2004.14973 [cs]*.
- Miyazawa, K.; Aoki, T.; Horii, T.; and Nagai, T. 2020. lamBERT: Language and Action Learning Using Multimodal BERT. *arXiv preprint arXiv:2004.07093*.
- Qi, Y.; Wu, Q.; Anderson, P.; Wang, X.; Wang, W. Y.; Shen, C.; and van den Hengel, A. 2020. REVERIE: Remote Embodied Visual Referring Expression in Real Indoor Environ-

ments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Savva, M.; Kadian, A.; Maksymets, O.; Zhao, Y.; Wijmans, E.; Jain, B.; Straub, J.; Liu, J.; Koltun, V.; Malik, J.; Parikh, D.; and Batra, D. 2019. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

Shridhar, M.; Thomason, J.; Gordon, D.; Bisk, Y.; Han, W.; Mottaghi, R.; Zettlemoyer, L.; and Fox, D. 2020. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *Computer Vision and Pattern Recognition (CVPR)*.

Siegel, A. W.; and White, S. H. 1975. The Development of Spatial Representations of Large-Scale Environments. In *Advances in Child Development and Behavior*, volume 10, 9–55. Elsevier.

Thomason, J.; Murray, M.; Cakmak, M.; and Zettlemoyer, L. 2019. Vision-and-Dialog Navigation. In *2019 Conference on Robot Learning (CoRL 2019)*.

Xia, F.; Li, C.; Chen, K.; Shen, W. B.; Martín-Martín, R.; Hirose, N.; Zamir, A. R.; Fei-Fei, L.; and Savarese, S. 2019. Gibson Env V2: Embodied Simulation Environments for Interactive Navigation. Technical report, Stanford University.

Xia, F.; Shen, W. B.; Li, C.; Kasimbeg, P.; Tchapmi, M. E.; Toshev, A.; Martín-Martín, R.; and Savarese, S. 2020. Interactive Gibson Benchmark: A Benchmark for Interactive Navigation in Cluttered Environments. *IEEE Robotics and Automation Letters* 5(2): 713–720.

Yan, C.; Misra, D.; Bennet, A.; Walsman, A.; Bisk, Y.; and Artzi, Y. 2018. Chalet: Cornell house agent learning environment. *arXiv preprint arXiv:1801.07357*.

Zhu, F.; Zhu, Y.; Chang, X.; and Liang, X. 2020. Vision-Language Navigation With Self-Supervised Auxiliary Reasoning Tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.