
Object-oriented Targets for Visual Navigation using Rich Semantic Representations

Jean-Benoit Delbrouck
Numediart
University of Mons, Belgium
jean-benoit.delbrouck@umons.ac.be

Stéphane Dupont
Numediart
University of Mons, Belgium
stephane.dupont@umons.ac.be

Abstract

When searching for an object humans navigate through a scene using semantic information and spatial relationships. We look for an object using our knowledge of its attributes and relationships with other objects to infer the probable location. In this paper, we propose to tackle the visual navigation problem using rich semantic representations of the observed scene and object-oriented targets to train an agent. We show that both allows the agent to generalize to new targets and unseen scene in a short amount of training time.

1 Introduction

In this paper, we focus on the problem of navigating a space to find and reach a given object using visual input as well as rich semantic information about the observations. Given that the agent has a semantic knowledge of the world, the attributes defining the objects he sees are grounded into the environment. We would like the agent to use its capacity to understand the attributes and relationships between objects to find the quickest pathway towards a given target. We focus on two training aspects to achieve this task

Firstly, the agent is pre-trained to describe all aspects (object-wise) of a scene using natural language. This ability relies on a strong semantic understanding of a visual scene and guarantees a strong grasp of the environment. He is also able to output the localization of the described objects within the visual frame as well as the confidence about the given inference. We hope this strong ability gives the agent the semantic and spatial knowledge required for the task. Indeed, objects with similar utility or characteristics are usually close to each other (i.e. when looking for a stove, you might infer that it stands on cooking plates, as well as a keyboard lies close to a computer).

In practice, the previous statement doesn't always hold true. Some spatial relationships between objects are specific to certain space configurations. For example, in a living-room, a TV is usually in front of (and therefore close by) of sofa, even though the two objects don't share characteristics. But this spatial relation between these two objects is not especially true for, let's say, a bedroom. To tackle this second problem, our model is built in two distinct layers: one is dedicated to every scene (global knowledge about the world) and the other is scene-type specific (namely bedroom, bathroom, living-room or kitchen). No model layer is dedicated to a specific scene (an instance of a bedroom, for example) as opposed to prior work [12], where a part of the model is reserved for each specific instance. In the latter contribution, the model tended to poorly transfer knowledge to unseen scene or objects.

For our experiments, we use the AI2-THOR framework [7] which provides near photo-realistic 3D scene of house's room. The framework features enables the setup to follow a reinforcement setting. The state of the agent in the environment changes as the agent take actions. The location of the agent is known at each step and can be randomized for training or testing. We dispose of 20 rooms each containing 5 targets (or frame) distributed along the four scene-types: kitchen, living-room, bedroom and bathroom.

2 Related work

Since the emergence of virtual environments, a few works are related to ours. Zhu et al. [12] proposed a deep reinforcement learning framework for target-driven visual navigation. As input to the model, only the target frame is given without any other priors about the environment. To transfer knowledge to unseen scene or object, a small amount of fine-tuning was required. We use their siamese network as baseline for our experiments. To adapt to unseen data, Gupta et al. [4] proposed a mapper and planner with semantic spatial memory to output navigation actions. The use of spatial memory for visual question answering in such interactive environments has been explored by [3]. Visually grounded navigation instruction have also been addressed by a plurality of works [1, 11, 5]. The closest work related to ours is probably from Yang et al. [10], they also use semantic priors for the virtual navigation task. Nevertheless, we differ in how the semantic representation is constructed. Their approach uses a knowledge graph built on Visual Genome [8] where the input of each node is based on the the current state (or current visual frame) and the word vector embedding of the target. Our idea is rather based on natural language as well as object localization (detailed in section 3.3). The task goal is also defined differently: their end state is reached when the target object is in the field of view and within a threshold of distance whilst we require the agent to go to a specific location.

3 Architecture

The following sections are structured as follows. First, we describe the main network architecture (§3.1), which is a siamese network (SN) used as a test-bed for our different experiments. The SN model is trained with random frames and then object-oriented frames (as explained in §3.2) to set up two baselines. We improve the SN architecture with an semantic component (§3.3). We explain how the semantic data are extracted from the agent observations (§3.3.1), and how we plug the component into our the baseline SN model (§3.3.2).

3.1 Network Architecture




Figure 1: SN model. The right grey box can be seen as the policy layer, there is one for every scene-type. It means that each room instance will share its policy parameters with other instances of the same scene-type.

This section describes the siamese network (SN) model. First, four history frames and the target frames (which is tiled 4 times to match history frames input size) are passed through a ResNet-50 pre-trained on ImageNet [2]. Both inputs (history and targets) produces 2048-d features for each frame. We freeze the ResNet parameters during training. The resulting history and target 8196-d output vectors is linearly projected into the 512-d embedding space by the same matrix W_1 . The CONCAT layer takes a 1024-d concatenated embedding and generates a 512-d joint representation with matrix W_2 that serves as input for each scene-type layer. This vector is subsequently passed through scene-type specific fully-connected layers $f_s, s \in \{bathroom, bedroom, kitchen, livingroom\}$, producing 4 policy outputs (actions) and a single value output. Matrices sizes of the models are $W_1 \in \mathbb{R}^{8196 \times 512}$, $W_2 \in \mathbb{R}^{1024 \times 512}$ and each f_s contains two matrices $W_{s_1} \in \mathbb{R}^{512 \times 512}$ and $W_{s_2} \in \mathbb{R}^{512 \times 5}$. All matrices have their subsequent bias and use ReLu as activation function.

As the shared layers parameters are used across different room types, it can benefit from learning with multiple goals simultaneously. We then use A3C reinforcement learning model [9] that learns by running multiple copies of training threads in parallel and updates the shared set of model parameters

in an asynchronous manner. The reward upon completion is set to 10. Each action taken decreases the reward by 0.01 hence favoring shorter paths to the target.

3.2 Object-oriented targets

We use the same data-set as used in Zhu et al. [12] for comparison. It’s composed of 20 rooms distributed along the four scene-types and for each room, five targets are chosen randomly. Our main assumption is that the use of semantic knowledge about objects improves the navigation task. As a first test-bed, we redefine the targets manually so that each target frame clearly contains an object. We make sure that an object picked appears on multiple other targets. To test this change, we make use of the model described in section 3.1.




Figure 2: Handpicked differences between randomly sampled target (top) and object-oriented target (bottom) across 4 scenes.

Observations 1, 2, 3 and 4 are part of the random dataset. They are very hard for the agent to locate. Indeed, frame one and two are mostly walls and have no discriminating component. The frame number three is loaded with too many information while frame number four is the view through a window. Observations 5, 6, 7 and 8 are the corresponding frames (from the same room) in the object-oriented dataset.

Results in section 4 will show that, after a certain amount of training, the SN model is able to learn with a moderate success to navigate to random targets. However, the training convergence (or the maximization of the reward) is much quicker when the model uses object-oriented targets. Moreover, providing object-oriented targets also improves the overall training so that it enables the agent to better generalize to other targets (random or object-oriented) and rooms. Both results are encouraging observations towards our intuition that the agent uses the objects and their relationship to navigate inside a room.

3.3 Semantic architecture

In this section, we explain how we collected data for our semantic knowledge database. We also show how we plug the semantics into the SN model presented in section 3.1.

3.3.1 Semantic knowledge

To build the semantics data of the observations, we use DenseCap [6], a vision system to both localize and describe salient regions in images in natural language. DenseCap is trained on Visual Genome [8], an image data-set that contains dense annotations of objects, attributes, and relationships within each image. Every frame available in the 20 scenes are passed through the trained DenseCap model to generate our semantic knowledge database. An output for a frame is as follows:

For each anchor box, the model predicts a confidence score and four scalars ($[x_{min}, y_{min}, x_{max}, y_{max}]$) which represent the predicted box coordinates. DenseCap also includes a language model to generate a rich sentence for each box. Regions with high confidence scores are more likely to correspond to ground-truth regions of interest, hence we keep 5 entries per frame having the top-5 confidence score. Even though DenseCap already has a language

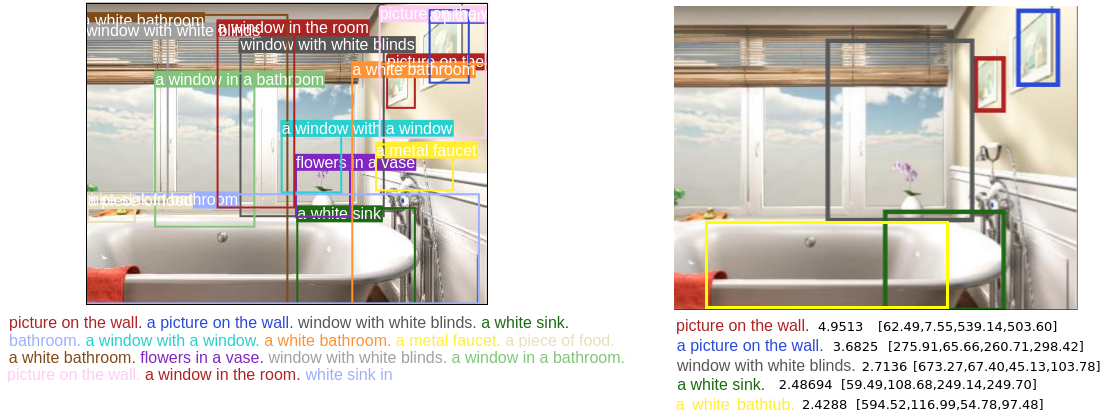



Figure 3: Output of DenseCap. We select the top-5 anchor boxes according to their confidence score.

model, we train an auto-encoder¹ with all the selected sentences amongst all scenes as training set. It allows us to perform an in-domain training (4915 sentences with a 477 words vocabulary) and to define ourselves the feature-size for one sentence. Once the training has converged, we extract a 64-d vector per sentence. A semantic knowledge of a single frame is thus a concatenation of five vectors (one per anchor box) of 69 dimensions: 64-d for the sentence, 4-d for the anchor coordinates and one scalar for the confidence score.

3.3.2 Semantic model



The semantic model (SSN) has two new inputs as shown in the figure on the left: the target semantics vectors and the history semantics. The target semantics is the semantic vector (of size $69 \times 5 = 345$) about the target frame and the history semantics is the concatenation of the semantic vectors of the past history frames (up to four). A new W_1' matrix encodes the semantic inputs and is of size 1380×512 . The two visual inputs (visual history frames and visual targets) are similar to the SN model presented in section 3.1.

4 Results

We now detail our results in two parts: first we illustrate the convergence of the reward from the Siamese Network (SN) being trained on random targets and the object-oriented targets from section 3.2. Finally, we describe the results of the semantic component compared with other models.

SN baselines Each graph of Figure 4 contains two training curves (or reward evolution) of a target. The orange curve depicts an instance of the SN model trained on the object-oriented targets, and in blue, a instance of the SN model trained on random targets. The number above each graph refers to the target observation presented in Figure 2.

As we see, the blue instance didn't converge for the random target #1 and #2. Indeed, both target are object-less: an empty corridor and a wall. The orange instance for target #5 and #6 has successfully converged. To compare their overall generalization ability, both blue and orange instances have been trained on the observation #3 and #4 (that are random). We notice that being trained on

¹<https://github.com/erickrf/autoencoder>

object-oriented target benefits the overall training: the orange instance converges quicker for random frames #3 and #4.




Figure 4: Evolution of the reward for the target observations ranging from 0 to 50 millions training frames.

Semantic We evaluate the semantic model performances on two sub-tasks:

- T1: generalization across targets (the scene is already seen, but not the targets)
- T2: generalization across scene (the scene has never been seen)

We compare the semantic model (SSN) with the Siamese Network model (SN) model and previous work [12]. We consider an episode successful if the target is reached in less than 1000 actions. The experiments are conducted with 20 scenes (or rooms) instances with 5 object-oriented targets per instance, so 100 targets total. After training, each target is evaluated for 100 episodes and then averaged.

For task T1, we train the models for 5 millions frames shared across all scenes instances and then evaluate on unseen targets in one instance of each scene-type. We report accuracy for each scene type.

For task T2, the models are trained for 10 millions frames shared along all scenes instances except one instance per scene-type. The latter are used for evaluation.

Finally, we also evaluate our SSN model trained on top-semantic targets. Per instance, we redefine the five targets with the 5 frames that have the highest confidence scores from DenseCap (within each instance). A target to reach is now a observation the agent has the best semantic information about. We call this experiment SSN_S.

T1	Bedroom		Bathroom		Kitchen		Living	
	E.L.	%	E.L.	%	E.L.	%	E.L.	%
R	963	6.6	941	9.4	988	6.4	968	4.8
[12]	903	19.9	691	71.4	885	23.9	921	16.2
SN	893	22	611	78.2	890	22.1	901	18.1
SSN	699	61.9	358	92.4	870	31.5	906	18.1
SSN_S	619	66.2	361	94.1	824	32.7	882	22.1

T2	Bedroom		Bathroom		Living	
	E.L.	%	E.L.	%	E.L.	%
R	941	8.6	928	10.7	991	2.8
SN	921	9.6	861	12.4	966	4.1
SSN	901	10.4	811	17.4	982	4.1
SSN_S	861	14.5	798	19.1	954	4.7

Table 1: Score table. R = Random, SN = Siamese Setwork model, SSN = semantic model, SSN_S = semantic model trained on top-semantic targets, E.L. = mean episode length (rounded), % = episode success rate.

5 Conclusion

For task 1, we see that the baseline Siamese Network (SN) using scene-type policies performs better overall than a scene-specific policy [12]. Also, our semantic model (SSN) is able to generalize to new targets within trained scene with a low amount of training frames. Importantly, training the semantic model on the top-semantic target (SSN_N) instead of object-oriented target (SSN) offers an even better generalization to new targets. Task 2 is harder but nonetheless, the semantic model still comes ahead. It is possible that more training frames are required to transfer knowledge to new scenes.

6 Acknowledgements

This work was partly supported by the Chist-Era project IGLU with contribution from the Belgian Fonds de la Recherche Scientifique (FNRS), contract no. R.50.11.15.F, and by the FSO project VCYCLE with contribution from the Belgian Walloon Region, contract no. 1510501.

References

- [1] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [2] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [3] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [4] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [5] Karl Moritz Hermann, Felix Hill, Simon Green, Fumin Wang, Ryan Faulkner, Hubert Soyer, David Szepesvari, Wojciech Marian Czarnecki, Max Jaderberg, Denis Teplyashin, et al. Grounded language learning in a simulated 3d world. *arXiv preprint arXiv:1706.06551*, 2017.
- [6] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [7] Eric Kolve, Roozbeh Mottaghi, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: An Interactive 3D Environment for Visual AI. *arXiv*, 2017.
- [8] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *Int. J. Comput. Vision*, 123(1):32–73, May 2017. ISSN 0920-5691. doi: 10.1007/s11263-016-0981-7. URL <https://doi.org/10.1007/s11263-016-0981-7>.
- [9] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1928–1937, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v48/mniha16.html>.
- [10] Wei Yang, Xiaolong Wang, Ali Farhadi, Abhinav Gupta, and Roozbeh Mottaghi. Visual semantic navigation using scene priors. *arXiv preprint arXiv:1810.06543*, 2018.

- [11] Haonan Yu, Haichao Zhang, and Wei Xu. Interactive grounded language acquisition and generalization in a 2d world. In *ICLR*, June 2018.
- [12] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. 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, 2017.