

# Look, Listen, and Act: Towards Audio-Visual Embodied Navigation

Chuang Gan<sup>\*1</sup>, Yiwei Zhang<sup>\*2</sup>, Jiajun Wu<sup>3</sup>, Boqing Gong<sup>4</sup>, Joshua B. Tenenbaum<sup>3</sup>

**Abstract**—A crucial ability of mobile intelligent agents is to integrate the evidence from multiple sensory inputs in an environment and to make a sequence of actions to reach their goals. In this paper, we attempt to approach the problem of Audio-Visual Embodied Navigation, the task of planning the shortest path from a random starting location in a scene to the sound source in an indoor environment, given only raw egocentric visual and audio sensory data. To accomplish this task, the agent is required to learn from various modalities, i.e., relating the audio signal to the visual environment. Here we describe an approach to audio-visual embodied navigation that takes advantage of both visual and audio pieces of evidence. Our solution is based on three key ideas: a visual perception mapper module that constructs its spatial memory of the environment, a sound perception module that infers the relative location of the sound source from the agent, and a dynamic path planner that plans a sequence of actions based on the audio-visual observations and the spatial memory of the environment to navigate toward the goal. Experimental results on a newly collected Visual-Audio-Room dataset using the simulated multi-modal environment demonstrate the effectiveness of our approach over several competitive baselines.

## I. INTRODUCTION

We humans perceive and navigate through the world by an integration of multiple sensory inputs, including but not limited to acoustic and visual representations. Extensive evidence from cognitive psychology and neuroscience suggests that humans, even in the young-child ages, are remarkably capable of capturing the correspondences between different signal modalities to arrive at a thorough understanding of natural phenomena [1], [2].

In this work, we aim to build a machine model to achieve such human competency on audio-visual scene analysis. Particularly, we focus on an under-explored problem: how to teach an intelligent agent, equipped with a camera and a microphone, to interact with environments and navigate to the sound source. This task is more complex and challenging than understanding static images. We call it *audio-visual embodied navigation*. The agent has to sequentially make decisions from multiple sensory observations in order to navigate through an environment.

Figure 1 illustrates an example of the audio-visual embodied navigation problem. An agent starts randomly from an indoor environment and provided a sound impulse (e.g., a phone is ringing). The agent then needs to figure out the shortest trajectory to navigate from its starting location to the source of the sound (the goal location). To accomplish this task, the agent must relate the sound to the visual environment




Fig. 1: Illustration of the proposed Audio-Visual Embodied Navigation. Given a sound source, the agent perceives and relates it with the visual environment so as to navigate towards the goal.

and meanwhile map from the raw sensory inputs to the actions of path planning. This task is very challenging, because it essentially asks the agent to reverse-engineer the causal effect of the signals it perceives. It has to not only understand the visual surrounding but also reason upon it for the potential sound source. If the sound source is in a different room, the agent additionally has to learn how to move to another room. Progress on this task will facilitate many real-world applications in home robots. For example, the sound-source-seeking ability can empower the robots to help humans find their cellphones or turn off faucets.

We tackle this problem by drawing insights from how humans complete this task. Imagine how a user seeks the sound source in a novel environment? S/he first gets a sense about the location, the surrounding, as well as an estimate about the distance and direction to sound source. As s/he moves around, the room layout and her/his memory about the sound and the room also play vital roles when s/he makes decisions for future actions. To this end, in this paper, we design an audio-visual network to let the machine agent imitate the navigation mechanism we humans use in real life. We systematically study two settings: explore-and-act and non-exploration. In the first setting, the agent is allowed to explore the environment up a certain number of steps before we alarm the sound; Thus it can build an incomplete spatial memory of the room from visual observations during the exploration and can refer back to this internal representation in resolving the subsequent navigation tasks. In the second setting, the agent has to build a dynamic spatial map during

<sup>\*</sup>indicates equal contributions. <sup>1</sup>MIT-IBM Watson AI Lab, <sup>2</sup>Tsinghua University, <sup>3</sup>Massachusetts Institute of Technology, <sup>4</sup>Google. Project page: <http://avn.csail.mit.edu>

the navigation. When the agent is presented a task of searching for the sound source, akin to humans, the agent must first infer the relative position of the sound source and then makes sequential actions to move towards the source based on its spatial memory and visual-audio observations. It can also update its spatial memory at each time step.

To sum up, our work makes the following contributions:

- We build a multi-modal navigation environment to facilitate research on audio-visual embodied navigation. The environment contains fairly complex apartments and has an integrated sound module that observes certain physical laws.
- We present a Visual-Audio Room (VAR) benchmark to systematically evaluate the performances of multi-modal navigation agents.
- We propose a framework for audio-visual embodied navigation and contrast it to several competitive baselines. Experiments verify the effectiveness of our approach and meanwhile reveal the key challenges in this task, shedding lights on future work.

## II. RELATED WORK

### A. Target-driven Navigation

Early work solves the navigation task by building a map of the scene using SLAM and then planing the path in this map [3]. More recently, deep learning-based methods have been used to plan actions directly from raw sensory data. Zhu et al. [4] investigate a reactive deep RL based navigation approach to find the picture of the target object in a discretized 3D environment. Gupta et al. [5] learn to navigate via a mapper and planner. Sadeghi et al. [6] study a RL approach that can solely use simulated data to teach an agent to fly in real. Mirowski et al. [7] improve the navigation results by jointly training with auxiliary tasks such as loop closure detection and depth estimations from RGB. Brahmhatt et al. [8] explore a CNN architecture to navigate large cities with street-view images. Wu et al. [9] combine deep RL with curriculum learning in a first-person shooting game setting. Yang et al. [10] propose to improve visual navigation with semantic priors. McLeod et al. [11] leverage past experience to improve future robot navigation in dynamically unknown environments. Katyal et al. [12] use generating networks to predict occupancy map representations for future robot motions. Mousavian et al. [13] demonstrate the use of semantic segmentation and detection masks for target driven visual navigation. The work by Savinov et al. [14] is inspirational to ours in terms of the use of memory in navigation. Unlike their topological graph memory, we use a key-value structure to better capture the incomplete nature of our agent’s internal knowledge about the environment and also introduce a dynamic spatial memory for the non-exploration setting.

There is also considerable work on vision-language based embodied navigation [15], [16], [17], [18], which map directly from language instructions and visual observations to actions. In contrast to all these approaches, the goal of embodied audio-visual navigation task is to find the sound source, thus

require the agent to integrate the visual and acoustic cues to plan a sequence of actions. Concurrent to our work, Chen et al. [19] also poses the task of audio-visual embodied navigation on the Habitat platform [20].

### B. Sound Localization

The research problem of sound localization, which aims to identify which regions in a video make the sound, has been studied for decades. Existing approaches fall into two broad categories: computation-based approaches and learning-based approaches. Early work measure the correlations between pixels and sounds using a Gaussian process model [21], subspace methods [22], canonical correlation analysis [23], hand-crafted motion [24], and segmentation [25]. More recently, researchers have proposed to train a deep neural network to see and hear many unlabeled videos to localize the objects that sound [26], [27], [28], [29], [30], [31], [32], [33], [34], [35]. Such methods address the task of localizing the regions on videos. In contrast, our goal is to find the sound source that might also be non-line-of-sight in the virtual environment. Our work is also related to acoustic based approaches for the sound source localization [36], [37], [38]. They typically require special devices (e.g. microphone arrays) to record the sounds. Our work goes beyond that, since the mobile agent needs to further connect the result of sound localization with the visual environment for the path planing.

## III. ENVIRONMENT

We build a multi-modal virtual environment on top of the AI2-THOR platform [39], which contains a set of scenes built by using the Unity Game engine. We further incorporate an spatial audio software development kit, Resonance Audio API [40], into the Unity game engine to support the audio-visual embodied navigation task. An agent equipped with a camera and a microphone can then leverage two different sensing modalities (egocentric RGB images and sound) to perceive and navigate through the scenes.

### A. Scene

The AI2-THOR platform provides near photo-realistic indoor scenes. It contains 120 scenes in four categories: *kitchens*, *bedrooms*, *living rooms*, and *bathrooms*. Each room category contains 30 rooms with diverse visual appearances. As the rooms in AI2-THOR are isolated from each other, they are not challenging enough for the area-goal navigation tasks. As shown in Figure 1, we instead manually concatenate several rooms into multi-room apartments. We were able to build seven apartments in total for this work. Similarly to Zhu et al. [4], we also discretize the apartments into grid worlds for the ease of quantitative evaluation. Particularly, each grid square is  $0.5 \times 0.5$  square meters. In total, there are about 150 to 200 grid cells for each apartment.

### B. Sound

The acoustic engine is implemented using the Resonance Audio API. This spatial audio SDK can create high fidelity spatial sound in the virtual environment and support acoustic

ray-tracing according to the geometry of object and house. The Resonance Audio API also supports the simulation of how sound waves interact with objects in various geometry and diverse materials by mapping visual materials to acoustic materials with realistic frequency-dependent sound reverberation properties. How sound waves interact with human ears can also be simulated accurately based on interaural time differences, interaural level differences, and the spectral effects. The scenes equipped with the Resonance Audio API create the illusion that sounds come from specific locations in the virtual world. Thus we can use this module to provide the agent with stereo sound frames. In our experimental setting, we mainly consider the continuous sound sources (e.g., ring tone and alarm alert).

### C. Problem Setup

We consider two problem setups in the visual-audio embodied navigation task: explore-and-act and non-exploration for seeking the sound sources. In the first setting, we let the agent conduct two stages of interactions with an environment, in spirit similar to the exploration setup in Savinov et al. [14]: exploration and goal-directed navigation. During the exploration stage, the agent randomly walks around the environment until its trajectory’s length reaches a given budget. This exploration experience can be utilized to build an internal world model of the environment for subsequent goal-directed navigation tasks. In the second setting, the agent has to simultaneously build a spatial memory of the environment while it navigates towards the goal.

In the testing phase, we randomly select a starting position for the agent and a target position for the sound source. At each time step, the agent takes an egocentric visual observation  $o_v$  and sound observation  $o_s$  about the environment and then draws an action from a set  $A$ . The action set  $A$  comprises the following: *Move Backward*, *Move Forward*, *Rotate Left*, *Rotate Right*, and *Stop*.

## IV. APPROACH

In this section, we introduce our algorithm for the audio-visual embodied navigation. As shown in Figure 2, it consists of three components: a visual perception mapper, an audio perception module, and a dynamic path planner.

In an attempt to tackle the visual-audio embodied navigation task, the agent can either use the visual perception mapper to retrieve from its spatial memory constructed during the exploration stage or use the occupancy map estimated during navigation to build a *partial* 2D spatial map (a graph with edges) of the environment. After that, the agent leverages the sound perception module to estimate the direction and distance of the sound source from its current location. Finally, the agent plans to find the shortest possible path to reach the goal based on the inference results from the visual and sound perception modules and the *partial* spatial map. In the middle of the navigation, the dynamic planner can also update the the estimated goal location as well as the agent’s navigation model based on any new visual and audio observations.

### A. Visual Perception Mapper

*a) Explore-and-Act Visual Mapper:* The exploration based visual mapper consists of a spatial memory and a non-parametric retrieval model. During the exploration, we use a key-value based spatial memory network [41] to encode the environment visited by the agent. Visual observation is stored in the key part, while the meta-data (e.g., location coordinates and actions taken by the agent) is stored in the value part. Concretely, we forward each first-person view RGB image through an ImageNet-pretrained ConvNet model, ResNet [42], to extract a visual feature vector, followed by  $\ell_2$  normalization. Each feature vector of the visual observation and the corresponding coordinates and orientation of the agent form a key-value pair (Figure 3 (a)).

In the goal-oriented navigation stage, we use the same feature extraction pipeline to obtain a query feature vector extracted from the agent’s first-person view. A non-parametric retrieval step is then conducted over the keys of the spatial memory. The top three memory slots that are closest to the query under the cosine distance are returned. We average the stored coordinates to estimate the location of the agent. Note that, in our experiments, we enforce the starting location of the agent has no overlap with the explored positions, so the localization of the agent in the early navigation stage has significant amount of uncertainty.

*b) Non-Exploration Visual Mapper:* Since the random exploration stage might not always be allowed in real applications, we further introduce a visual perception mapper that can construct a dynamic spatial memory on the fly while the agent navigates through an environment. We construct the spatial memory as a top-down 2D occupancy map  $M \in R^{N \times N}$  about the grid world environment (Figure 3 (b)). Each entry of the memory represents a grid cell in the environment, taking values in the range  $[0, 1]$  and indicating the agent’s belief about whether the grid cell can be traversed. We initialize the value to 0.5 and the starting position of the agent at the center of the memory  $M$ . Since the agent only has access to the visual observations of the nearby area, the visual perception mapper produces a local occupancy map in front of the agent at each time step. The agent then updates the local occupancy map using Bayes’ rule [43].

In order to infer if a cell is transversable or not, we train a convolutional neural network to map from RGB images to the agent’s believes. We collect training data by utilizing RGB images and groundtruth 2D occupancy maps in two training apartments. We feed a first-person RGB image into the network to predict the free space probability in a small window size ( $5 \times 5$  grids in our experiments).

### B. Sound Perception Module

Motivated by the fact that humans can mentally infer the sound locations, here we introduce a sound perception module used for estimating the coordinates of the goal, i.e., the location of the sound source in our environment.

To make the network easy to learn, we use relative locations to estimate the coordinates of the goal. We define the relative location in a grid world as follows. The relative location of

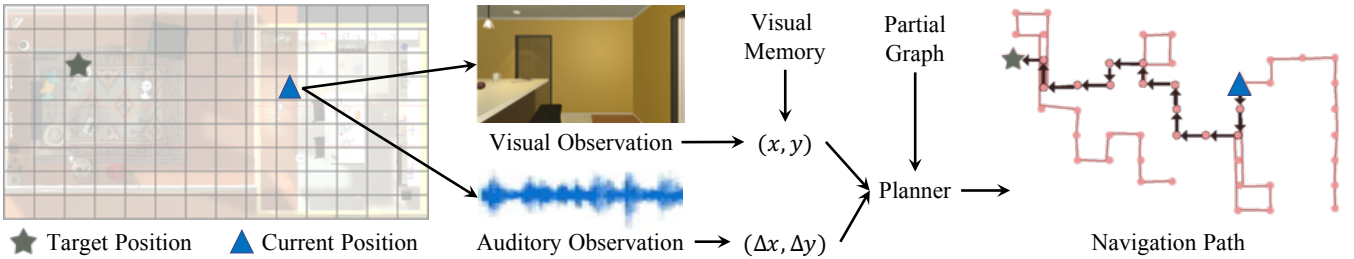


Fig. 2: During navigation, the agent uses a visual perception module to localize itself and a sound perception module to estimate the sound source location. Finally, the agent connects the estimation with the partial graph built during exploration, and plans a path to reach the goal.

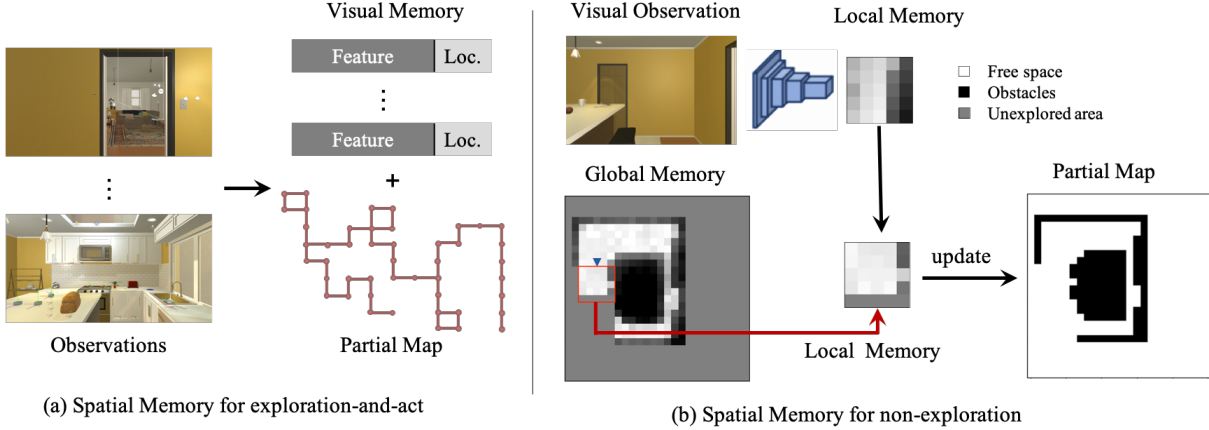


Fig. 3: The spatial memory used for exploration-and-act and non-exploration based navigation.

the audio source from the agent’s perspective, denoted by  $(x, y)$ , represents that the audio source is located  $x$  meters to the right of the agent and  $y$  meters in front of the agent. The relative position can thus be calculated by jointly considering the orientation of the agent and the absolute coordinates of the sound source and the agent. In the absolute coordinate system, the east axis is the positive  $x$  axis and the north axis is the positive  $y$  axis. Assuming that the absolute coordinates of the sound source and the agent in the grid world are  $(16, 18)$  and  $(4, 8)$ , respectively, the relative location of the sound source is then  $(12, 10)$  if the orientation of the agent is towards the north. If the agent faces the west, the relative location of the sound source becomes  $(10, -12)$ .

We collect training data by using all the recorded audio clips in two training apartments. To pre-process the stereo sound, we first re-sample the sound to 16 HZ and then convert it into a spectrogram through Short-Time Fourier Transform (STFT). We feed the spectrogram into a five-layer convolutional network. The training objective is to estimate the relative sound location with respect to the agent. Specifically, we train the sound perception module with the mean square error (MSE) loss in a supervised way. The absolute coordinate of the target can be calculated from the relative location and the agent’s coordinates. The sound perception module also contains noises, since the layout and the surface materials of the room influence the sound. We also train a separate sound classification model to decide if the agent reaches the goal; if so, it will take a *stop* action.

### C. Dynamic Path Planner

The agent employs a dynamic path planner to navigate in the environment upon observing changes and meanwhile

updating the memory as a result of those actions.

For the explore-and-act setting, the agent can construct a *partial graph*  $G = (V, E)$  of the room based on a sequence of the agent’s exploration trajectory  $(x_0, x_1, \dots, x_n)$  and actions  $(a_1, a_2, \dots, a_n)$ . Denote by  $V$  and  $E$  the nodes and edges, respectively. Each node  $v_i \in V$  in the graph stores a location (coordinates) in the environment. The edge exists between two nodes  $v_i$  and  $v_j$  if the nodes correspond to consecutive time steps.

For the non-exploration based approach, the agent can use the memory to construct a *partial graph* of the environment. According to the spatial memory, each grid can be divided into three types: free spaces, obstacles, and unexplored areas. The *partial graph*  $G = (V, E)$  is generated by assuming the unexplored and unoccupied grids are transversable, converting the transversable grids into nodes, and connecting adjacent nodes via edges. When a new obstacle is detected, the corresponding node and edge(s) will be deleted.

The dynamic planner outputs a path (sequence of actions) that transits the agent to the desired target location based on the graph of the environment and the estimated target coordinates. At each time step  $t$ , the agent tries to find a shortest path from the graph node of the agent itself  $x_{agent}^t$  to the graph node  $x_{graph}^t$  that is nearest to the estimated coordinates of the sound source  $x_{target}^t$ . We use the Dijkstra’s algorithm to solve this problem. Specifically, we can use this algorithm to obtain a shortest path denoted as  $p_{min}^t = \{a_0^t, a_1^t, \dots, a_m^t\}$ . We select  $a_0^t$  as the action to take at time  $t$ .

## V. EXPERIMENT

In this section, we first describe the dataset and environment we created for the sound-source-seeking audio-visual




Fig. 4: Trajectories of ours (red) against those of A3C (V+A+M+Mapper) (blue). Our model finds more efficient routes to the targets.

embodied navigation task, followed by demonstrating the benefits of our method against several competing baselines.

#### A. The Visual-Audio Room Dataset

In order to systematically evaluate the multi-modal navigation performances and facilitate future study on this research direction, we collect a new Visual-Audio Room (VAR) benchmark. Our ultimate goal is to enable robots to navigate in indoor environments by fusing visual and audio observations. As it is very challenging to conduct thorough and controlled experiment with physical robots, we instead use a 3D simulated environments, built on the top of AI-Thor platform [4]. There are seven apartments used in the experiment. We split them to two apartments for training and five apartments for testing. We consider three recorded audio categories: ring tone, alert alarm, and clock clicking as the sound source in the work. To generate navigation data, we pre-extract the first person view RGB images in four orientations over all the grids of the seven apartments. For the training apartments, we randomly pick up ten locations to put the sound source and then record the sound the agent hears in four orientations of all locations in the room. For the five testing apartments, we pick five grid locations to put the sound source and also record the sound the agent hears in all locations. In total, we collect 3,728 RGB images, and record 75,720 audio clips.

#### B. Experiment Setup

**Setup.** We aim to test the agent’s ability to navigate in a new environment. During evaluation, for each of the five possible locations of the sound source in each room, we randomly pick 20 locations and orientations of the agent. Therefore, we conduct 100 testing episodes for each sound. We set the time duration of each sound at each time step as 1 second. The number of random walk steps in the explore-and-act setting are set as 400 .

**Evaluation metrics.** We use two metrics to evaluate the models: success rates and Success weighted by Path Length (SPL) [44]. *SPL* measures the agent’s navigation performance across all the testing scenarios as  $\frac{1}{N} \sum_{i=1}^N S_i \frac{l_i}{\max(p_i, l_i)}$ , where  $N$  denotes the number of testing episodes,  $S_i$  is a binary indicator of success in the testing episode  $i$ ,  $P_i$  represents path length, and  $l_i$  indicates the shortest distance to the goal.

#### C. Baseline

We consider seven baseline methods for evaluations.

- **Random Walk.** This is the simplest heuristic for the navigation task. The agent randomly samples an action from the action space at each time step.
- **Greedy Search (A).** This is a sound-only baseline. The agent navigates in the environment only using the sound perception module. The agent moves greedily towards the sound source without observing the structure and layout of the apartment.
- **A3C (V)** [45]. Asynchronous advantage actor-critic (A3C) is a state-of-the-art Deep RL based approach for navigation. This is a goal-oriented, vision-only navigation baseline without memory.
- **A3C (V+A).** This is a goal-oriented, audio-vision navigation baseline without memory. In our setting, the input of the A3C model is an audio-visual representation of the current state. For a fair comparison, we concatenate the feature representations extracted from the visual and sound perception networks used in our approach.
- **A3C (V+A+Mapper).** In this setting, the input of the A3C model is an audio-visual representation of the current state and the explicit 2D occupancy maps used in our non-exploration based approach. The top-down occupancy map is concatenated with the audio-visual representation.
- **A3C (V+A+M).** [14]. Similar to the baseline used in Savinov et al. [14], we have also implemented an A3C equipped with LSTM memory as a baseline for our task. First, the agent navigates the environment for the same number of random walking steps in the exploration mode without a given goal. There is no reward for reaching a goal. At test time, we feed the exploration sequence to the LSTM agent and then let it perform goal-directed navigation without resetting the LSTM state. We expect that the LSTM module can implicitly build an internal representation of them.
- **A3C (V+A+M+Mapper).** We further combine the explicit and implicit memory to training the A3C.

#### D. Results

Table I summarizes the results for the three different sounds sources in the five testing apartments. The proposed audio-

Approach	Sound	Room 1	Room 2	Room 3	Room 4	Room 5	Average
Random Search	Ring	6/3.5	10/6.8	8/4.6	5/2.6	8/5.1	7.4/4.5
	Alarm	5/3.3	9/5.7	7/4.2	7/4.1	7/3.9	7/4.2
	Clock	8/5.2	9/6.6	6/3.4	4/2.1	7/4.4	6.8/4.3
Greedy Search (A)	Ring	12/10.5	15/13.5	13/11.8	10/8.2	12/11.4	12.4/11.1
	Alarm	9/8.2	13/11.2	10/8.4	8/7.1	11/10.2	10.2/9.0
	Clock	13/12.0	15/13.4	12/11.0	12/9.6	11/10.1	12.6/11.2
A3C (V)	Ring	5/4.7	12/11.5	10/9.1	8/7.8	10/8.9	9/8.4
	Alarm	7/6.0	11/9.3	9/8.9	7/6.2	8/7.5	8.4/7.6
	Clock	7/6.1	8/7.5	9/7.8	9/8.6	11/9.8	8.8/8.0
A3C (V+A)	Ring	18/15.7	25/18.5	15/11.7	14/11.5	20/14.9	18.4/14.5
	Alarm	19/15.9	22/17.1	16/14.5	14/12.5	19/14.6	18/14.9
	Clock	18/13.0	27/23.6	18/13.8	16/13.8	22/16.1	20.2/16.1
A3C (V+A+M)	Ring	27/18.4	36/25.7	33/22.2	29/20.8	31/21.2	31.2/21.7
	Alarm	24/16.5	37/25.0	29/20.4	31/21.4	33/23.2	30.8/21.3
	Clock	29/20.3	34/25.6	27/19.3	34/24.3	28/19.9	30.4/21.9
A3C (V+A+Mapper)	Ring	24/17.9	30/21.3	21/15.6	19/14.3	24/17.4	23.6/17.3
	Alarm	22/16.6	29/20.1	23/16.9	20/15.2	25/17.9	23.8/17.3
	Clock	25/18.5	31/21.7	22/16.2	20/14.9	24/17.2	24.4/17.7
A3C(V+A+M+Mapper)	Ring	30/20.7	38/25.7	35/23.8	33/23.3	33/23.8	33.8/23.5
	Alarm	28/19.1	39/26.4	33/22.5	34/23.5	36/24.4	34.0/23.2
	Clock	31/21.3	37/24.9	31/22.5	36/24.9	32/23.5	33.4/23.4
Ours (no Exp.)	Ring	69/42.2	55/37.4	53/35.4	56/38.6	59/40.3	58.4/38.8
	Alarm	67/41.7	53/35.9	55/36.8	56/37.9	58/39.7	57.8/38.4
	Clock	70/43.9	58/39.8	56/37.7	54/36.4	61/41.6	59.8/40.0
Ours (w/ Exp.)	Ring	<b>78/58.0</b>	<b>66/47.1</b>	<b>62/43.4</b>	<b>68/55.0</b>	<b>74/51.3</b>	<b>69.6/51.0</b>
	Alarm	<b>76/57.0</b>	<b>62/45.5</b>	<b>65/45.7</b>	<b>66/45.1</b>	<b>69/54.4</b>	<b>67.6/49.5</b>
	Clock	<b>83/60.1</b>	<b>70/56.7</b>	<b>69/48.7</b>	<b>62/51.1</b>	<b>70/55.1</b>	<b>70.8/54.4</b>

TABLE I: Success rates (%) / Success weighted by Path Length (SPL) of different audio-visual navigation methods.

visual agent performs better than all baselines across all metrics. In the exploration-and-act setting, our model (w/Exp.) remarkably achieves over a 65% success rate in average for every sound source, almost doubling the best-performing baseline. Figure 4 presents the results in two testing rooms in more detail, by plotting the trajectory of our approach against that of the best performing baseline (A3C (V+A+M+Mapper)). In the non-exploration based setting, our approach (no Exp.) can also achieve nearly 60% success rate, even outperform the strongest baseline with explorations. These results indicate that our system is better to leverage the layout of the room and find the shortest path to reach the goal. For instance, the agent tends to come to the door when the sound is in another room, and thus efficiently reaches the goal. More qualitative results are available in the supplementary material.

From Table I, it is not surprising to find that the success rates of the random walk-based approaches are very low, because it is very challenging to conduct navigation in our environment without any prior knowledge. Greedy search is slightly better than random search, as the agent can anticipate the location of the goal from the sound. However, the success rates are still below 15%, because the agent has no visual prior knowledge of the room layout and misunderstands the shortest geometric distance as the shortest path, which does not always hold in our setting. These results demonstrate that it is important to connect visual and sound information for the audio-visual navigation task.

The performance of the A3C based approaches is significantly weaker than those previously reported in literature [45]. The key reason is that we focus on generalizations—putting agents in previously unseen environments. This is a much more challenging, but also more realistic scenarios. We also find that the performance gap is still large between our approach and A3C, which uses both implicit spatial memory from exploration and explicit spatial memory during navigation. We speculate that our approach can disentangle the visual perception and sound perception from path planing, thus providing better generalizability in new testing environments.

## VI. CONCLUSION

We tackle a novel problem, audio-visual embodied navigation, in a multi-modal virtual environment. The agent can use the internal structured representation of the environment to efficiently navigate to targets in previously unseen environments. Our multi-modal virtual environment consists of complex apartments with a sound module and realistic layouts. We demonstrate that the state-of-the-art deep reinforcement learning approaches meet difficulties in generalization. Our approach generalizes to new targets and new environments with remarkable results.

**Acknowledgement:** This work is in part supported by ONR MURI N00014-16-1-2007, the Center for Brain, Minds, and Machines (CBMM, NSF STC award CCF-1231216), and IBM Research.

## REFERENCES

- [1] H. McGurk and J. MacDonald, "Hearing lips and seeing voices," *Nat.*, vol. 264, no. 5588, p. 746, 1976.
- [2] K. P. Körding, U. Beierholm, W. J. Ma, S. Quartz, J. B. Tenenbaum, and L. Shams, "Causal inference in multisensory perception," *PLOS ONE*, vol. 2, no. 9, p. e943, 2007.
- [3] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [4] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *ICRA*, 2017.
- [5] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik, "Cognitive mapping and planning for visual navigation," in *CVPR*, 2017.
- [6] F. Sadeghi and S. Levine, "Cad2rl: Real single-image flight without a single real image," in *RSS*, 2017.
- [7] P. Mirowski, R. Pascanu, F. Viola, H. Soyer, A. J. Ballard, A. Banino, M. Denil, R. Goroshin, L. Sifre, K. Kavukcuoglu *et al.*, "Learning to navigate in complex environments," in *ICLR*, 2017.
- [8] S. Brahmabhatt and J. Hays, "Deepnav: Learning to navigate large cities," in *CVPR*, 2017.
- [9] Y. Wu and Y. Tian, "Training agent for first-person shooter game with actor-critic curriculum learning," in *ICLR*, 2017.
- [10] W. Yang, X. Wang, A. Farhadi, A. Gupta, and R. Mottaghi, "Visual semantic navigation using scene priors," in *ICLR*, 2019.
- [11] S. McLeod and J. Xiao, "Navigating dynamically unknown environments leveraging past experience," in *ICRA*, 2019.
- [12] K. Katyal, K. Popek, C. Paxton, P. Burlina, and G. D. Hager, "Uncertainty-aware occupancy map prediction using generative networks for robot navigation," in *ICRA*, 2019.
- [13] A. Mousavian, A. Toshev, M. Fiser, J. Kosecka, A. Wahid, and J. Davidson, "Visual representations for semantic target driven navigation," in *ICRA*, 2019.
- [14] N. Savinov, A. Dosovitskiy, and V. Koltun, "Semi-parametric topological memory for navigation," in *ICLR*, 2018.
- [15] D. Fried, R. Hu, V. Cirić, A. Rohrbach, J. Andreas, L.-P. Morency, T. Berg-Kirkpatrick, K. Saenko, D. Klein, and T. Darrell, "Speaker-follower models for vision-and-language navigation," in *NeurIPS*, 2018.
- [16] H. Yu, H. Zhang, and W. Xu, "Interactive grounded language acquisition and generalization in a 2d world," in *ICLR*, 2018.
- [17] H. Mei, M. Bansal, and M. R. Walter, "Listen, attend, and walk: Neural mapping of navigational instructions to action sequences," in *AAAI*, 2016.
- [18] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. van den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *CVPR*, 2018.
- [19] C. Chen, U. Jain, C. Schissler, S. V. A. Gari, Z. Al-Halah, V. K. Ithapu, P. Robinson, and K. Grauman, "Audio-visual embodied navigation," *arXiv:1912.11474*, 2019.
- [20] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik *et al.*, "Habitat: A platform for embodied AI research," in *ICCV*, 2019.
- [21] J. R. Hershey and J. R. Movellan, "Audio vision: Using audio-visual synchrony to locate sounds," in *NeurIPS*, 2000.
- [22] J. W. Fisher III, T. Darrell, W. T. Freeman, and P. A. Viola, "Learning joint statistical models for audio-visual fusion and segregation," in *NeurIPS*, 2001.
- [23] E. Kidron, Y. Y. Schechner, and M. Elad, "Pixels that sound," in *CVPR*, 2005.
- [24] Z. Barzelay and Y. Y. Schechner, "Harmony in motion," in *CVPR*, 2007.
- [25] H. Izadinia, I. Saleemi, and M. Shah, "Multimodal analysis for identification and segmentation of moving-sounding objects," *IEEE TMM*, vol. 15, no. 2, pp. 378–390, 2013.
- [26] R. Arandjelović and A. Zisserman, "Objects that sound," in *ECCV*, 2018.
- [27] A. Senocak, T.-H. Oh, J. Kim, M.-H. Yang, and I. S. Kweon, "Learning to localize sound source in visual scenes," in *CVPR*, 2018.
- [28] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," in *SIGGRAPH*, 2018.
- [29] T. Afouras, J. S. Chung, and A. Zisserman, "The conversation: Deep audio-visual speech enhancement," in *Interspeech*, 2018.
- [30] A. Owens and A. A. Efros, "Audio-visual scene analysis with self-supervised multisensory features," in *ECCV*, 2018.
- [31] R. Gao, R. Feris, and K. Grauman, "Learning to separate object sounds by watching unlabeled video," in *ECCV*, 2018.
- [32] H. Zhao, C. Gan, A. Rouditchenko, C. Vondrick, J. McDermott, and A. Torralba, "The sound of pixels," in *ECCV*, 2018.
- [33] H. Zhao, C. Gan, W.-C. Ma, and A. Torralba, "The sound of motions," in *ICCV*, 2019.
- [34] A. Rouditchenko, H. Zhao, C. Gan, J. McDermott, and A. Torralba, "Self-supervised audio-visual co-segmentation," in *ICASSP*, 2019.
- [35] C. Gan, H. Zhao, P. Chen, D. Cox, and A. Torralba, "Self-supervised moving vehicle tracking with stereo sound," in *ICCV*, 2019.
- [36] I. An, M. Son, D. Manocha, and S.-e. Yoon, "Reflection-aware sound source localization," in *ICRA*, 2018.
- [37] H. L. Van Trees, *Optimum array processing: Part IV of detection, estimation, and modulation theory*. John Wiley & Sons, 2004.
- [38] A. Zunino, M. Crocco, S. Martelli, A. Trucco, A. Del Bue, and V. Murino, "Seeing the sound: A new multimodal imaging device for computer vision," in *CVPR Workshop*, 2015.
- [39] E. Kolve, R. Mottaghi, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi, "Ai2-thor: An interactive 3d environment for visual ai," *arXiv:1712.05474*, 2017.
- [40] M. Gorzel, A. Allen, I. Kelly, J. Kammerl, A. Gungormusler, H. Yeh, and F. Boland, "Efficient encoding and decoding of binaural sound with resonance audio," in *AES*, 2019.
- [41] A. Miller, A. Fisch, J. Dodge, A.-H. Karimi, A. Bordes, and J. Weston, "Key-value memory networks for directly reading documents," in *EMNLP*, 2016.
- [42] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [43] H. P. Moravec, "Sensor fusion in certainty grids for mobile robots," in *Sensor devices and systems for robotics*. Springer, 1989, pp. 253–276.
- [44] P. Anderson, A. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva *et al.*, "On evaluation of embodied navigation agents," *arXiv:1807.06757*, 2018.
- [45] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *ICML*, 2016.