

# Visually-Grounded Planning without Vision: Language Models Infer Detailed Plans from High-level Instructions

Peter A. Jansen

School of Information, University of Arizona, Tucson, AZ  
pajansen@email.arizona.edu

## Abstract

The recently proposed ALFRED challenge task aims for a virtual robotic agent to complete complex multi-step everyday tasks in a virtual home environment from high-level natural language directives, such as “put a hot piece of bread on a plate”. Currently, the best-performing models are able to complete less than 5% of these tasks successfully. In this work we focus on modeling the translation problem of converting natural language directives into detailed multi-step sequences of actions that accomplish those goals in the virtual environment. We empirically demonstrate that it is possible to generate gold multi-step plans from language directives alone without any visual input in 26% of unseen cases. When a small amount of visual information is incorporated, namely the starting location in the virtual environment, our best-performing GPT-2 model successfully generates gold command sequences in 58% of cases. Our results suggest that contextualized language models may provide strong visual semantic planning modules for grounded virtual agents.

## 1 Introduction

Simulated virtual environments with steadily increasing fidelity are allowing virtual agents to learn to perform high-level tasks that couple language understanding, visual planning, and embodied reasoning through sensorimotor grounded representations (Gordon et al., 2018; Puig et al., 2018; Wijmans et al., 2019). The ALFRED challenge task recently proposed by Shridhar et al. (2020) requires a virtual robotic agent to complete everyday tasks (such as “put cold apple slices on the table”) in one of 120 interactive virtual home environments by generating and executing complex visually-grounded semantic plans that involve movable objects, irreversible state changes, and an egocentric viewpoint. Integrating natural language task directives with one



Figure 1: An example of the ALFRED grounded language task. In this work, we focus on visual semantic planning – from the textual directive alone (top), our model predicts a visual semantic plan of  $\{command, argument\}$  tuples (captions) that matches the gold plan without requiring visual input (images).

of the most complex interactive virtual agent environments to date is challenging, with the current best performing systems successfully completing less than 5% of ALFRED tasks in unseen environments<sup>1</sup>, while common baseline models generally complete less than 1% of tasks successfully.

In this work we explore the visual semantic planning task in ALFRED, where the high-level natural language task directive is converted into a detailed sequence of actions in the AI2-THOR 2.0 virtual environment (Kolve et al., 2017) that will accomplish that goal (see Figure 1). In contrast to previous approaches to visual semantic planning (e.g. Zhu et al., 2017; Fried et al., 2018; Fang et al., 2019), we explore the performance limits of this task solely using goals expressed in natural language as input – that is, *without* visual input from the virtual environment. The contributions of this

<sup>1</sup><https://leaderboard.allenai.org/alfred/>

work are:

1. We model visual semantic planning as a sequence-to-sequence translation problem, and demonstrate that our best-performing GPT-2 model can translate between natural language directives and sequences of gold visual semantic plans in 26% of cases without visual input.
2. We show that when a small amount of visual input is available – namely, the starting location in the virtual environment – our best model can successfully predict 58% of unseen visual semantic plans.
3. Our detailed error analysis suggests that repairing predicted plans with correct locations and fixing artifacts in the ALFRED dataset could substantially increase performance of this and future models.

## 2 Related Work

Models for completing multi-modal tasks can achieve surprising performance using information from only a single modality. The Room-to-Room (R2R) visual language navigation task (Anderson et al., 2018) requires agents to traverse a discrete scene graph and arrive at a destination described using natural language. In ablation studies, Thomason et al. (2019) found that models using input from a single modality (either vision or language) often performed nearly as good as or better than their multi-modal counterparts on R2R and other visual QA tasks. Similarly, Hu et al. (2019) found that two state-of-the-art multi-modal agents performed significantly worse on R2R when using both linguistic and visual input instead of a single modality, while also showing that performance can improve by combining separate-modality models into mixture-of-expert ensembles.

Where R2R requires traversing a static scene graph using locomotive actions, ALFRED is a dynamic environment requiring object interaction for task completion, and has a substantially richer action sequence space that includes 8 high-level actions. This work extends these past comparisons of unimodal vs. multimodal performance by demonstrating that strong performance on visual semantic planning is possible in a vastly more complex virtual environment using language input alone, through the use of generative language models.

## 3 Models and Embeddings

We approach the task of converting a natural language directive into a *visual semantic plan* – a series of commands that achieve that directive in a virtual environment – as a purely textual sequence-to-sequence translation problem, similar to conversion from Text-to-SQL (e.g. Yu et al., 2018; Guo et al., 2019). Here we examine two embedding methods that encode language directives and decode command sequences.

**RNN:** A baseline encoder-decoder network for sequence-to-sequence translation tasks (e.g. Bahdanau et al., 2015), implemented using recurrent neural networks (RNNs). One RNN serves as an encoder for the input sequence, here the tokens representing the natural language directive. A decoder RNN network with attention uses the context vector of the encoder network to translate into output sequences of command triples representing the visual semantic plan. Both encoder and decoder networks are pre-initialized with 300-dimensional GLoVe embeddings (Pennington et al., 2014).

**GPT-2:** The OpenAI GPT-2 transformer model (Radford et al., 2019), used in a text generation capacity. We fine-tune the model on sequences of natural language directives paired with gold command sequences separated by delimiters (i.e. “<Directive> [SEP] <CommandTuple<sub>1</sub>> [CSEP] <CommandTuple<sub>2</sub>> [CSEP] ... [CSEP] <CommandTuple<sub>N</sub>> [EOS]”). During evaluation we provide the prompt “<Directive> [SEP]”, and the model generates a command sequence until producing the end-of-sequence (EOS) marker. We make use of nucleus sampling (Holtzman et al., 2020) to select only tokens from the set of most likely tokens during generation, with  $p = 0.9$ , but do not make use of top-K filtering (Fan et al., 2018) or penalize repetitive n-grams, which are commonly used in text generation tasks, but are inappropriate here for converting to the often repetitive (at the scale of bigrams) command sequences. For tractability we make use of the GPT-2 Medium pre-trained model, which contains 24 layers, 16 attention heads, and 325M parameters. During evaluation, task directives are sorted into same-length batches to prevent generation artifacts from padding, and maintain high generation quality.<sup>2</sup>

<sup>2</sup>Negative results not reported for space: We hypothesized that separating visual semantic plans into variablized action-sequence templates and variable-value assignments represented as separate decoders would help models learn to

Model	Triple Components			Full Triples	Entire Visual Semantic Plans	
	Command	Arg1	Arg2		Full Sequence	Full Minus First
<i>Strict Scoring</i>						
RNN	89.6%	64.8%	58.4%	60.2%	17.1%	43.6%
GPT-2	90.8%	69.9%	63.8%	65.8%	22.2%	53.4%
<i>Permissive Scoring</i>						
RNN	89.6%	70.6%	61.4%	65.9%	23.6%	26.1%
GPT-2	90.8%	73.8%	65.1%	69.4%	26.1%	58.2%

Table 1: Average prediction accuracy on the unseen test set broken down by triple components, full triples, and full visual semantic plans. *Full Sequence* accuracy represents the proportion of predicted visual semantic plans that perfectly match gold plans. *Full Minus First* represents the same, but omitting the first tuple, typically a  $\{goto, location\}$  that moves the agent to the starting location in the virtual environment (see description in text).

Model	Goto	Pickup	Put	Cool	Heat	Clean	Slice	Toggle	Avg.
RNN	59	81	60	<b>77</b>	69	<b>83</b>	67	91	66
GPT-2	<b>63</b>	<b>84</b>	<b>66</b>	72	<b>77</b>	82	<b>70</b>	<b>94</b>	<b>69</b>

Table 2: Average triple prediction accuracy on the test set broken down into each of the 8 possible ALFRED commands. Values represent percentages. *Goto* has an  $N$  of 24k, *Pick up* an  $N$  of 11k, and *Put* an  $N$  of 10k. All other commands occur approximately 1000 times in the test dataset.

## 4 Experiments

**Dataset:** The ALFRED dataset contains 6,574 gold command sequences representing visual semantic plans, each paired with 3 natural language directives describing the goal of those command sequences (e.g. “*put a cold slice of lettuce on the table*”) authored by mechanical turkers. High-level command sequences range from 3 to 20 commands (average 7.5), and are divided into 7 high-level categories (such as *examine object in light*, *pick two objects then place*, and *pick then cool then place*). Commands are represented as triples that pair one of 8 actions (*goto*, *pickup*, *put*, *cool*, *heat*, *clean*, *slice*, and *toggle*) with up to two arguments, typically the object of the action (such as “*slicing lettuce*”) and an optional receptacle (such as “*putting a spoon in a mug*”). Arguments can reference 58 possible objects (e.g. *butter knife*, *chair*, or *apple*) and 26 receptacles (e.g. *fridge*, *microwave*, or *bowl*). To prevent knowledge of the small unseen test set for the full task, here we redivide the large training set into three smaller train, development, and test sets of 7,793, 5,661, and 7,571 gold-

separate the general formula of action sequences with specific instances of objects in action sequences, which has been shown to help in Text-to-SQL translation (Guo et al., 2019). Pilot experiments with both RNNs and transformer models yielded slightly lower results than vanilla models. Language modeling: In addition to GPT-2 we also piloted XLNET, but perplexity remained high even after significant fine-tuning.

directive/command-sequence pairs, respectively.

**Processing Pipeline:** Command sequences are read in as sequences of  $\{command, arg1, arg2\}$  triples, converted into natural language using completion heuristics (e.g. “ $\{put, spoon, mug\}$ ”  $\rightarrow$  “*put the spoon in the mug*”, and augmented with argument delimiters to aid parsing (e.g. “*put <arg1> the spoon <arg2> in the mug*”). Input directives are tokenized, but receive no other pre-processing. Generated strings from all models are post-processed for common errors in sequence-to-sequence models, including token doubling, completing missing bigrams (e.g. “*pick <arg1>*”  $\rightarrow$  “*pick up <arg1>*”), and heuristics for adding missing argument tags. Post-processed output sequences are then parsed and converted back into  $\{command, arg1, arg2\}$  tuples for evaluation.

**Evaluation Metrics:** Performance in translating between natural language directives and sequences of command triples is evaluated in terms of accuracy at the command-element (*command*, *argument1*, *argument2*), triple, and full-sequence level. Because our generation includes only textual input and no visual input for a given virtual environment, commands may be generated that reference objects that do not exist in a scene (such as generating an action to toggle a “*lamp*” to examine an object, when the environment specifically contains a “*desk lamp*”). As such we include two scoring metrics: a *strict* metric that requires exact matching of each token in an argument to be counted as correct, and a *permissive* metric that requires matching only a single token within an argument to be correct.

<i>Strict Scoring</i>	butter knife $\neq$ knife
<i>Permissive Scoring</i>	desk lamp = lamp

All accuracy scoring is binary. Triples receive a score of one if all elements in a given gold and predicted triple are identical, and zero oth-

Prop.	Error Class Description	Example Errors
<i>Incorrect Arguments</i>		
45%	Predicted wrong location	<i>Predicted wrong location:</i> (G) ... slice lettuce, put knife on <b>countertop</b> , put lettuce in fridge, ...
4%	Predicted wrong object	(P) ... slice lettuce, put knife in <b>microwave</b> , put lettuce in fridge, ...
<i>Incorrect Triples</i>		
22%	Offset due to extra/missing actions	<i>Predicted extra (not harmful) action<sup>†</sup>, and introduced offset error<sup>‡</sup></i> Instructions: Put a mug with a spoon in the sink.
22%	Predicted extra (incorrect) actions	(G) ... pick up mug, <b>put mug in sink basin<sup>‡</sup></b>
12%	Predicted missed actions	(P) ... pick up mug, <b>go to sink basin<sup>†</sup>, put mug in sink basin<sup>‡</sup></b>
7%	Predicted extra (not harmful) actions	
5%	Order of actions swapped	
<i>Instruction Errors</i>		
17%	Gold Instructions Incorrect	<i>Gold Instructions Incomplete:</i> Instructions: Put a heated mug in the microwave.
13%	Gold Instructions Incomplete	(G) ... go to microwave, heat mug, <b>go to cabinet, put mug in cabinet</b>

Table 3: (left) Common classes of prediction errors in the GPT-2 model, and their proportions in 100 predictions from the development set. (right) Example errors, where (G) and (P) represent subsets of gold and predicted visual semantic plans, respectively.

erwise. Full-sequence scoring directly compares  $\langle \text{CommandTriple}_i \rangle$  for each  $i$  in the gold and predicted sequences, and receives a score of one only if all triples are identical and in identical locations  $i$ , and zero otherwise.<sup>3</sup>

#### 4.1 Results

Performance of the embedding models is reported in Table 1, broken down by triple components, full triples, and full sequences. Both models achieve approximately 90% accuracy in predicting the correct commands, in the correct location  $i$  in the sequence. Arguments are predicted less accurately, with the RNN model predicting 65% and 58% of first and second arguments correctly, respectively. The GPT-2 model increases performance on argument prediction by approximately +5%, reaching 70% and 64% under strict match scoring. Permissive scoring, allowing for partial matches between arguments (e.g. “lamp” and “desk lamp” are considered equivalent) further increases argument scoring to approximately 74% and 65% in the best model. Scoring by complete triples in the correct location  $i$  shows a similar pattern of performance, with the best-scoring GPT-2 model achieving 66% accuracy using strict scoring, and 69% under permissive scoring, with triple accuracy broken down by command shown in Table 2.

Fully-correct predicted sequences of commands that perfectly match gold visual semantic plans using only the text directives as input, – i.e. *without*

<sup>3</sup>Tuning and Computational Resources: RNN models required approximately 100k epochs of training to reach convergence over 12 hours, requiring 8GB of GPU RAM. GPT-2 models asymptoted performance at 25 epochs, requiring 6 hours of training and 16GB of GPU RAM. All experiments were conducted using an NVIDIA Titan RTX.

visual input from the virtual environment – occur in 17% of unseen test cases with the RNN model, and 22% of cases with the GPT-2 model, highlighting how detailed and accurate visual plans can be constructed from text input alone in a large subset of cases. In analyzing the visual semantic plans, the first command is typically to move the virtual agent to a starting location that contains the first object it must interact with (for example, moving to *the countertop*, where a potato is resting in the initialized virtual environment, to begin a directive about *slicing, washing, and heating a potato slice*). If we supply the model with this single piece of visual information from the environment, full-sequence prediction accuracy for all models more than doubles, increasing to 53% in the strict condition, and 58% with permissive scoring, for the best-performing GPT-2 model.

#### 4.2 Error Analysis

Table 3 shows an analysis of common categories of errors in 100 directive/visual semantic plan pairs randomly drawn from the development set that were not answered correctly by the best-performing GPT-2 model that includes the starting location for the first step. As expected, a primary source of error is the lack of visual input in generating the visual plans, with the most common error, *predicting the wrong location in an argument*, occurring in 45% of errors.<sup>4</sup> Conversely, predicting the wrong object to interact with occurred in only 4% of errors,

<sup>4</sup>An unexpected source of error is that our GPT-2 planner frequently prefers to store used cutlery in either the fridge or microwave – creating a moderate fire hazard. Interestingly, this behavior appears learned from the training data, which frequently stores cutlery in unusual locations. Disagreements on discarded cutlery locations occurred in 15% of all errors.

as this information is often implicitly or explicitly supplied in the text directive. This suggests augmenting the model with object locations from the environment could mend prediction errors in nearly half of all errorful plans.

The GPT-2 model predicted additional (incorrect) actions in 22% of errorful predictions, while missing key actions in 12% of errors, causing offset errors in sequence matching that reduced overall performance in nearly a quarter of cases. In a small number of cases, the model predicted extra actions that were not harmful to completing the goal, or switched the order of sets of actions that could be completed independently (such as picking up and moving two different objects to a single location). In both cases the virtual agent would likely have been successful in completing the directive if following these plans.

A final significant source of error includes inconsistencies in the crowdsourced text directives or gold visual semantic plans themselves. In 17% of errors, the gold task directive had a mismatch with the objects referenced in the gold commands (e.g. the directive referenced a *watering can*, where the gold annotation references a *tea pot*), and automated scoring marked the predicted sequence as incorrect. Similarly, in 13% of cases, the task directive failed to mention one or more subtasks (e.g. the directive is “*turn on a light*”, but the gold command sequence also includes first retrieving a specific object to examine in the light). This suggests that nearly one-third of errors may be due to issues in the evaluation data, and that overall visual semantic plan generation performance may be significantly higher.

## 5 Data Dependence and Few-Shot Learning

To examine how performance varies with the amount of training data available, we randomly downsampled the amount of training data to 25%, 10%, and 1% of its original size. This analysis, shown in Figure 2, demonstrates that relatively high performance on the visual semantic prediction task is still possible with comparatively little training data. When only 10% of the original training data is used, average prediction accuracy reduces by 24%, but still reaches 44%. In the *few-shot* case (1% downsampling), where each of the 7 ALFRED tasks observes only 4 gold command sequences each (for a total of 12 natural language directives

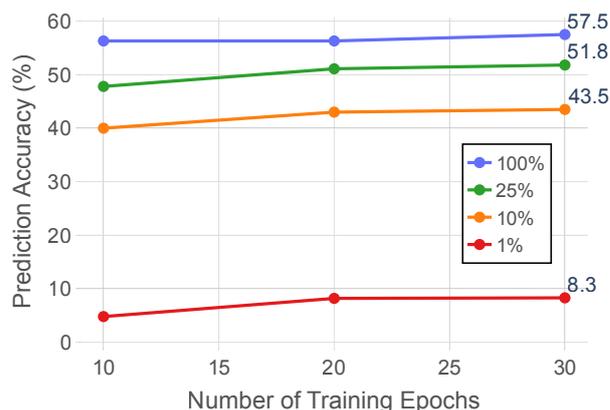


Figure 2: Average prediction accuracy as a function of training set size (100%, 25%, 10%, or 1% of the full training set) for the GPT-2 model on the test set. Even with a large reduction in training data, the model is still able to accurately predict a large number of visual semantic plans. Performance represents the permissive scoring metric in the “*full minus first*” condition in Table 1.

per task) during training, the GPT-2 model is still able to generate an accurate visual semantic plan in 8% of cases. Given that large pre-trained language models have been shown to encode a variety of commonsense knowledge *as-is*, without fine-tuning (Petroni et al., 2019), it is possible that some of the model’s few-shot performance on ALFRED may be due to an existing knowledge of similar common everyday tasks.

## 6 Conclusion

We empirically demonstrate that detailed gold visual semantic plans can be generated for 26% of unseen task directives in the ALFRED challenge using a large pre-trained language model *without* visual input from the simulated environment, where 58% can be generated if starting locations are known. We envision these plans may be used either *as-is*, or as an initial “hypothetical” plan of how the model believes the task might be solved in a generic environment, that is then modified based on visual or other input from a specific environment to further increase overall accuracy.

We release our planner code, data, predictions, and analyses for incorporation into end-to-end systems at: <http://github.com/cognitiveailab/alfred-gpt2/>.

## References

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*, pages 3674–3683.
- 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)*.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.
- Kuan Fang, Alexander Toshev, Li Fei-Fei, and Silvio Savarese. 2019. Scene memory transformer for embodied agents in long-horizon tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 538–547.
- Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. 2018. Speaker-follower models for vision-and-language navigation. In *Advances in Neural Information Processing Systems*, pages 3314–3325.
- Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. 2018. Iqa: Visual question answering in interactive environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4089–4098.
- Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. 2019. Towards complex text-to-sql in cross-domain database with intermediate representation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4524–4535.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. The curious case of neural text degeneration. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Ronghang Hu, Daniel Fried, Anna Rohrbach, Dan Klein, Trevor Darrell, and Kate Saenko. 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*, pages 6551–6557, Florence, Italy. Association for Computational Linguistics.
- Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. 2017. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473.
- Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. 2018. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8494–8502.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. [Alfred: A benchmark for interpreting grounded instructions for everyday tasks](#). In *Computer Vision and Pattern Recognition (CVPR)*.
- Jesse Thomason, Daniel Gordon, and Yonatan Bisk. 2019. [Shifting the baseline: Single modality performance on visual navigation & QA](#). 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)*, pages 1977–1983, Minneapolis, Minnesota. Association for Computational Linguistics.
- Erik Wijmans, Samyak Datta, Oleksandr Maksymets, Abhishek Das, Georgia Gkioxari, Stefan Lee, Irfan Essa, Devi Parikh, and Dhruv Batra. 2019. Embodied question answering in photorealistic environments with point cloud perception. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921.
- Yuke Zhu, Daniel Gordon, Eric Kolve, Dieter Fox, Li Fei-Fei, Abhinav Gupta, Roozbeh Mottaghi, and Ali Farhadi. 2017. Visual semantic planning using deep successor representations. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 483–492.