

Are You Looking? Grounding to Multiple Modalities in Vision-and-Language Navigation

Ronghang Hu¹, Daniel Fried¹, Anna Rohrbach¹, Dan Klein¹, Trevor Darrell¹, Kate Saenko²

¹University of California, Berkeley ²Boston University

{ronghang, dfried, anna.rohrbach, klein}@cs.berkeley.edu,
trevor@eecs.berkeley.edu, saenko@bu.edu

Abstract

Vision-and-Language Navigation (VLN) requires grounding instructions, such as *turn right and stop at the door*, to routes in a visual environment. The actual grounding can connect language to the environment through multiple modalities, e.g. *stop at the door* might ground into visual objects, while *turn right* might rely only on the geometric structure of a route. We investigate *where* the natural language empirically grounds under two recent state-of-the-art VLN models. Surprisingly, we discover that visual features may actually *hurt* these models: models which only use route structure, ablating visual features, outperform their visual counterparts in unseen new environments on the benchmark Room-to-Room dataset. To better use all the available modalities, we propose to decompose the grounding procedure into a set of expert models with access to different modalities (including object detections) and ensemble them at prediction time, improving the performance of state-of-the-art models on the VLN task.

1 Introduction

The Vision-and-Language Navigation (VLN) task (Anderson et al., 2018) requires an agent to navigate to a particular location in a real-world environment, following complex, context-dependent instructions written by humans (e.g. *go down the second hallway on the left, enter the bedroom and stop by the mirror*). The agent must navigate through the environment, conditioning on the instruction as well as the visual imagery that it observes along the route, to stop at the location specified by the instruction (e.g. *the mirror*).

Recent state-of-the-art models (Wang et al., 2018; Fried et al., 2018b; Ma et al., 2019) have demonstrated large gains in accuracy on the VLN task. However, it is unclear which *modality* these

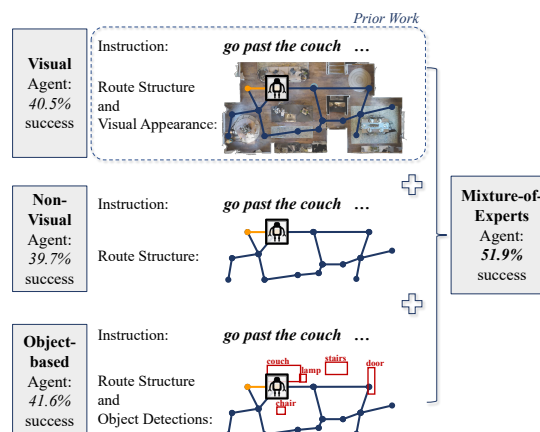


Figure 1: We factor the grounding of language instructions into visual appearance, route structure, and object detections using a mixture-of-experts approach.

substantial increases in task metrics can be attributed to, and, in particular, whether the gains in performance are due to stronger grounding into visual context or e.g. simply into the discrete, geometric structure of possible routes, such as turning left or moving forward (see Fig. 1, top vs. middle).

First, we analyze to what extent VLN models ground language into visual appearance and route structure by training versions of two state-of-the-art models *without* visual features, using the benchmark Room-to-Room (R2R) dataset (Anderson et al., 2018). We find that while grounding into route structure is useful, the models with visual features fail to learn generalizable visual grounding. Surprisingly, when trained without visual features, their performance on unseen environments is comparable or even better.

We hypothesize that the low-level, pixel-based CNN features in the visual models contribute to their failure to generalize. To address this, we introduce a high-level object-based visual representation to ground language into visual context in a more generalizable way, using the symbolic output of a pretrained object detection system. For example, while a concept *table* could ground into visual

appearance of a specific table in a given environment, detecting tables and other objects in scenes, mapping them into symbols, and grounding the text mentions into these symbols should generalize better to unseen environments.

Finally, inspired by the complementary errors of visual and non-visual agents, we decompose the grounding process through a mixture-of-experts approach. We train separate visual and non-visual agents, encouraging each one to focus on a separate modality, and combine their predictions as an ensemble (see Fig. 1). Our mixture-of-experts outperforms the individual agents, and is also better than the ensembles of multiple agents of the same modality (e.g. both visual or both non-visual).

Adding our object representation and mixture-of-experts approach to both state-of-the-art models improves their success rate by over 10% (absolute) in novel environments, obtaining a 51.9% success rate on the *val-unseen* split of the benchmark R2R dataset (Anderson et al., 2018).

2 Related work

Vision and Language Navigation. Vision-and-Language Navigation (VLN) (Anderson et al., 2018; Chen et al., 2019) unites two lines of work: first, of following natural language navigational instructions in an environmental context (MacMahon et al., 2006; Vogel and Jurafsky, 2010; Tellex et al., 2011; Chen and Mooney, 2011; Artzi and Zettlemoyer, 2013; Andreas and Klein, 2015; Mei et al., 2016; Fried et al., 2018a; Misra et al., 2018), and second, of vision-based navigation tasks (Mirowski et al., 2017; Zhu et al., 2017; Yang et al., 2019; Mirowski et al., 2018; Cirik et al., 2018) that use visually-rich real-world imagery (Chang et al., 2017).

A number of methods for the VLN task have been recently proposed. Wang et al. (2018) use model-based and model-free reinforcement learning to learn an environmental model and optimize directly for navigation success. Fried et al. (2018b) use a separate instruction generation model to synthesize new instructions as data augmentation during training, and perform pragmatic inference at test time. Most recently, Ma et al. (2019) introduce a visual and textual co-attention mechanism and a route progress predictor.

These approaches have significantly improved performance on the VLN task, when evaluated by metrics such as *success rate*. However, it

is unclear where the high performance comes from. In this paper, we find that agents without any visual input can achieve competitive performance, matching or even outperforming their vision-based counterparts under two state-of-the-art model models (Fried et al., 2018b; Ma et al., 2019). We also explore two approaches to make the agents better utilize their visual inputs.

The role of vision in vision-and-language tasks.

In several vision-and-language tasks, high performance can be achieved without effective modeling of the visual modality. Devlin et al. (2015) find that image captioning models can exploit regularity in the captions, showing that a nearest-neighbor matching approach can achieve competitive performance to sophisticated language generation models. Hendricks et al. (2018) and Rohrbach et al. (2018) find that neural captioning models often ground object mentions into incorrect objects due to correlations in the training data, and can hallucinate non-existing objects.

Recent work has also investigated single-modality performance in vision-and-language embodiment tasks. Anand et al. (2018) find that state-of-the-art results can be achieved on the EmbodiedQA task (Das et al., 2018) using an agent without visual inputs. Work concurrent to ours evaluates the performance of single-modality models for several embodied tasks including VLN (Thomason et al., 2019), finding that high performance can be achieved on the R2R dataset using a non-visual version of the baseline model (Anderson et al., 2018). In this paper, we show that the same trends hold for two recent state-of-the-art architectures (Ma et al., 2019; Fried et al., 2018b) for the VLN task; we also analyze to what extent object-based representations and mixture-of-experts methods can address these issues.

3 State-of-the-art VLN models do not use vision effectively

We experiment with the benchmark Room-to-Room (R2R) dataset (Anderson et al., 2018) for the Vision-and-Language navigation task, which consists of a set of annotated instructions for routes through *environments* from the Matterport3D dataset (Chang et al., 2017). Each environment is a building, such as a house or office, containing a set of *viewpoints*: physical locations in the environment, each with an associated panoramic image. Viewpoints are connected

in a connectivity graph determined by line-of-sight in the physical environment. See the top row of Fig. 1 for a top-down environment illustration.

In the VLN task, a virtual agent is placed at a particular viewpoint in an environment, and is given a natural language instruction (written by a human annotator) to follow. At each timestep, the agent receives the panoramic image for the viewpoint it is currently located at, and either predicts to move to one of the adjacent connected viewpoints, or to stop. When the agent predicts the stop action, it is evaluated on whether it has correctly reached the end of the route that the human annotator was asked to describe.

In this work, we analyze two recent VLN models, which typify the visual grounding approaches of VLN work: the panoramic “follower” model from the Speaker-Follower (SF) system of Fried et al. (2018b) and the Self-Monitoring (SM) model of Ma et al. (2019). These models obtained state-of-the-art results on the R2R dataset. Both models are based on the encoder-decoder approach (Cho et al., 2014) and map an instruction to a sequence of actions in context by encoding the instruction with an LSTM, and outputting actions using an LSTM decoder that conditions on the encoded instruction and visual features summarizing the agent’s environmental context. Compared to the SF model, the SM model introduces an improved visual-textual co-attention mechanism and a progress monitor component. We refer to the original papers for details on the two models.

To analyze the models’ visual grounding ability, we focus on their core encoder-decoder components. In our experiments, we use models trained without data augmentation, and during inference predict actions with greedy search (i.e. without beam search, pragmatic, or progress monitor-based inference). For SF, we use the publicly released code. For SM, we use a reimplementaion without the progress monitor, which was shown to be most important for search in inference (Ma et al., 2019).

We investigate how well these models ground instructions into visual features of the environment, by training and evaluating them without access to the visual context: setting their visual feature vectors to zeroes during training and testing. We compare performance on the validation sets of the R2R dataset: the *val-seen* split, consisting of the same environments as in training, and the *val-*

#	model arch.	train. appr.	vis. feat.	SR on val-seen	SR on val-unseen
1	SF	stud.-forc.	no vis.	29.7	31.7
2			RN	53.3	29.0
3		teach.-forc.	no vis.	34.1	35.2
4			RN	40.4	29.0
5	SM	stud.-forc.	no vis.	36.1	39.7
6			RN	62.8	40.5
7		teach.-forc.	no vis.	34.3	32.2
8			RN	44.0	32.8

Table 1: Success rate (SR) of the vision-based full agent (“RN”, using ResNet) and the non-visual agent (“no vis.”, setting all visual features to zero) on the R2R dataset under different model architectures (Speaker-Follower (SF) (Fried et al., 2018b) and Self-Monitoring (SM) (Ma et al., 2019)) and training schemes.

unseen split of novel environments. Since we aim to evaluate how well the agents *generalize* to the unseen environments, we focus on the val-unseen split. For both the SF and SM models, we train two versions of the agents, using either the student-forcing or teacher-forcing approaches of Anderson et al. (2018)¹, and select the best training snapshot on the val-seen split.² The results are shown in Table 1. In each block, the two rows show the agent’s performance (under the specific model architecture and training approach) with or without access to the visual features (“RN”: ResNet-152 network (He et al., 2016), “no vis.”: non-visual).

While visual features improve performance on environments seen during training, we see that for the SF architecture the non-visual agent (lines 1 and 3) outperforms the visual agent (lines 2 and 4) on unseen environments under both student-forcing and teacher-forcing training. For SM, the non-visual agent (lines 5 and 7) has a success rate very close to the visual agent (lines 6 and 8). This indicates that these models do not learn generalizable visual perception, so that the visual features may actually *hurt* them in unseen environments.

4 Object representation for better grounding and generalization

In both the SF and SM architectures, the agents use visual features from a pretrained ResNet-152 CNN (He et al., 2016). As the training data for the R2R dataset contains only 61 distinct environments, the agents may overfit to the appearance of the training environments and thus struggle to gen-

¹During training, the agent either follows the ground-truth actions (teacher-forcing) or samples actions from its own prediction (student-forcing). See supplemental for more details.

²Following previous work, we use success rate (SR) as our evaluation metric, where an episode is considered successful if the agent stops within 3 meters of the goal location.

eralize. For example, for the instruction *go down the staircase*, a model may learn to ground *staircase* into a specific staircase in a given training environment, and fail to generalize to staircases with different appearances or in different contexts in unseen environments. We thus propose an *object-based representation*, where object detection results from a pretrained large-scale object detector are used as the environment representation. The object-based representation is intended to prevent overfitting to training scenes and to transfer to new environments better than CNN features.

Both the SF and SM models represent the visual appearance at each location with a set of visual features $\{x_{img,i}\}$, where $x_{img,i}$ is a vector extracted from an image patch at a particular orientation i using a CNN. Both models also use a visual attention mechanism to extract an attended visual feature $x_{img,att}$ from $\{x_{img,i}\}$. For our object-based representation, we use a Faster R-CNN (Ren et al., 2015) object detector trained on the Visual Genome dataset (Krishna et al., 2017). We construct a set of vectors $\{x_{obj,j}\}$ representing detected objects and their attributes. Each vector $x_{obj,j}$ (j -th detected object in the scene) is a concatenation of summed GloVe vectors (Pennington et al., 2014) for the detected object label (e.g. *door*) and attribute labels (e.g. *white*) and a location vector from the object’s bounding box coordinates. We then use the same visual attention mechanism as in Fried et al. (2018b) and Ma et al. (2019) to obtain an attended object representation $x_{obj,att}$ over these $\{x_{obj,j}\}$ vectors. We either substitute the ResNet CNN features $x_{img,att}$ (“RN”) with our object representation $x_{obj,att}$ (“Obj”), or concatenate $x_{img,att}$ and $x_{obj,att}$ (“RN+Obj”). Then we train the SF model or the SM model using this object representation, with results shown in Table 2.³

For SF (lines 1–4), object representations substantially improve generalization ability: using either the object representation (“Obj”) or the combined representation (“RN+Obj”) obtains higher success rate on unseen environments than using only the ResNet features (“RN”), and the combined representation (“RN+Obj”) obtains the highest overall performance. For SM (lines 5–8),

³For each model and setting, we use the best training mechanism as found in Table 1, where student-forcing is used in all experiments except line 1 (where teacher-forcing is used to obtain the best performance for the non-visual agent under the SF architecture). See supplemental for more details.

#	model arch.	vis. feat.	SR on val-seen	SR on val-unseen
1		no vis.	34.1	35.2
2	SF	RN	53.3	29.0
3		Obj	38.5	33.5
4		RN+Obj	47.8	39.8
5		no vis.	36.1	39.7
6	SM	RN	62.8	40.5
7		Obj	48.8	41.6
8		RN+Obj	59.2	39.5

Table 2: Success rate (SR) of agents with different visual inputs on the R2R dataset (“RN”: ResNet CNN, “Obj”: objects, “no vis.”: no visual representation). Models: Speaker-Follower (SF) (Fried et al., 2018b) and Self-Monitoring (SM) (Ma et al., 2019).

the model that uses only the object representation achieves the best performance (line 7). Here the success rates across the four settings are closer, and the improvement from object representation is smaller than for SF. However, in Sec. 5 we find that object representation can be combined with other inputs to further improve the performance.

5 Mixture-of-experts makes better use of all available information

While the agent with CNN visual features does not outperform its non-visual counterpart (Sec. 3) on average, it often succeeds on individual instructions where the non-visual model fails, indicating the visual and non-visual modalities are complementary. To encourage grounding into both modalities, we ensemble visual and non-visual models in a mixture-of-experts approach.

5.1 Separate Training

We first ensemble the models from Sec. 3 and Sec. 4 at test time (after training them separately) by combining their predictions at each timestep.⁴

Lines 9–22 of Table 3 show ensembles of two models. Compared to single-model performance (line 1–8 in Table 2), an ensemble of a visual and a non-visual agent outperforms the individual agents for both the SF and the SM models. The best performing setting is the combination of “RN” and “no vis.” (non-visual) in line 20 under the SM model. While it is unsurprising that the mixture-of-experts can boost performance, it is interesting to see that the best mixture in line 20 outperforms mixtures of two agents of the same type (two non-visual agents in line 16, two visual agents in line 17, trained from distinct random parameter initializations), confirming that two agents

⁴We combine model predictions at each timestep by averaging action logits across models, which in early experiments slightly outperformed averaging action probabilities.

#	model arch.	mix.-of-exp. comb.	SR on val-unseen
9		(no vis., no vis.)	35.1
10		(RN, RN)	32.1
11	SF	(Obj, Obj)	35.4
12	(mixture of 2 models)	(RN+Obj, RN+Obj)	43.3
13		(RN, no vis.)	39.5
14		(Obj, no vis.)	38.4
15		(RN+Obj, no vis.)	43.1
16		(no vis., no vis.)	41.0
17		(RN, RN)	43.5
18	SM	(Obj, Obj)	45.2
19	(mixture of 2 models)	(RN+Obj, RN+Obj)	42.2
20		(RN, no vis.)	46.9
21		(Obj, no vis.)	43.4
22		(RN+Obj, no vis.)	46.4
23	SM (3-way mix.)	(RN, Obj, no vis.)	49.5
24	SM	(RN, no vis.)	48.3
25	(joint training)	(RN, Obj, no vis.)	51.9

Table 3: Success rate (SR) of different mixture-of-experts ensembles. Models: Speaker-Follower (SF) (Fried et al., 2018b) and Self-Monitoring (SM) (Ma et al., 2019); “RN”: ResNet CNN, “Obj”: objects, “no vis.”: no visual representation.

with access to different modalities can complement each other, especially in the SM model.

We also experiment with a 3-way mixture in the SM model, combining a visual agent with ResNet CNN features, a visual agent with object features, and a non-visual agent (line 23). This mixture outperforms all the 2-way mixtures by a noticeable margin, showing that the CNN and object-based visual representations are also complementary.

5.2 Joint Training

Finally, given the success of this simple test-time ensemble, we also explore *jointly training* these models by building a *single agent* which uses a single instruction encoder shared between multiple (visual and non-visual) jointly-trained decoders. During joint training, each decoder is supervised to predict the true actions, applying the same loss function as in separate training. During testing, actions are predicted by averaging logits from the separate decoders as in Sec. 5.1. We experiment with jointly training the agents in each of the two best-performing combinations (RN, no vis.) and (RN, Obj, no vis.) under the SM architecture (line 24 and 25 of Table 3). From line 24 vs. 20 and line 25 vs. 23, joint training gives higher performance than training each model separately and combining them only at test time. Overall, we obtain 51.9% final success rate on the val-unseen split (line 25), which is over 10% (absolute) higher than the SF or SM baselines using a single decoder with CNN features (line 2 and 6 in Table 2).⁵

⁵On the val-seen split, we maintain performance comparable to the original SM model.

6 Discussion

The success of non-visual versions of two recent state-of-the-art VLN models, often outperforming their vision-based counterparts in unseen environments on the benchmark R2R dataset, shows that these models do not use the visual inputs in a generalizable way. Our intuition is that, while language has rich, high-level symbolic meaning, which can be easily matched to the modality of the route structures, pixel-based visual representations, even those extracted via CNNs, are a lower-level modality which require more data to learn, and so a model trained on both modalities may learn to mostly rely on the route structure. This is also supported by the results in Table 3 (line 23 vs. line 20), where adding higher-level object representations improves the success rate by 2.6%.

Notably, an agent in the R2R environment is only able to move to a discrete set of locations in the environment, and at each point in time it only has a small number of actions available, determined by the environment’s connectivity graph (i.e., moving to the adjacent locations). These constraints on possible routes help explain our findings that language in the VLN instructions often grounds into geometric route structure in addition to visual context along the route. For example, if an instruction says *turn left at the couch*, and the route structure only allows the agent to turn left at a single location, it may not need to perceive the couch. Other instructions, such as *go straight for 5 meters and stop* may also be carried out without access to visual perception.

The improvement of our mixture-of-experts approach over single models suggests that it is challenging to learn to ground language into multiple modalities in one model. The “RN+Obj” model (Table 2, line 8) has access to *the same* information as our best result in Table 3, line 25, but obtains much lower success rate (39.5% vs. 51.9%). Thus, splitting the prediction task across several models, where each has access to a different input modality, is an effective way to inject an inductive bias that encourages the model to ground into each of the modalities.

Acknowledgements. This work was partially supported by Berkeley AI Research, the NSF and DARPA XAI. DF is supported by a Tencent AI Lab Fellowship.

References

- Ankesh Anand, Eugene Belilovsky, Kyle Kastnerand, Hugo Larochelle, and Aaron Courville. 2018. Blindfold baselines for embodied qa. *arXiv preprint arXiv:1811.05013*.
- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. 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)*.
- Jacob Andreas and Dan Klein. 2015. Alignment-based compositional semantics for instruction following. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Yoav Artzi and Luke Zettlemoyer. 2013. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association for Computational Linguistics*, 1(1):49–62.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niebner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. 2017. Matterport3d: Learning from rgb-d data in indoor environments. In *2017 International Conference on 3D Vision (3DV)*, pages 667–676. IEEE.
- David L. Chen and Raymond J. Mooney. 2011. Learning to interpret natural language navigation instructions from observations. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*.
- Howard Chen, Alane Shur, Dipendra Misra, Noah Snaveley, and Yoav Artzi. 2019. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Volkan Cirik, Yuan Zhang, and Jason Baldridge. 2018. Following formulaic map instructions in a street simulation environment. In *Visually Grounded Interaction and Language (ViGIL) Workshop, NeurIPS*.
- Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. 2018. Embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jacob Devlin, Hao Cheng, Hao Fang, Saurabh Gupta, Li Deng, Xiaodong He, Geoffrey Zweig, and Margaret Mitchell. 2015. Language models for image captioning: The quirks and what works. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 2: Short Papers*.
- Daniel Fried, Jacob Andreas, and Dan Klein. 2018a. Unified pragmatic models for generating and following instructions. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. 2018b. Speaker-follower models for vision-and-language navigation. In *Advances in Neural Information Processing Systems (NIPS)*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. 2018. Women also snowboard: Overcoming bias in captioning models. In *European Conference on Computer Vision*, pages 793–811. Springer.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73.
- Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan Al-Regib, Zsolt Kira, Richard Socher, and Caiming Xiong. 2019. Self-monitoring navigation agent via auxiliary progress estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Matt MacMahon, Brian Stankiewicz, and Benjamin Kuipers. 2006. Walk the talk: Connecting language, knowledge, and action in route instructions. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*.
- Hongyuan Mei, Mohit Bansal, and Matthew Walter. 2016. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*.

- Piotr Mirowski, Matt Grimes, Mateusz Malinowski, Karl Moritz Hermann, Keith Anderson, Denis Teplyashin, Karen Simonyan, Andrew Zisserman, Raia Hadsell, et al. 2018. Learning to navigate in cities without a map. In *Advances in Neural Information Processing Systems (NIPS)*.
- Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, et al. 2017. Learning to navigate in complex environments. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Dipendra Misra, Andrew Bennett, Valts Blukis, Eyvind Niklasson, Max Shatkhin, and Yoav Artzi. 2018. Mapping instructions to actions in 3d environments with visual goal prediction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2667–2678.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R Walter, Ashis Gopal Banerjee, Seth J Teller, and Nicholas Roy. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *AAAI*, volume 1, page 2.
- Jesse Thomason, Daniel Gordon, and Yonatan Bisk. 2019. Shifting the baseline: Single modality performance on visual navigation & qa. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Adam Vogel and Dan Jurafsky. 2010. Learning to follow navigational directions. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 806–814. Association for Computational Linguistics.
- Xin Wang, Wenhan Xiong, Hongmin Wang, and William Yang Wang. 2018. Look before you leap: Bridging model-free and model-based reinforcement learning for planned-ahead vision-and-language navigation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 37–53.
- Wei Yang, Xiaolong Wang, Ali Farhadi, Abhinav Gupta, and Roozbeh Mottaghi. 2019. Visual semantic navigation using scene priors. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. 2017. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 3357–3364. IEEE.