Transformer-based Localization from Embodied Dialog with Large-scale Pre-training

Meera Hahn Georgia Institute of Technology meerahahn@gatech.edu

Abstract

We address the challenging task of Localization via Embodied Dialog (LED). Given a dialog from two agents, an Observer navigating through an unknown environment and a Locator who is attempting to identify the Observer's location, the goal is to predict the Observer's final location in a map. we develop a novel LED-Bert architecture and present an effective pretraining strategy. We show that a graphbased scene representation is more effective than the top-down 2D maps used in prior work. Our approach outperforms previous baselines.

1 Introduction

A key goal of AI is to develop embodied agents that can effectively communicate with humans and other agents using natural language. A basic capability for these agents is the ability to perceive and navigate through an environment and respond to instructions and questions about the space they are in. The recently-introduced Where Are You? (WAY) dataset (Hahn et al., 2020) provides a setting for developing such a multi-modal and multi-agent paradigm. This dataset (collected via AMT) contains episodes of a localization scenario in which two agents communicate via turn-taking natural language dialog: An Observer agent moves through an unknown environment, while a Locator agent attempts to identify the Observer's location in a map.

The Observer produces descriptions such as '*I*'m in a living room with a gray couch and blue armchairs. Behind me there is a door.' and can respond to instructions and questions provided by the Locator: 'If you walk straight past the seating area, do you see a bathroom on your right?' Via this dialog (and without access to the Observer's view of the scene), the Locator attempts to identify the Observer's location on a map (which is not available to the Observer). This is a complex task for which a successful localization requires accurate situational James M. Rehg Georgia Institute of Technology rehg@gatech.edu

grounding and the production of relevant questions and instructions.

One of the benchmark tasks supported by WAY is *Localization via Embodied Dialog (LED)*. This task involves the development of a machine learning model that takes the dialog and a representation of the map as inputs, and outputs the final location of the Observer agent. LED is a first step towards developing a Locator agent. Two basic issues are to identify an effective map representation and deep learning architecture. The LED baseline from (Hahn et al., 2020) uses 2D images of top down (birds-eye view) floor maps to represent the environment and an (x,y) location for the Observer.

This paper provides a new solution to the LED task with two key components. First, we propose to model the environment using the first person view (FPV) panoramic navigation graph from Matterport (Anderson et al., 2018a), as an alternative to top-down maps. Second, we introduce a novel visiolinguistic transformer model, LED-Bert, which scores the alignment between navigation graph nodes and dialogs. LED-Bert is an adaption of ViLBERT (Lu et al., 2019) for the LED task, and we show that it outperforms all prior baselines. A key challenge is the small size of the WAY dataset (approximately 6K episodes), which makes it challenging to use transformer-based models given their reliance on large-scale training data. We address this challenge by developing a pretraining approach (based on (Majumdar et al., 2020)) that yields an effective visuolinguistic representation.

Contributions: To summarize:

- 1. We demonstrate an LED approach using navigation graphs to represent the environment.
- 2. We present LED-Bert, a visiolinguistic transformer model which scores alignment between graph nodes and dialogs, and we develop an effective pretraining strategy.
- 3. We show that LED-Bert outperforms all baselines, increasing accuracy at 0m by 8.21 abso-

lute percent on the test split.

2 Related Work

ViLBERT ViLBERT (Lu et al., 2019) is a multimodal transformer that extends the BERT architecture (Devlin et al., 2018) to learn joint visiolinguistic representations. Similar multi-modal transformer models exist (Li et al., 2020, 2019; Su et al., 2019; Tan and Bansal, 2019; Zhou et al., 2020). ViLBERT is constructed of two transformer encoding streams, one for visual inputs and one for text inputs. Both of these streams use the standard BERT-BASE (Devlin et al., 2018) backbone. The input for the text stream is then text tokens, identical to BERT. The visual tokens for the visual stream are a sequence of image regions which are generated by an object detector pretrained on Visual Genome. The input to ViLBERT is then a sequence of visual and textual tokens which are not concatenated and only enter their respective streams. The two streams then interact using coattention layers which are implemented by swapping the key and value matrices between the visual and textual encoder streams for certain layers. Coattention layers are used to attend to one modality via a conditioning on the other modality, thus allowing for attention over image regions given the corresponding text input and vise versa.

Vision-and-Language Pre-training Prior work has experimented with utilizing dual-stream transformer based models that have been pretrained with self-supervised objectives and transferring them to downstream multi-modal tasks with large success. This has been seen for tasks such as Visual Question Answering (Antol et al., 2015), Commonsense Reasoning (Zellers et al., 2019), Natural Language Visual Reasoning (Suhr et al., 2018), Image-Text Retrieval (Lee et al., 2018), and Visual-Dialog (Murahari et al., 2020) and Vision Language Navigation (Majumdar et al., 2020). Specifically VLN-Bert and VisDial + BERT adapt the ViLBERT architecture and utilize a pretraining scheme which inspired our approach to training LED-Bert.

3 Approach

3.1 Environment Representation

A key challenge in the LED task is that environments often have multiple rooms with extremely similar attributes, i.e. multiple bedrooms with the same furniture. Therefore a successful model must be able to visually ground fine-grained attributes. Strong generalizability is also required in order to generalize to unseen test environments. The LED baseline in (Hahn et al., 2020) approaches localization as a language-conditioned pixel-to-pixel prediction task – producing a probability distribution over positions in a top-down view of the environment, illustrated in Part A, in the Supplementary, Figure 2. This choice is justified by the fact that it mirrors the observations that human Locators had access to during data collection, allowing for a straightforward comparison. However, this does not address the question of what representation is optimal for localization.

We propose to use a navigation-graph map representation derived from the panoramic-RGB graphs of the Matterport environments (Chang et al., 2017), illustrated in Part B, in the Supplementary, Figure 2. The Observer agent traverses these same navigation graphs during data collection, which may result in a strong alignment between the dialog and the nodes. Using this approach, the LED task can be framed as a prediction problem over the possible nodes in the navigation graph. At inference time, this can be accomplished by producing an alignment score between each node in the test environment and the test dialog, and then returning the node with the highest score as the predicted observer location.

3.2 Adapting ViLBERT for LED

To formalize the graph based LED task, we consider a function f that maps a node location n and a dialog x to a compatibility score f(n, x). We model f(n, x) using a visiolinguistic transformerbased model we denote as LED-Bert, shown in the Supplementary, Figure 1. The architecture of LED-Bert is structurally similar to ViLBERT and VLN-Bert (Majumdar et al., 2020), but with some key differences due to our need to ground dialog and fine-tune on the relatively small WAY dataset. This allows us to initialize the majority of the LED-Bert using pretrained weights. We represent each panoramic node I as a set of image regions $r_1, ..., r_k$. Let an dialog x be a sequence of tokens $w1, ..., w_L$. Then for a given dialog-node pair for LED-Bert as the input sequence:

 $\langle IMG \rangle r_1, ..., r_K \langle CLS \rangle w_1, ..., w_L \langle SEP \rangle$ where IMG, CLS, and SEP are special tokens.

Transformer models are by nature invariant to sequence order and they only model interactions between inputs as a function of their values (Vaswani et al., 2017). This leads to the standard practice of adding positional embeddings for each input token to re-introduce order information. For the dialog tokens we simply use an index sequence order encoding. However the panoramic node visual tokens have a more complicated positional encoding, as the panorama is broken up into image regions. The visual positional information is very important for encoding spatial relationships between objects and for scene understanding as a whole. For instance consider the a question the Locator might as 'Are you located to the right of the blue couch?' This question will require information about which part of the panorama the couch is located in. To address this, we follow the VLN-Bert (Majumdar et al., 2020) strategy of encoding the spatial location of each image region, in terms of its location in the panorama (top-left and bottom-right corners in normalized coordinates as well as area of the image covered), its elevation relative to the horizon and all angles are encoded as $[\cos(\theta), \sin(\theta)]$. The resulting 11-dimensional vector S is projected into 2048 dimensions using a learned projection W^S .

3.3 Training Procedure for LED-Bert

LED-Bert can be trained from scratch using the WAY dataset however due to the small size (6k episodes) of the WAY dataset and since large-transformer models have been shown to work best on large amounts of data we follow the 4 stage pretraining procedure of prior work (Majumdar et al., 2020; Murahari et al., 2020; Lu et al., 2019). These works do extensive pretraining for multi-modal transformers using large scale web-data. The pipeline for pretraining has 4 stages and is also visualized in the Supplementary, Figure 1.

Stage 1-3 are the same as (Majumdar et al., 2020), and we replace the 4th stage with fine-tuning for node localization over the WAY dataset. To train LED-Bert for localization, we consider the task as a selection task over the possible nodes in the graph, on average there are 117.32 nodes, with the largest environment containing 345 nodes. We run LED-Bert on each node-dialog pair and extract the final representations for each stream, denoted as $h_{\rm CLS}$ and $h_{\rm IMG}$, using these we compute a compatibility score by doing element-wise multiplication of the two vectors and passing them through a single linear layer. The scores are normalized via a softmax layer and then supervised using a cross-entropy loss against a one-hot vector

with a mass at the ground truth node.

4 **Experiments**

4.1 Baselines

We propose strong baseline methods to compare the LED-Bert architecture against. All approaches will use the panoramic based maps to ensure the same the prediction space of all panoramic nodes in an environment.

Human Performance: Uses the average performance of AMT Locator workers from the WAY dataset. We snap the human prediction over the top down map to the nearest node.

Random: Selects a random node from the test environment as the predicted location for each episode.

Joint Embedding: This baseline learns a common embedding space between the dialogs and corresponding node locations. Each panoramic node is represented by 36 image patches and image features are extracted for each patch. Visual features are extracted using a ResNet152 (He et al., 2016) pretrained on Places 365 (Zhou et al., 2017). We experiment with three types of joint embedding architectures. All models encode the dialog in the same way and is described below.

Graph Convolutional Network Both the jointembedding baselines and LED-Bert discard edge information. We propose to experiment with Graph Convolutional Network (GCN) (Zhang et al., 2019) to model the LED task using the navigation graph as input which incorporates edge information. In the graph representation input to the model, nodes attributes are visual features and edge attributes contain the pose transformation between connected nodes. The goal of the GCN architecture is to model the relational information between the nodes of the graph and the localization dialog in order to produce a probability distribution of localization likelihood over the nodes.

Dialog Encoding: The Locator and Observer messages are tokenized using a standard toolkit (Loper and Bird, 2002). The dialog is represented as a single sequence with identical 'start' and 'stop' tokens surrounding each message, and then encoded using a single-layer bidirectional LSTM. Word embeddings are initialized using GloVe (Pennington et al., 2014) and fine tuned end-to-end. In the first model called the 'late-fusion model', the LSTM has a 2048 dimension hidden state and the node features are down-sampled using self attention to

	val-seen			val-unseen			test		
Method	$\text{LE}\downarrow$	Acc@0m↑	Acc@5m↑	$\text{LE}\downarrow$	Acc@0m \uparrow	Acc@5m↑	$\text{LE}\downarrow$	Acc@0m \uparrow	Acc@5m↑
Human Locator	6.00	47.87	77.38	3.20	56.13	83.42	5.89	44.92	75.00
Random Node	20.8	0.33	10.82	18.61	1.9	11.05	20.93	0.92	11.00
Late Fusion	12.56	17.38	47.54	12.87	7.77	34.37	15.86	8.92	32.75
Attention Model	9.83	18.36	56.07	10.93	10.54	41.11	14.96	6.92	34.42
Attention over History Model	11.64	21.64	49.18	11.44	10.02	43.18	14.98	7.14	33.68
Graph Convolutional Network	10.95	19.67	59.13	9.10	8.64	46.99	14.32	9.46	35.10
LED-BERT	9.04	25.57	60.66	8.82	21.07	52.5	11.12	17.67	51.67

Table 1: Comparison of the LED-BERT model with baselines and human performance on the LED task. We report average localization error (LE) and accuracy at k meters (all \pm standard error).

be of size 2048, the visual and dialog features are fused through late fusion passed through a twolayer MLP and softmax and the output is a prediction over the possible nodes in the environment. In the 'attention model', the visual and dialog features are fused instead through top-down bottom up attention, the final layers of the model are also an MLP and softmax. In the 'attention over history model', there are two separate LSTMs. One to encode dialog history and another to encode the current message. The dialog history based attention is applied over the visual features, then the encoded current message and visual features are fused through late fusion followed by an MLP and softmax.

4.2 Metrics

We propose to evaluate the localization error (LE) of our models using geodesic distance instead of euclidean distance as used in (Hahn et al., 2020). Geodesic distance is more meaningful than euclidean distance for determining error across rooms and across floors in multi-story environments. To discern the precision of the models, we report a binary success metric that places a threshold k on the LE. Accuracy (Acc) at 0 meters indicates the correct node was predicted. Accuracy at k meters indicates that the node predicted was within k meters of the true node.

4.3 Results

Table 1 shows the performance of our LED-BERT model and relevant baselines on the val-seen, val-unseen, and test splits of the WAY dataset.

Human and No-learning Baselines. Humans succeed 44.92% of the time in test environments at 0 meters, this shows it is a difficult task.

Attention and History increase performance. Adding bottom-up and top-down attention increases performance, additionally separating the encoders for the current message from the dialog history further increases performance.

Graph Networks see slight improvement. Graph networks see slight increase in performance on the test split. Graph networks do not have a straight forward large-scale web-data pretraining schema for this task.

LED-BERT outperforms all baselines. LED-BERT significantly outperforms the other crossmodal modeling baselines in terms of both accuracy and localization error – improving the best baseline, Graph Convolutional Network (GCN), by an absolute 7.54% (test) to 12.43% (val-seen and val-unseen). There remains a gap between our model and human performance – especially on novel environments (-% vs -% on test).

5 Conclusion

In summary, we propose a viso-linguistic transformer, LED-BERT, for the LED task and instantiate a new version approach which does localization over the navigation graph. We demonstrate a pre-training schema for LED-BERT which utilizes large scale web-data as well as other multi-modal embodied AI task data to learn the visual grounding required for successful localization's in LED. We show LED-BERT is able to achieve SOTA performance and outperform other learned baselines by a significant margin.

6 Supplementary

6.1 Training Procedure for LED-Bert

LED-Bert can be trained from scratch using the WAY dataset however due to the small size (6k episodes) of the WAY dataset and since large-transformer models have been shown to work best on large amounts of data we follow the 4 stage pretraining procedure of prior work (Majumdar et al., 2020; Murahari et al., 2020; Lu et al., 2019).

These works do extensive pretraining for multimodal transformers using large scale web-data. The pipeline for pretraining has 4 stages and is also visualized in Figure 1.

6.2 Environment Representation

The LED baseline in (Hahn et al., 2020) approaches localization as a language-conditioned pixel-topixel prediction task – producing a probability distribution over positions in a top-down view of the environment, illustrated in Part A, Figure 2. In this paper we used a navigation-graph map representation derived from the panoramic-RGB graphs of the Matterport environments (Chang et al., 2017), illustrated in Part B, Figure 2

References

- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. 2018a. Visionand-Language Navigation: Interpreting visuallygrounded navigation instructions in real environments. In Computer Vision and Pattern Recognition.
- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. 2018b. Visionand-language navigation: Interpreting visuallygrounded navigation instructions in real environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In Proceedings of the IEEE international conference on computer vision, pages 2425–2433.
- Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. 2017. Matterport3d: Learning from rgb-d data in indoor environments. arXiv preprint arXiv:1709.06158.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Meera Hahn, Jacob Krantz, Dhruv Batra, Devi Parikh, James M Rehg, Stefan Lee, and Peter Anderson. 2020. Where are you? localization from embodied dialog. *arXiv preprint arXiv:2011.08277*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*.
- Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. Stacked cross attention for

image-text matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 201–216.

- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Edward Loper and Steven Bird. 2002. Nltk: the natural language toolkit. *arXiv preprint arXiv:0205028*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Ad*vances in neural information processing systems, 32.
- Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra. 2020. Improving vision-and-language navigation with imagetext pairs from the web. In *European Conference on Computer Vision*. Springer.
- Vishvak Murahari, Dhruv Batra, Devi Parikh, and Abhishek Das. 2020. Large-scale pretraining for visual dialog: A simple state-of-the-art baseline. In *European Conference on Computer Vision*. Springer.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2019. VI-bert: Pre-training of generic visual-linguistic representations. arXiv preprint arXiv:1908.08530.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2018. A corpus for reasoning about natural language grounded in photographs. arXiv preprint arXiv:1811.00491.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*.



Training Curriculum

Figure 1: We propose the LED-BERT for the LED task. The model is pretrained in 3 stages over different datasets before being fine-tuned over the node-dialog pairs of the WAY dataset (Hahn et al., 2020). The language stream of the model is first pretrained on English Wikipedia and the BooksCorpus (Zhu et al., 2015) datasets. Second, both streams of the model are trained on the Conceptual Captions (Sharma et al., 2018) dataset. Third, both streams are train on the path-instruction pairs of the Room2Room dataset (Anderson et al., 2018b). Finally we fine-tune the model over the node-dialog pairs of the WAY dataset (Hahn et al., 2020).

- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Si Zhang, Hanghang Tong, Jiejun Xu, and Ross Maciejewski. 2019. Graph convolutional networks: a comprehensive review. *Computational Social Networks*.
- Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Places: A 10 million image database for scene recognition. *Transactions on Pattern Analysis and Machine Intelligence*.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. Unified visionlanguage pre-training for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence.*
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*.



Figure 2: Examples of the types of map representations of the Matterport3D (Chang et al., 2017) indoor environments which can be used for the Localization via Embodied Dialogue task. Part A shows the top down floor maps used in the original LED paper. Part B shows an overlay of the navigation graph of panoramic nodes over the top down map, note the lines represent traversability between nodes and the circles represent the panoramic node location. Part C shows examples of the FPV panoramic nodes in different environments. Note each of these images are mapped to a node in a connectivity graph for the respective environment.