

VideoNavQA: Bridging the Gap between Visual and Embodied Question Answering

Cătălina Cangea
Catalina.Cangea@cst.cam.ac.uk
Eugene Belilovsky
eugene.belilovsky@umontreal.ca
Aaron Courville
aaron.courville@umontreal.ca

University of Cambridge
Cambridge, United Kingdom
Mila, Université de Montréal
Montréal, Canada
Mila, CIFAR Fellow
Montréal, Canada

Abstract

Embodied Question Answering (EQA) is a recently proposed task, where an agent is placed in a rich 3D environment and must act based solely on its egocentric input to answer a given question. The desired outcome is that the agent learns to combine capabilities such as scene understanding, navigation and language understanding in order to perform complex reasoning in the visual world. However, initial advancements combining standard vision and language methods with imitation and reinforcement learning algorithms have shown EQA might be too complex and challenging for these techniques. In order to investigate the feasibility of EQA-type tasks, we build the VideoNavQA dataset that contains pairs of questions and videos generated in the House3D environment. The goal of this dataset is to assess question-answering performance from nearly-ideal navigation paths, while considering a much more complete variety of questions than current instantiations of the EQA task. We investigate several models, adapted from popular VQA methods, on this new benchmark. This establishes an initial understanding of how well VQA-style methods can perform within this novel EQA paradigm.

1 Introduction

The Embodied Question Answering [6] (EQA) and Interactive Question Answering [10] (IQA) tasks were introduced as a means to study the capabilities of agents in rich, realistic environments, requiring both *navigation and reasoning* to achieve success. Each of these skills typically needs a different approach [12, 13, 21] that should nevertheless be smoothly integrated with the rest of the system leveraged by the agent. These abilities are assessed via placing the agent at a random location in a house environment and asking it a question—successful completion of the task involves the agent knowledgeably exploring the environment and reasoning about visual stimuli.

Initial attempts at solving the EQA task [6, 7] have thus combined typically used vision (convolutional neural networks for object detection) and language (question encoding or program generation) techniques with imitation learning and reinforcement learning. However, these



Figure 1: High-level description of the VideoNavQA task and our approach: the VQA system receives an example containing a video of a trajectory inside the house environment and a question, processes the input and produces an answer.

approaches either suffer from potentially weaker performance than when using a language-only model [11] or are preceded by additional hand-engineered steps (manually defined sub-goals, imitation learning on pre-computed expert trajectories). Indeed, even for simple single object questions, the agent is often unable to advance meaningfully towards the target [12], producing visual streams that cannot be used to answer the query. This suggests that EQA is a highly challenging task which might not be easily approached from this angle—further investigation is required to realistically assess the gap between current state-of-the-art and desired human level performance on the EQA performance.

Single-image Visual Question Answering is only now starting to tackle complex reasoning questions, even in limited settings [13, 16]—it is unclear if existing methods can handle the rich video streams produced in our task. On the other hand, EQA introduces a navigation component, which can often lead to video inputs which are uninformative with respect to the question. A natural issue arises: can the desired tasks be solved with current methods, if we assume the agent is given correct visual streams (ones that can be used to provide an answer)? In our attempt to answer this question, we propose a dataset that decouples the visual reasoning from the navigation aspects in the EQA problem. We introduce the VideoNavQA task, illustrated in Figure 1. While removing the navigation and action selection requirements from EQA, we increase the difficulty of the visual reasoning component via a much larger question space, tackling the sort of complex reasoning questions that make QA tasks challenging [16]. By designing and evaluating several VQA-style models on the dataset, we *establish a novel way of evaluating EQA feasibility* given existing methods, while highlighting the difficulty of the problem even in the most ideal setting.

2 Related work

To the best of our knowledge, there are no previous works that position themselves at the intersection of these two paradigms. Instead, VQA and EQA have been tackled from separate angles, with multiple interesting developments in each case.

Das *et al.* [9] proposed the EQA-v1 dataset containing 4 types of questions (`location`, `color`, `color_room`, `preposition`) that always refer to a single object, the navigation goal of the agent. They train a model using imitation and reinforcement learning, while revealing that RL finetuning often results in overshooting the goal. A subsequent improvement is achieved via hierarchical policy learning with neural modular control [17]—however, this

	Houses	Samples
Train	622	84990
Validation	65	8755
Test	56	7587

Table 1: Dataset split statistics. The three sets of environments are disjoint.

EQAv1 (Q types: 4)	What room is the <OBJ> located in? What color is the <OBJ> in the <ROOM>?
VideoNavQA (Q types: 28)	Are both <attr1> <OBJ1> and <attr2> <OBJ2> <color>? How many <attr> <OBJ> are in the <ROOM>? Is there <art> <attr> <OBJ>?

Table 2: Some examples of the type of question templates found in EQAv1 versus the more complex existence, counting and comparison questions found in VideoNavQA.

approach uses hand-crafted sub-policies. Anand *et al.* [10] study the performance of question-only baselines on EQA-v1, which give better results when the agent is spawned more than 10 steps away, concluding that existing EQA methods struggle to exploit (and are often impeded by) the environment. A more recent extension of EQA, multi-target EQA [39], requires reaching multiple goals to answer the question (for example, comparing the sizes of two objects in different rooms). The authors build the MT-EQA dataset containing 6 types of questions and tackle the task by first decomposing each question into small sub-goals, which are more easily achievable by a model similar to the one used in previous works. Another EQA variant based on photorealistic environments [66] uses point clouds instead of RGB input; the authors “port” three of the EQA-v1 questions to Matterport3D. Although in this work we focus on the House3D environment and EQA, we also note that the IQA task [10] for the AI2Thor [18] environment was introduced with similar goals. However, this task is defined in single room settings, hence the negative effect of poor navigation is less severe.

Visual question answering has been extensively tackled over the past few years, with numerous datasets being released, various algorithms designed and general studies carried out [4, 20]. The VQA dataset [4] is one such example, containing free-form and open-ended questions about real-life images and abstract scenes that require natural language answers. CLEVR [16] uses a functional program-based question representation to generate a vast range of questions from synthetic scene graphs that contain 3–10 objects with restricted variability (3 types, 2 sizes, 2 materials, 8 colors). More recently, the GQA dataset [14] has been proposed to address some of the issues of current datasets, including biases in the answer distribution. Widely used VQA models include *stacked attention networks* [58], which use the question embedding as a query for attending over the visual input, *multimodal compact bilinear pooling* [8], which fuses the text and visual embeddings via multiplying them in the Fourier space, *feature-wise linear modulation* [26], which selects informative visual feature maps via question-based conditioning, its *multi-hop* extension [52], which conditions feature maps by iteratively attending over the language input, and *compositional attention networks* (MAC) [13], which can reason about the visual input in a multi-step fashion using the memory-attention-control cell. Graph neural network-based approaches have also been used for VQA tasks lately [23, 25, 34], operating on the relational representation of objects in the image. The model that most closely resembles the ones we propose is explored in a multi-turn QA setting [24], where the system is provided with a dialogue (set of question-answer pairs) and a video—the question encoding is used to perform both per-frame conditioning and attention over the hidden states of the LSTM which encodes all the video frames.

Several datasets have been proposed that also consider **video question answering** in settings such as movies [19, 27, 33]. They are often provided with rich per-frame annotations (for example, subtitles in movies). However, these tasks largely focus on identifying actions or other dynamic behaviors and are distant from the aims of VideoNavQA. Our task considers




Figure 2: VideoNavQA dataset distribution. (Left:) Proportions for each question category. (Middle:) Question lengths (maximum is 56). (Right:) Video lengths (maximum is 140).

indoor navigation trajectories that exhibit rich visual data at each time step. This poses an additional challenge for video QA systems—these are now required to *isolate the relevant information from a large pool of visual concepts* present in each frame and *perform more advanced reasoning to answer the extensive variety of questions* designed.

3 The VideoNavQA dataset

We have constructed a new benchmark to study the capabilities of VQA approaches in a variant of the EQA task, where the agent is required to answer a question while having access to a near-optimal trajectory—that is, the trajectory taken corresponds to natural trajectories with sufficient information in the video input to answer the question. Our task can be seen as complementary to the Habitat Challenge [29], where the focus is on navigation instead of question answering. We use the House3D virtual environment [57] to generate approximately 101,000 pairs of videos and questions; the dataset contains 28 types of questions belonging to 8 categories (see Figure 2), with 70 possible answers. Each question type is associated with a template that facilitates programmatic generation using ground truth information extracted from the video. The complexity of the questions in the dataset is far beyond that of other similar tasks using this generation method (such as CLEVR): the questions involve single or multiple object/room existence, object/room counting, object color recognition and localization, spatial reasoning, object/room size comparison and equality of object attributes (color, room location). The full list of question types and dataset counts can be found in the supplementary material.¹

3.1 Visual information

Environments and scene representation The House3D environment is based on indoor scenes from the SUNCG dataset [51]. Table 1 shows the number of houses used in each of the three dataset splits, along with the corresponding number of examples (question-video pairs). Each house appears only in a single split and in no more than 150 videos, to ensure that there is a large and consistent variability in the visual information across the dataset.

Video generation and ground truth extraction Each example in the VideoNavQA dataset is constructed by first generating the video component. We use the underlying grid representation of the environments to compute the shortest path between arbitrary locations in

¹The dataset and code can be found at <https://github.com/catalinal17/VideoNavQA>.




Figure 3: (Left:) Concat-CNN3D processes the entire video. (Right:) Concat-CNN2D aggregates frame features via an LSTM. Both merge the result with the question embedding.

two different rooms of the house. To obtain the video, we render the shortest-path trajectory inside House3D. This corresponds to what an agent would see in the EQA setting while exploring the house, with the added benefit that its trajectory is already sensible from a navigation perspective. The ground truth information is then obtained by parsing each video frame. We use the SUNCG semantic rendering to identify the objects that are visible. By linking them via depth rendering to the current room the agent is in, or to an adjacent one, we index the objects and rooms that are seen on the trajectory and therefore in the video.

3.2 Questions

Functional form representation Similarly to other synthetic benchmark datasets (EQA-v1 [6], CLEVR [16]), we choose to generate questions according to functional, template-style representations (e.g. “How many $\langle attr \rangle$ $\langle obj_type-pl \rangle$ are in the $\langle room_type \rangle$?”). This facilitates the instantiation of ground truth tags, once the video for the corresponding trajectory has been generated and analyzed. Moreover, we can easily execute the corresponding program to determine the answer—this amounts to performing a series of basic operations such as `filter()`, `count()`, `get_attr()` on the ground truth.

Generation The question generation process starts by randomly choosing one of the 28 templates to be instantiated. A valid question will always have tags instantiated with ground truth values. For example, if there is a $\langle room_type \rangle$ tag and we have only seen a kitchen and a living room on our trajectory, then the set of possible instantiations is $\{kitchen, living\ room\}$. Using this principle, we build sets of possible values for each tag in the template. In order to generate a valid $(question, answer)$ pair, we randomly assign each tag a value from its set, then run the template functional program to compute whether the question is valid and can be answered using the ground truth. To illustrate the process, consider the template “What color is the $\langle attr \rangle$ $\langle obj_type \rangle$?” with the associated program:

```
input(objs) → filter(obj_type) → filter(attr) → unique() → get_attr(color)
```

We first select the set of objects seen on the trajectory from our ground truth representation. Next, we filter by the instantiated object type, then by the instantiated attribute (enforced to not be a color during the tag value assignment). Finally, we ensure that the result is unique (i.e. that the question is unambiguous) and retrieve the color of the object as the answer.

4 Models

We have designed the VideoNavQA dataset to address the EQA challenge from an alternative perspective, requiring a smaller degree of fusion among different classes of methods. In this section, we describe the architectures used to establish more realistic expectations on EQA performance and obtain initial results on VideoNavQA: several essential baselines and new models inspired by previous successes in visual question answering and computer vision.

Language models Question-only models have been surprisingly effective in EQA, often performing better than the complex approaches initially suggested [10]. We include two simple yet powerful language modelling approaches in our evaluation: a *1-layer LSTM* [10] and an *bag-of-words (BoW)* [27] model. Both of them contain an initial embedding layer where a separate representation is learned for each of the 134 vocabulary words. The BoW then averages all representations of the question words and feeds the result to a linear layer, while the LSTM encodes the entire sequence into a vector. These baselines demonstrate the inherent biases that exist in the environment distribution and place a lower bound on the desired performance of models that can usefully exploit visual information in this setting.

Vision models Learning to answer a question based only on the visual input is not expected to perform better than being biased towards the most frequently occurring answer (approximately 66% of the VideoNavQA questions require a binary answer). To illustrate the expected behavior, we use a per-frame processing VGG-style [30] convolutional neural network with an LSTM classifier (V-CNN2D) and a C3D-like [35] CNN (V-CNN3D).

Concatenation models Based on single-modality baseline results, we integrate the LSTM model with each of the two vision models, in order to derive a joint representation of the obtained features. We achieve this by concatenating the final question representation with their respective video representations (Concat-CNN2D/3D; see Figure 3), then passing them through an MLP and obtaining the most likely answer via softmax.

FiLM-based per-frame reasoning Feature-wise linear modulation (FiLM) [26] has previously achieved state-of-the-art results on visual question answering tasks, directly using the question embedding to scale and shift the feature maps of a CNN pipeline taking the image as input. Here, we *extend FiLM to address the additional VideoNavQA temporal dimension*. As per Figure 4, each frame is processed independently by a fixed number of ResBlocks [26]. FiLM conditioning ensures that visual information relevant for answering the question will be propagated to the final frame representation. We thereby obtain a series of visual features that we then aggregate via two mechanisms:

- *attention*: we apply a linear transformation to the features from each frame, then use recurrent attention [9, 5] to obtain an attended, final encoding. (FiLM AT)
- *global max-pooling*: we apply a 1×1 convolution, followed by a feature-wise max-pooling operation over all frames that produces the final representation. (FiLM GP)

The motivation for these two classifiers is related to having a rich representation of the video across time, from which the visual reasoning method needs to select the information that is required to answer the question, typically contained in only a few frames. Attention and global max-pooling are generally effective at picking the essential features from a sequence and thus are reasonable choices for our model design.

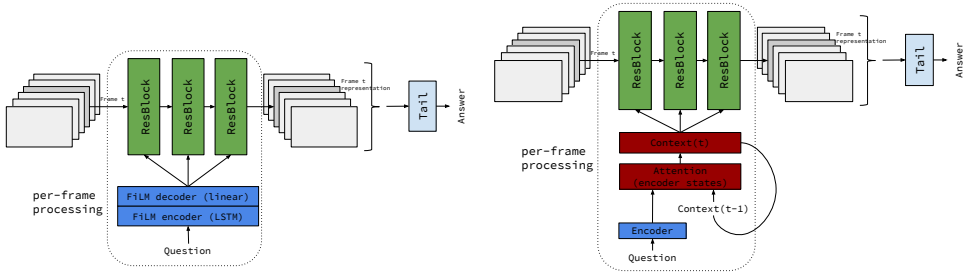


Figure 4: (Left:) Per-frame FiLM model. Video frames are processed separately by the ResBlocks, then all features are aggregated by the classifier to answer the question. (Right:) Temporal Multi-hop model. Each video frame is processed by the ResBlocks: the FiLM parameters are computed from the current attention context, which is initialized with the one from the previous frame. Temporal summarization is achieved via global max-pooling.

Multi-hop temporal processing The multi-hop extension of FiLM [52] modulates feature maps at a particular level in the CNN hierarchy (i.e. the output of a certain ResBlock) by attending over the hidden states of the question encoder. The attention mechanism is initialized with the context vector computed for the previous level, which makes this approach scale better to settings with a longer input sequence such as a dialogue.

VideoNavQA requires the incorporation of an additional dimension in the visual reasoning process. We propose a *temporal multi-hop* model (see Figure 4) to condition the video features using the question input. We compute FiLM parameters for all ResBlocks at frame t by initializing the attention context with the one from frame $t - 1$, attending over question encoder hidden states, obtaining the new context vector and passing it through a linear layer. This way, the FiLM parameters at a certain time step are based on what has already been computed for previous frames, thus modelling the temporal structure of the input.

5 Experiments

5.1 Setup

We evaluate all of the models described in the previous section on the VideoNavQA task and produce an initial expectation of the feasibility of EQA, using accuracy to compare the overall performance between different architectures. More in-depth analysis is then carried out for each of the models, in order to determine their respective strengths across different question categories and types. All models are trained with the Adam optimizer [17], by monitoring the validation accuracy. Hyperparameters are reported after evaluating different combinations of layer hidden sizes, number of residual blocks and classifier dimensions.

Language-only models The recurrent model has an embedding layer of 512 units and an LSTM with 128 hidden units, whereas the BoW model uses an embedding size of 128. We use a batch size of 1024 and learning rates of $5e^{-5}$ and $1e^{-5}$, respectively.

Video-only models The V-CNN3D model has 3 {convolutional, max-pool, batch normalization (BN)} blocks with 64, 128, and 128 output feature maps, respectively, kernel size (1, 2, 2) for the first block and (4, 4, 4) otherwise. The classifier has 2 linear layers with 2048

and 128 hidden units respectively. BN [15] and ReLU [9] activations are used for each layer. For the V-CNN2D model, we use a VGG configuration with 5 {convolutional, max-pool} blocks with 16, 32, 64, 128 and 128 channels, respectively, and kernel size 2.

FiLM-extended models The models that process each frame in a FiLM fashion have 4 (GP) or 5 (AT) ResBlocks with 1024 channels, preceded by a 1×1 convolution on the input. The attention mechanism for FiLM AT has 128 hidden units, whereas the global max-pooling classifier obtains 32 channels via the 1×1 convolution. The temporal multi-hop model has 3 ResBlocks and 64 tail channels. It is trained with a learning rate of $5e^{-5}$ and a batch size of 16, whereas the other two use $1e^{-4}$ and 32.

Video dimensionality The videos have an initial size of $140 \times 3 \times 160 \times 208$. In order to make learning feasible time-wise and allow the exploration of models with bigger capacity, we extract features from an object detector pre-trained on a set of 2000 frames that are not part of the dataset, which we initialize with the output from the 10th layer of a Faster R-CNN [28]. The object detector has 3 {conv, conv, BN, max-pool} blocks with 512 output maps for each layer; the classifier consists of a 1024-unit linear layer followed by softmax. No significant difference in validation accuracy is noticed when training the Concat-CNN3D model with and without pre-trained features. We also adopt a randomized sub-sampling strategy to reduce the maximum video length from 140 to 35. On each training iteration, every section of 4 frames is reduced to 1 by randomly picking out a frame from the chunk—this enables models to become more robust and invariant to changes between nearby frames.

5.2 Results

Model performance As expected, the video-only baselines can only predict the most prevalent answer (*Yes/True*). The LSTM proves to be the most powerful language model, achieving 7.5% more than the bag-of-words. Surprisingly, Concat-CNN2D manages to outperform the other models, obtaining an accuracy of 64.47%. The overall performance of Concat-CNN3D and FiLM AT is roughly similar, whereas FiLM GP achieves approximately 0.7% less than Concat-CNN2D. The temporal multi-hop falls about 0.3% behind FiLM GP, nevertheless outperforming the LSTM by 7%.

Overall, these results suggest that the models presented manage to exploit the visual context of the environment, improving by a significant margin over the language-only capabilities of the LSTM and BoW. This finding represents an initial validation of the feasibility of the VideoNavQA task, ensuring that the manner in which we constructed the dataset allows the visual reasoning process to be generalizable across new environments. A more focused investigation of novel VQA techniques is further required to accommodate the rich dimensionality and contextual information encountered in the videos.

Analysis by question category Figure 5 provides a more detailed analysis of the model performances on the test set across the eight question categories, in order to reveal the capabilities developed on VideoNavQA sub-tasks. All models score highest on *existence* questions, with FiLM GP having an edge over others, whereas *identifying object types* seems to be the hardest task: Concat-CNN3D performs worst, whereas FiLM AT achieves the highest accuracy, closely followed by Multi-hop. FiLM AT also does best at *identifying the room locations of objects*, this result being consistent with the idea of focusing on relevant features. All models achieve roughly 50% on the *counting* task, with slight improvements brought by

Model	Accuracy All	Yes/No	Other	Num
BoW	49.02	57.67	30.57	40.21
LSTM	56.49	68.36	35.27	38.90
V-CNN3D	33.29	-	-	-
V-CNN2D	33.62	-	-	-
Concat-CNN3D	64.00	72.99	49.12	49.10
Concat-CNN2D	64.47	73.50	49.20	49.59
FiLM-GP	63.79	72.91	47.71	50.00
FiLM-AT	64.08	72.93	49.54	49.26
Temporal multi-hop	63.53	71.81	49.54	50.16

Table 3: Results for all models reported in terms of accuracy on the VideoNavQA test set. We also use standard VQA reporting of Yes/No, other, and number categories [2].




Figure 5: Comparative performance of the models on each question category. Questions from the top row categories have binary answers, whereas the ones from the bottom row are answered by integer counts, room types, colors and object types, respectively.

Multi-hop and FiLM GP. When *comparing object attributes* (i.e. room location, color), most models achieve around 73%, with Concat-CNN2D obtaining a 1% improvement. *Identifying colors* works well for both concatenative models, with Multi-Hop achieving the next best result of 58% and FiLM AT performing worst, just under 56%. FiLM GP does relatively well at *comparing object/room sizes*, followed by FiLM AT, whereas Multi-Hop obtains a weaker performance of around 71%. Alternatively, when *comparing counts*, FiLM GP does much worse than all the other models and Concat-CNN2D considerably outperforms them, with a performance gap of 5% between the two architectures.

6 Conclusion

We have introduced the VideoNavQA task and dataset as a means to gain a better understanding of what is achievable in the EQA domain, where the environment context is visually rich—here, suitable methods need to be integrated accordingly to build an agent that can reason, navigate and act. This task represents an alternative view of the EQA paradigm: the navigation aspect is made trivial by providing nearly-optimal trajectories to the agent, but the reasoning difficulty becomes much higher, having 28 question types across 8 categories.

By reporting initial results from a variety of single-modality and integrative models, we establish a lower bound on the achievable performance in the VideoNavQA task. Since results are not widely different across architectures, we note the need for more investigation into whether there exist other video QA models that can better isolate the required visual information, bridging the gap to human performance. By first tackling the VideoNavQA task effectively, the research community will allow more clear and measured progress when attempting to integrate navigation into the EQA task, gradually reaching a better solution.

Acknowledgements

We wish to thank Ankesh Anand and Ethan Perez for the useful discussions over the course of this project. CC is supported by DREAM CDT and was funded by Mila during the time in Montréal. EB is funded by IVADO. We also thank the University of Cambridge Research Computing Services for providing HPC cluster resources.

References

- [1] Ankesh Anand, Eugene Belilovsky, Kyle Kastner, Hugo Larochelle, and Aaron Courville. Blindfold Baselines for Embodied QA. *arXiv preprint arXiv:1811.05013*, 2018.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [4] Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron Courville. Systematic Generalization: What Is Required and Can It Be Learned? *arXiv preprint arXiv:1811.12889*, 2018.
- [5] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [6] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied Question Answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [7] Abhishek Das, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Neural Modular Control for Embodied Question Answering. *arXiv preprint arXiv:1810.11181*, 2018.
- [8] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*, 2016.

- [9] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323, 2011.
- [10] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. IQA: Visual question answering in interactive environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4089–4098, 2018.
- [11] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [12] Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. Learning to reason: End-to-end module networks for visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 804–813, 2017.
- [13] Drew A Hudson and Christopher D Manning. Compositional Attention Networks for Machine Reasoning. *arXiv preprint arXiv:1803.03067*, 2018.
- [14] Drew A Hudson and Christopher D Manning. GQA: a new dataset for compositional question answering over real-world images. *arXiv preprint arXiv:1902.09506*, 2019.
- [15] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [16] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2901–2910, 2017.
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Her-rasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: An Interactive 3D Environment for Visual AI. *arXiv*, 2017.
- [19] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. TVQA: Localized, Compositional Video Question Answering. In *EMNLP*, 2018.
- [20] Mateusz Malinowski, Carl Doersch, Adam Santoro, and Peter Battaglia. Learning visual question answering by bootstrapping hard attention. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–20, 2018.
- [21] Dmytro Mishkin, Alexey Dosovitskiy, and Vladlen Koltun. Benchmarking Classic and Learned Navigation in Complex 3D Environments. *arXiv preprint arXiv:1901.10915*, 2019.
- [22] Jonghwan Mun, Paul Hongsuck Seo, Ilchae Jung, and Bohyung Han. MarioQA: Answering Questions by Watching Gameplay Videos. In *ICCV*, 2017.

- [23] Medhini Narasimhan, Svetlana Lazebnik, and Alexander Schwing. Out of the box: Reasoning with graph convolution nets for factual visual question answering. In *Advances in Neural Information Processing Systems*, pages 2654–2665, 2018.
- [24] Dat Tien Nguyen, Shikhar Sharma, Hannes Schulz, and Layla El Asri. From FiLM to Video: Multi-turn Question Answering with Multi-modal Context. *arXiv preprint arXiv:1812.07023*, 2018.
- [25] Will Norcliffe-Brown, Stathis Vafeias, and Sarah Parisot. Learning conditioned graph structures for interpretable visual question answering. In *Advances in Neural Information Processing Systems*, pages 8334–8343, 2018.
- [26] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [27] Mengye Ren, Ryan Kiros, and Richard Zemel. Exploring models and data for image question answering. In *Advances in neural information processing systems*, pages 2953–2961, 2015.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [29] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A Platform for Embodied AI Research. *arXiv preprint arXiv:1904.01201*, 2019.
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [31] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic Scene Completion from a Single Depth Image. *Proceedings of 30th IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [32] Florian Strub, Mathieu Seurin, Ethan Perez, Harm De Vries, Jérémie Mary, Philippe Preux, Aaron Courville, and Olivier Pietquin. Visual Reasoning with Multi-hop Feature Modulation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784–800, 2018.
- [33] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. MovieQA: Understanding Stories in Movies through Question-Answering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [34] Damien Teney, Lingqiao Liu, and Anton van den Hengel. Graph-structured representations for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2017.
- [35] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.

- [36] Erik Wijmans, Samyak Datta, Oleksandr Maksymets, Abhishek Das, Georgia Gkioxari, Stefan Lee, Irfan Essa, Devi Parikh, and Dhruv Batra. Embodied Question Answering in Photorealistic Environments with Point Cloud Perception. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [37] Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. Building generalizable agents with a realistic and rich 3D environment. *arXiv preprint arXiv:1801.02209*, 2018.
- [38] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29, 2016.
- [39] Licheng Yu, Xinlei Chen, Georgia Gkioxari, Mohit Bansal, Tamara L. Berg, and Dhruv Batra. Multi-Target Embodied Question Answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.