

Semantic Audio-Visual Navigation in Continuous Environments

Yichen Zeng^{1,2} Hebaixu Wang^{1,2} Meng Liu^{2,3,*} Yu Zhou^{4,2} Chen Gao^{2,5,*} Kehan Chen^{6,7} Gongping Huang^{1,*}
¹Wuhan University ²Zhongguancun Academy ³Shandong Jianzhu University
⁴Nankai University ⁵Tsinghua University ⁶CASIA ⁷UCAS

Abstract

Audio-visual navigation enables embodied agents to navigate toward sound-emitting targets by leveraging both auditory and visual cues. However, most existing approaches rely on precomputed room impulse responses (RIRs) for binaural audio rendering, restricting agents to discrete grid positions and leading to spatially discontinuous observations. To establish a more realistic setting, we introduce Semantic Audio-Visual Navigation in Continuous Environments (SAVN-CE), where agents can move freely in 3D spaces and perceive temporally and spatially coherent audio-visual streams. In this setting, targets may intermittently become silent or stop emitting sound entirely, causing agents to lose goal information. To tackle this challenge, we propose MAGNet, a multimodal transformer-based model that jointly encodes spatial and semantic goal representations and integrates historical context with self-motion cues to enable memory-augmented goal reasoning. Comprehensive experiments demonstrate that MAGNet significantly outperforms state-of-the-art methods, achieving up to a 12.1% absolute improvement in success rate. These results also highlight its robustness to short-duration sounds and long-distance navigation scenarios. The code is available at <https://github.com/yichenzeng24/SAVN-CE>.

1. Introduction

Embodied navigation requires agents to autonomously reach targets in previously unseen environments using sensory inputs. Most prior work has focused on either egocentric vision [24, 51, 57, 58] or incorporating textual instructions as an additional modality [5, 15, 16, 19, 22, 26, 49]. However, visual perception is often insufficient in indoor environments, where targets may lie outside the agent’s field of view or lack distinctive visual cues. For instance, an agent may need to respond to an elderly person falling in another room, locate a ringing phone in the house, or turn off a stove with boiling water. In these situations, au-




Figure 1. Illustration of the three navigation tasks: (a) the agent is restricted to discrete grid points and the sound-emitting goal is placed arbitrarily; (b) the goal, which emits a creaking sound for a limited duration, is semantically grounded in a chair; (c) the agent moves freely using fine-grained actions and the goal sound is available only within a short temporal window.

ditory perception provides critical complementary information, enabling agents to infer the locations and categories of otherwise invisible targets.

Building on this motivation, audio-visual navigation (AVN) [9, 20] enables agents to navigate toward sound-emitting goals in unmapped environments using audio-visual cues. Agents do not have access to explicit goal information such as coordinates, object categories, or textual instructions. Semantic audio-visual navigation (SAVN) [10] further extends this task by grounding short-duration audio signals in visual objects rather than arbitrary locations. However, both tasks rely on precomputed room impulse responses (RIRs), which demand terabytes of storage for binaural audio rendering. This dependence further confines agents to discrete grid positions (1 m resolution) and four fixed orientations [9], reducing task realism and hindering free exploration, as illustrated in Fig. 1(a) and (b).

To bridge this gap, we introduce SAVN-CE (*Semantic Audio-Visual Navigation in Continuous Environments*), a new task that allows agents to move freely in continuous 3D environments with fine-grained actions. In this setting, agents can perceive temporally and spatially coherent observations, which enhances their ability to reason about how their position and orientation evolve relative to the goal as they move. Unlike previous tasks, SAVN-CE features a highly dynamic goal condition, where the goal neither emits sound at the beginning nor persists until the end of an episode. Consequently, agents must first explore the en-

*Corresponding Author.

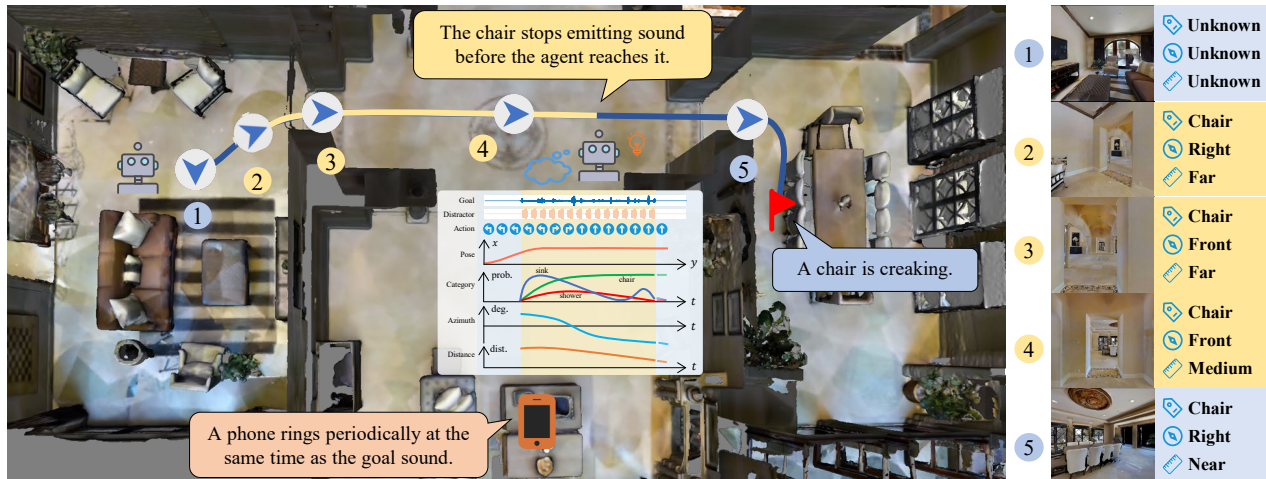


Figure 2. Overview of the proposed SAVN-CE framework. ① The agent is randomly initialized without prior knowledge of the environment or the goal. ② It explores the environment until the goal object (a chair) starts emitting sound. ③ Leveraging multimodal cues, the agent infers the goal’s semantic category, azimuth, and distance, and navigates toward it while avoiding obstacles and acoustic distractors (e.g., a ringing phone). As the agent approaches the goal (④ → ⑤), the sound emission ceases. Elements highlighted in yellow and blue denote sound-emitting and silent periods, respectively. ⑤ During the silent period, the agent maintains goal tracking by integrating historical goal representations with current self-motion cues (e.g., the previous action and the current pose), successfully localizing and reaching the creaking chair despite the presence of visually similar objects and distractor sounds.

environment without any goal information and, once the goal begins to emit sound, execute long-horizon navigation toward it while avoiding obstacles, as illustrated in Fig. 1(c).

The core challenge of SAVN-CE lies in accurately inferring both the spatial location and semantic category of the goal from partial sensory observations. During the sound-emitting period, the goal may temporarily become silent due to large temporal gaps (e.g., intermittent creaking sounds) and a finer simulation step (0.25 s in our setting), making it difficult to continuously estimate its position and category. Moreover, sound semantics are often ambiguous within short temporal windows and become distinguishable only when observed across longer temporal contexts. The situation is further exacerbated once the goal stops emitting sound, resulting in a sustained loss of goal information.

Taking these challenges into consideration, we propose MAGNet (*Memory-Augmented Goal descriptor Network*), a multimodal transformer-based architecture designed for reliable goal inference and efficient navigation. MAGNet jointly encodes spatial and semantic goal representations and integrates historical context with self-motion cues for continuous goal inference. By capturing temporal dependencies between past information and current sensory observations, MAGNet enables agents to navigate effectively toward the goal even after auditory signals are no longer available, as illustrated in Fig. 2.

For a fair comparison, we adapt and retrain existing audio-visual navigation methods in continuous environments and systematically evaluate them on our SAVN-CE dataset, which is built upon Matterport3D [8]. Experi-

mental results show that MAGNet significantly outperforms prior methods, particularly in challenging scenarios involving short-duration sounds and long-distance navigation.

In summary, our contributions are threefold:

- We introduce SAVN-CE, which extends semantic audio-visual navigation to continuous environments, bringing the task closer to realistic scenarios.
- We propose MAGNet, which leverages historical context and self-motion cues for robust goal reasoning and efficient navigation, thereby addressing the challenge of goal information loss when the goal sound is absent.
- Extensive experiments demonstrate that MAGNet significantly outperforms existing methods, achieving up to a 12.1% absolute improvement in success rate.

2. Related Work

Sound event localization and detection. Sound event localization and detection (SELD) [2] unifies two fundamental auditory perception tasks: sound event detection [7, 52] and sound source localization [1, 23, 39], thereby providing a comprehensive spatiotemporal representation of acoustic scenes. It jointly estimates the temporal boundaries, categories, and spatial positions of sound events within a unified vector framework [45, 46]. Recent advances have extended SELD to handle more challenging scenarios, including multiple and moving sound sources [46], and have also incorporated visual cues [6] to further enhance performance. However, most existing approaches rely on either simulated datasets [2, 28] or single-room recordings [47],

leaving realistic, acoustically complex multi-room environments largely underexplored.

Semantic audio-visual navigation. Although subsequent studies have extended AVN to more complex settings, such as sound-attacker [55], multi-goal [30], and moving-source scenarios [54], the standard formulation still exhibits several key limitations: 1) the goal emits sound continuously throughout the episode; 2) the goal’s position is placed arbitrarily in the scene, lacking any visual embodiment; and 3) agents are confined to predefined grid locations where precomputed RIRs are available. To mitigate the first two issues, SAVN [10] was introduced to enable agents to navigate toward semantically grounded sound-emitting objects. Subsequent efforts have further enriched this framework by incorporating language instructions [36, 38], leveraging large language models [53], or handling multiple sound sources [44]. However, the third limitation remains largely unsolved [13, 30], and no prior work has explored it within the SAVN paradigm, despite the capability of SoundSpaces 2.0 [12] to support continuous navigation in realistic 3D environments. The principal difficulty lies in the complexity of the training procedure, which is further aggravated by the limited simulation speed resulting from the high computational cost of binaural audio rendering.

VLN in continuous environments. Vision-and-language navigation in continuous environments (VLN-CE) was introduced by Krantz et al. [32] to eliminate the discrete graph assumptions of the original VLN setting [5]. In VLN-CE, agents must execute low-level actions within 3D continuous environments to reach a goal by following language instructions and visual observations, without access to global topology, oracle navigation, or perfect localization. Consequently, these requirements make the task substantially more challenging than discrete VLN, resulting in much lower absolute performance. To address these challenges, recent research has focused on waypoint prediction for long-range planning [3, 14, 27, 31, 33] and obstacle-avoidance strategies to prevent agents from getting stuck [3, 56]. In contrast, SAVN-CE centers on inferring goal information from partial sensory observations even after the goal sound has ceased.

3. SAVN in Continuous Environments

We propose SAVN-CE, a new task that requires embodied agents to navigate toward semantic sound-emitting goals in continuous 3D indoor environments using fine-grained, low-level actions. Unlike prior settings that rely on precomputed RIRs, which consume terabytes of storage, SAVN-CE allows agents to move freely in continuous spaces with temporally and spatially coherent audio rendered dynamically. This formulation demands that agents effectively integrate multimodal sensory inputs and perform long-horizon reasoning to reach the goal.

Simulator. SAVN-CE is implemented on SoundSpaces 2.0 [12], an extension of Habitat [41, 48] that supports continuous audio rendering within realistic Matterport3D scenes [8]. The simulator operates at a 16 kHz audio sampling rate with a 0.25 s simulation step, corresponding to 4,000 audio samples per step. Since the reverberation time of binaural RIRs rendered in Matterport3D scenes is typically much longer than a single simulation step, considering only the RIRs of the current and previous steps is insufficient to model long-tail reverberation effects (as done in SoundSpaces 2.0). To address this issue, following [17, 42], we convolve the source sound with the current-step binaural RIRs and accumulate the residual responses from all previous steps to generate temporally coherent audio.

Actions and observations. The agent’s action space comprises four discrete actions: *MoveForward 0.25 m*, *TurnLeft 15°*, *TurnRight 15°*, and *Stop*, consistent with VLN-CE [32]. Observations include binaural audio waveforms (mimicking human spatial hearing), egocentric RGB-D images (128×128 pixels, 90° field-of-view), and the agent’s pose relative to its initial position and orientation.

Dataset construction. Each episode is defined by: 1) the scene, 2) the agent’s initial location and orientation, 3) the goal’s location and semantic category, and 4) the onset time and duration of the goal sound. When a distractor is included, its location and category are also specified, and it shares identical temporal boundaries with the goal sound. The onset time of the goal sound is uniformly sampled from [0, 5] s, while its duration follows a Gaussian distribution with a mean of 15 s and a standard deviation of 9 s. We adapt the dataset from SAVi [10] to construct our SAVN-CE dataset, adopting the same 21 semantic categories as goal objects and 102 periodic sounds from SoundSpaces [9] as distractor candidates, disjoint from the goal categories. This setup ensures temporal diversity and acoustic ambiguity, making the task more challenging and realistic. Our dataset contains 0.5M/500/1,000 episodes for train/val/test, respectively. These splits use disjoint sets of scenes and source sounds, requiring agents to generalize to both unseen environments and unheard sounds. In the test split, the average number of oracle actions is 78.49, substantially higher than 26.52 in the discrete setting.

Success criterion. An episode is considered successful if the agent issues the *Stop* action within 1 m of the target sound source. Stopping near the distractor or another instance of the same semantic category is regarded as a failure. Following prior work in audio-visual navigation [9, 10], each episode is limited to at most 500 actions.

4. Method

To tackle the challenges of SAVN-CE, particularly in maintaining goal awareness when the goal sound becomes intermittent or silent, we propose MAGNet, a multimodal

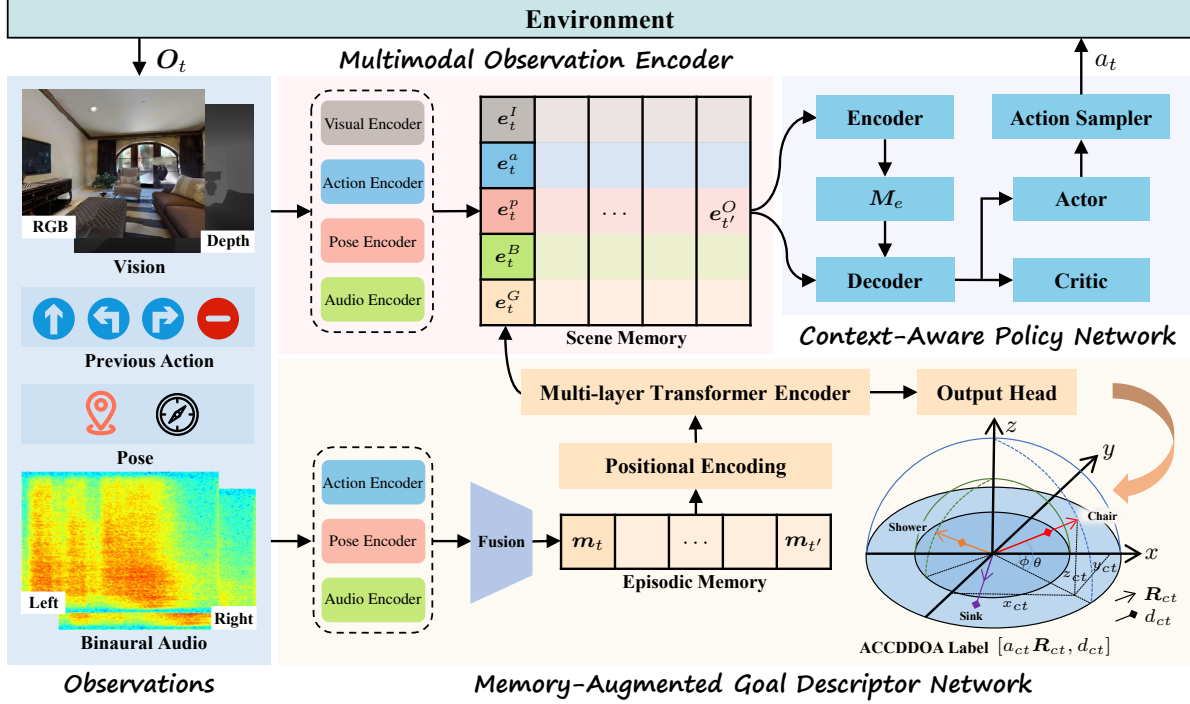


Figure 3. Overall architecture of MAGNet. The *multimodal observation encoder* extracts multimodal features from current sensory inputs and updates the scene memory accordingly. The *memory-augmented goal descriptor network* infers spatial and semantic representations of the goal by integrating auditory cues, self-motion cues, and historical goal embeddings stored in the episodic memory, thereby ensuring temporally consistent inference even after the goal sound ceases entirely. Conditioned on the latest scene memory embeddings, the *context-aware policy network* attends to the encoded memory M_e to predict the next action, enabling continuous navigation toward the goal.

transformer-based architecture to enable robust goal reasoning and efficient navigation. As illustrated in Fig. 3, MAGNet consists of three modules: 1) Multimodal Observation Encoder, which transforms multimodal inputs into compact embeddings and stores them in a long-term scene memory [10, 18]; 2) Memory-Augmented Goal Descriptor Network, which fuses auditory cues, egocentric motion information, and episodic memory to maintain a stable goal representation, ensuring persistent tracking even after the sound ceases; and 3) Context-Aware Policy Network, which attends to the aggregated scene memory to predict the next action of the agent.

4.1. Multimodal Observation Encoder

At time step t , the agent receives multimodal observations $O_t = \{I_t, a_{t-1}, p_t, B_t\}$, where I_t denotes the RGB-D images, a_{t-1} indicates the action taken at the previous time step, $p_t = [x_t, y_t, \theta_t, t]$ is the agent’s current pose, and B_t represents the binaural audio input. The observation encoder module consists of four modality-specific encoders (visual, action, pose, and audio), which process their corresponding inputs into embeddings. The concatenation of these embeddings constitutes the observation representation: $e_t^O = [e_t^I, e_t^a, e_t^p, e_t^B, e_t^G]$, where e_t^G is the goal em-

bedding described in Sec. 4.2. The scene memory maintains the most recent N_s encoded observations:

$$M_{s,t} = \{e_{t'}^O \mid t' \in [\max(0, t - N_s + 1), t]\}. \quad (1)$$

The visual encoder processes RGB and depth images using two independent ResNet-18 backbones [25], producing a concatenated embedding e_t^I . The previous action is mapped into the action embedding e_t^a via an embedding layer. The agent’s pose p_t is normalized to $[x_t/d, y_t/d, \sin(\theta_t), \cos(\theta_t), t/t_{\max}]$, where d is a distance scale factor and t_{\max} denotes the maximum episode length. The normalized pose is projected into the pose embedding e_t^p through a fully connected layer.

The binaural waveforms are transformed into complex spectrograms using a short-time Fourier transform (STFT) with a 512-point FFT and a hop length of 160 samples. To jointly encode spatial and semantic acoustic cues, we compute four complementary channels: the mean magnitude spectrogram, the sine and cosine components of the interaural phase difference, and the interaural level difference [34, 44]. The audio embedding e_t^B is extracted from these features by the audio encoder, which comprises three convolutional layers followed by a fully connected layer. See Supp. for details of the acoustic feature extraction.

4.2. Memory-Augmented Goal Descriptor Network

When the goal sound is intermittent or completely silent, maintaining a stable goal representation is essential for reliable navigation. To this end, we propose a memory-augmented goal descriptor network (GDN) that fuses binaural features, self-motion cues, and episodic memory to model temporal continuity and spatial dynamics explicitly. Unlike prior approaches that rely solely on current binaural audio, our design captures the evolving spatial relationship between the agent and the goal over time by combining auditory and self-motion cues, while leveraging episodic memory to maintain temporal continuity.

Self-motion cues, including the agent’s previous action and current pose, are crucial for estimating how the goal’s relative position changes as the agent moves. Specifically, the *TurnLeft* and *TurnRight* actions decrease and increase the goal’s azimuth relative to the agent by 15° , respectively, while the *MoveForward* action affects both azimuth and distance depending on the agent’s current position. The agent’s pose \mathbf{p}_t , encoding its translation and orientation relative to the initial state, further enhances spatial reasoning in binaural sound source localization [21, 34].

Since real-world acoustic events are often intermittent or ambiguous, the network must reason over historical auditory context rather than isolated observations. To capture such temporal dynamics, we adopt an episodic memory module [37], which stores goal-relevant embeddings from past steps. At time step t , the GDN receives the binaural audio \mathbf{B}_t , the previous action a_{t-1} , and the current pose \mathbf{p}_t . These encoders follow the same structures as in Sec. 4.1, except that the audio encoder produces a higher-dimensional embedding to better capture spatial and semantic goal information. The three embeddings are fused through a multi-layer perceptron (MLP) into a unified representation: $\mathbf{m}_t = \text{MLP}[e_t^a, e_t^p, e_t^B]$, which is then appended to the episodic memory:

$$\mathbf{M}_{g,t} = \{\mathbf{m}_{t'} \mid t' \in [\max(0, t - N_g + 1), t]\}, \quad (2)$$

where N_g is the episodic memory capacity. The collected episodic memory is augmented with positional encodings to preserve temporal order and then processed by a transformer encoder with two output branches.

The first branch projects the encoder output through a fully connected layer to obtain the goal embedding: $e_t^G = \text{FC}([\text{Encoder}(\mathbf{M}_{g,t})]_t)$, where the subscript t denotes the encoder output corresponding to the current time step. The second branch, in contrast, outputs goal descriptions in the activity-coupled Cartesian distance and direction-of-arrival (ACCDDOA) format [35, 45] using an MLP output head, which is employed for loss computation and network optimization during training. As illustrated in the bottom-right corner of Fig. 3, the ACCDDOA-formatted goal descriptions are formulated as: $\mathbf{y}_{ct} = [a_{ct}\mathbf{R}_{ct}, d_{ct}]$, where c and

t denote the category and time step indices, respectively. Here, $\mathbf{R}_{ct} = [x_{ct}, y_{ct}, z_{ct}] \in [-1, 1]^3$ represents the unit-norm direction-of-arrival (DOA) vector, $a_{ct} \in \{0, 1\}$ indicates the sound activity status (0 for inactive and 1 for active), and $d_{ct} > 0$ is the normalized distance.

By integrating self-motion dynamics and temporally accumulated goal embeddings, the memory-augmented GDN preserves consistent goal representations even in the absence of auditory input. The fine-grained action space further limits the positional changes between consecutive steps, ensuring stable and coherent goal tracking throughout long-horizon navigation in continuous environments.

4.3. Context-Aware Policy Network

The context-aware policy network employs a transformer-based encoder-decoder architecture to facilitate temporally informed decision-making by integrating both historical and current observations. At each time step t during an episode, the encoder processes the accumulated scene memory $\mathbf{M}_{s,t}$ to capture temporal dependencies across past observations, yielding an encoded representation $\mathbf{M}_e = \text{Encoder}(\mathbf{M}_{s,t})$. The decoder then generates a context-aware latent state representation:

$$\mathbf{s}_t = \text{Decoder}(\mathbf{M}_e, \text{MLP}(e_t^O)), \quad (3)$$

which serves as a compact summary of both historical and current sensory information. This representation is passed separately to an actor and a critic, each implemented as a fully connected layer that predicts the action distribution and the state value, respectively. Finally, an action sampler selects the next action a_t from the predicted distribution, enabling the agent to execute coherent and contextually grounded actions throughout the episode in a partially observable continuous environment.

4.4. Training Strategy

For stable and efficient training, each iteration consists of a 150-step rollout with the current policy network, followed by updates to both the GDN and the policy network using the collected experiences.

The GDN is trained online in a supervised manner using complete episodes. To construct these episodes, unfinished episodes from the previous iteration are merged with newly collected ones, ensuring full temporal continuity. Episodes shorter than 30 steps are discarded to guarantee that the goal has emitted sound by the end of the episode. The training procedure leverages oracle ACCDDOA labels with a mean squared error (MSE) loss and the Adam [29] optimizer at a learning rate of 1×10^{-3} . To maintain temporal causality, the encoder employs causal attention, preventing information leakage from future time steps.

The policy network is trained with decentralized distributed proximal policy optimization (DD-PPO) [50], fol-

lowing the two-stage paradigm used in SAVi [10]. It is optimized with the standard PPO loss [43] using the Adam [29] optimizer with a learning rate of 2.5×10^{-4} . The reward function has three components: a success reward of +10 for reaching the goal, an intermediate reward proportional to the change in geodesic distance to the goal, and a small time penalty of -0.01 per step to encourage efficient navigation.

5. Experiments

5.1. Experimental Setup

Baselines. We use the following methods for comparison:

- **Random:** A non-learning policy that samples actions according to the action distribution in the train split [32].
- **ObjectGoal:** A policy where the agent receives RGB-D observations and the ground-truth goal category, without access to a perfect *Stop* action.
- **AV-Nav:** A policy that leverages audio-visual input and employs a GRU to encode past observations [9].
- **SMT + Audio:** A policy that integrates audio with other sensory observations through a scene memory transformer (SMT) [18], without explicit goal inference.
- **SAVi:** A policy designed for the SAVN task, employing a GDN to independently infer the goal’s position and category from audio cues, and aggregating estimates over time using a weighting factor λ [10].
- **Oracle:** A variant of our architecture without the audio encoder or GDN, but with access to oracle ACCDDOA labels. Oracle1 has access only during the sound-emitting period, whereas Oracle2 retains access until the episode ends, providing upper-bound performance when goal information is lost or maintained after sound ceases.

Evaluation metrics. Navigation performance is evaluated using the following metrics: 1) **SR:** the fraction of successful episodes; 2) **SPL:** success weighted by the ratio of the shortest path length to the agent’s trajectory length [4]; 3) **SNA:** success weighted by the ratio of the oracle number of actions to the agent’s executed actions, which penalizes collisions and unnecessary rotations [11]; 4) **DTG:** the average geodesic distance to the goal at the end of the episode; 5) **SWS:** the fraction of successful episodes in which the agent reaches the goal while it is silent [10]. To assess the agent’s susceptibility to distractors, we introduce a diagnostic metric, **DSR** (Distractor Success Rate), which measures the fraction of episodes in which the agent reaches the distractor incorrectly. DTG is measured in meters, whereas all other metrics are reported as percentages.

Implementation details. All learning-based approaches are trained for up to 240M steps, with early stopping applied when no improvement is observed on the validation split for 24M consecutive steps. Training details follow those described in Sec. 4.4. The maximum capacities of the scene and episodic memories are set to $N_s = 150$ and $N_g = 128$,

respectively. The scale factor d is set to 20 m to normalize the agent’s pose and goal positions. For models using bin-aural audio, training the full 240M steps takes roughly 14 days on 128 CPU threads and 4 NVIDIA A800 GPUs.

5.2. Experimental Results

We consider two experimental scenarios: 1) *Clean Environments*, where only the goal emits sound; and 2) *Distracted Environments*, where an additional distractor emits sound simultaneously. Following prior work, we select the check-point achieving the highest SPL on the validation split for evaluation. Tab. 1 summarizes the navigation performance averaged over 1,000 test episodes and five independent runs. Overall, MAGNet consistently outperforms all baselines by a large margin across all metrics under *Clean Environments*, while the performance gain decreases under *Distracted Environments* due to interference from distractor sounds.

Performance under Clean Environments. Consistent with intuition, Random performs poorly, while ObjectGoal yields only marginal improvement, indicating that merely knowing the goal category without access to audio cues is insufficient for successful navigation. AV-Nav and SMT + Audio show notable improvements but remain limited, reflecting their constrained capacity to exploit audio cues. Among all baselines, SAVi achieves the best performance by explicitly inferring the goal’s position and category through its GDN. Nevertheless, MAGNet surpasses SAVi by a substantial margin across all metrics, highlighting the effectiveness of our memory-augmented GDN and multi-modal fusion strategy in enabling reliable goal reasoning and efficient navigation (see SWS and SPL metrics).

Performance under Distracted Environments. All methods exhibit degraded performance compared to *Clean Environments*, indicating that distractor sounds significantly hinder agents from focusing on goals. Although an additional ground-truth goal category is provided to help agents distinguish the goal from the distractor, they still struggle to locate it accurately. In particular, despite its excellent capability to reason about sound-emitting objects through the memory-augmented GDN, MAGNet achieves only modest performance gains while exhibiting the highest DSR when acoustically similar distractors are present.

Comparison with Oracle settings. To further contextualize performance, we compare MAGNet with its oracle variants. As shown in Tab. 1, Oracle2 achieves a substantial improvement over Oracle1, indicating that the limited duration of the goal sound significantly constrains navigation performance. This finding underscores the importance of accurately updating goal information once the goal sound ceases. Although MAGNet effectively maintains goal estimates after silence, a performance gap remains relative to Oracle1 and further widens against Oracle2, suggesting considerable room for improvement in the proposed GDN.

Table 1. Navigation performance on the SAVN-CE dataset under *Clean* and *Distracted Environments*.

	<i>Clean Environments</i>					<i>Distracted Environments</i>					
	SR \uparrow	SPL \uparrow	SNA \uparrow	DTG \downarrow	SWS \uparrow	SR \uparrow	SPL \uparrow	SNA \uparrow	DTG \downarrow	SWS \uparrow	DSR
Random	0.3	0.2	0.1	15.7	0.2	0.4	0.3	0.1	15.8	0.4	1.2
ObjectGoal	0.8	0.6	0.2	14.4	0.4	1.1	0.8	0.4	14.4	0.8	0.8
AV-Nav	21.3	17.8	13.1	10.7	4.0	13.5	10.9	7.6	12.5	3.0	6.7
SMT + Audio	24.8	21.0	16.8	10.1	5.3	16.5	13.7	9.4	11.8	5.0	6.3
SAVi	25.6	21.2	17.3	10.1	6.0	18.5	14.9	11.1	10.4	5.4	6.4
MAGNet (Ours)	37.7	32.9	27.4	8.0	10.6	19.3	16.5	13.4	10.6	4.8	7.8
Oracle1	41.4	37.8	31.0	6.3	13.0	41.5	38.0	31.0	6.3	11.2	0.5
Oracle2	75.0	63.7	51.9	4.2	48.4	74.3	64.0	51.8	4.5	45.4	0.8




Figure 4. Navigation trajectories of different methods under *Clean Environments*.

5.3. Visualization of Navigation Trajectories

We visualize the navigation trajectories of different methods under *Clean Environments* in Fig. 4. Intuitively, before the goal starts emitting sound, agents may either approach or diverge from it. In the latter case (first row of Fig. 4), additional actions are required within the limited sound-emitting window to compensate for earlier incorrect movements, leading to inefficient exploration. As shown in the second row of Fig. 4, agents tend to lose their way once the goal sound ceases, ultimately failing to reach the goal. In contrast, our method successfully completes the navigation task by leveraging historical context and self-motion cues to reliably maintain and update goal information. See Supp. for more visualizations and video demos.

5.4. Factors Affecting Navigation Success

To gain deeper insight into the factors that determine navigation success in SAVN-CE, we analyze two key variables that strongly influence SR: 1) **Action Ratio**: the ratio between the oracle number of actions required to reach the goal and the number of actions available during the goal’s sound-emitting period; 2) **Geodesic Distance**: the geodesic




Figure 5. Impact of (a) action ratio and (b) geodesic distance on the cumulative success rates of different methods under *Clean Environments* (solid) and *Distracted Environments* (dashed).

distance between the agent’s initial position and the goal. A higher action ratio indicates a larger portion of actions must be executed without the goal sound, making the episode inherently more challenging. Similarly, a larger geodesic distance increases task difficulty, as sound attenuation and complex spatial layouts hinder accurate goal inference.

The results presented in Fig. 5 reveal consistent trends across all evaluated methods. As either the action ratio or the geodesic distance increases, the cumulative success rate gradually saturates, indicating that most successful episodes

Table 2. Comparison of GDN performance for SAVi and MAGNet on the test split. *Sounding* and *Silent* denote metrics measured during the sound-emitting and silent periods, respectively. Note that a low LE_{CD} accompanied by an extremely low $F_{<20^\circ}$ is not informative.

Period	Method	Clean Environments					Distracted Environments				
		$ER_{<20^\circ}\downarrow$	$F_{<20^\circ}\uparrow$	$LE_{CD}\downarrow$	$LR_{CD}\uparrow$	$RDE\downarrow$	$ER_{<20^\circ}\downarrow$	$F_{<20^\circ}\uparrow$	$LE_{CD}\downarrow$	$LR_{CD}\uparrow$	$RDE\downarrow$
Sounding	SAVi	0.882	0.148	39.71	0.499	0.270	1.123	0.080	43.12	0.515	0.281
	MAGNet	0.762	0.290	36.77	0.601	0.117	0.896	0.119	47.84	0.733	0.319
Silent	SAVi	0.995	0.006	18.05	0.015	0.242	1.761	0.027	55.58	0.240	0.166
	MAGNet	0.905	0.140	48.83	0.368	0.166	0.892	0.143	50.65	0.508	0.207

Table 3. Ablation study of MAGNet on the SAVN-CE dataset under *Clean Environments*.

GDN	Memory	Self-motion	Navigation Metrics					SELD Metrics				
			SR \uparrow	SPL \uparrow	SNA \uparrow	DTG \downarrow	SWS \uparrow	$ER_{<20^\circ}\downarrow$	$F_{<20^\circ}\uparrow$	$LE_{CD}\downarrow$	$LR_{CD}\uparrow$	$RDE\downarrow$
\times	\times	\times	32.4	27.9	23.5	9.2	6.3	-	-	-	-	-
\checkmark	\times	\checkmark	33.9	29.8	24.8	8.8	8.9	0.953	0.082	26.59	0.122	0.128
\checkmark	\checkmark	\times	34.3	30.4	25.6	8.5	7.8	0.855	0.193	36.15	0.453	0.134
\checkmark	\checkmark	\checkmark	37.7	32.9	27.4	8.0	10.6	0.816	0.240	31.30	0.494	0.119

occur with longer sound durations and closer goal distances. Compared with the baselines, MAGNet consistently achieves higher upper bounds on the cumulative success rate, demonstrating superior robustness to short-duration sounds and long-distance navigation scenarios. This highlights the capability of the memory-augmented GDN to maintain stable goal representations.

5.5. Evaluation of the Goal Descriptor Network

Following standard SELD metrics [35, 40], we jointly assess the localization and detection performance of GDNs using the error rate and F1-score within a 20° localization tolerance ($ER_{<20^\circ}$ and $F_{<20^\circ}$), as well as the localization error, localization recall, and relative distance error conditioned on correctly detected events (LE_{CD} , LR_{CD} , and RDE). The results, macro-averaged across goal categories, are summarized in Tab. 2. See Supp. for a detailed analysis of the results and visualizations.

As shown in Tab. 2, SELD metrics notably decline under *Distracted Environments* compared with *Clean Environments*, indicating that distractor sounds substantially increase the difficulty of accurate goal localization and categorization. Our model outperforms SAVi on most metrics, demonstrating its superior capability to jointly infer spatial and semantic goal information. When the goal sound ceases, both methods experience a performance degradation; however, our memory-augmented GDN remains competitive, suggesting that it effectively preserves and updates goal representations even without auditory input. These observations reinforce the motivation presented in Sec. 4.2.

5.6. Ablation Study

To further quantify the contribution of each component in MAGNet, we conduct a series of ablation experiments. The results summarized in Tab. 3 report both navigation and

SELD performance under four configurations: 1) without the GDN, 2) with the GDN but no memory ($N_g = 1$), 3) with the GDN but without self-motion cues, and 4) the full memory-augmented GDN.

Removing the GDN leads to a notable degradation in navigation performance; however, the remaining modules still outperform all baselines, indicating their effectiveness in capturing informative features for navigation. Introducing either episodic memory or self-motion cues individually improves both navigation and SELD performance, highlighting their respective roles in modeling temporal context and spatial reasoning. Further improvements are achieved when all components are combined, demonstrating their complementary contributions within the memory-augmented GDN for robust navigation and SELD. We also observe that improved SELD performance consistently translates into better navigation outcomes.

6. Conclusion

In this paper, we introduce SAVN-CE, a new task that extends semantic audio-visual navigation to continuous environments, where agents navigate freely toward semantically grounded sound-emitting goals while avoiding distractors. To address the challenge of losing goal information when the goal becomes silent, we propose MAGNet, which employs a memory-augmented goal descriptor network for robust goal reasoning and efficient navigation. Comprehensive experiments demonstrate that MAGNet substantially outperforms existing state-of-the-art methods, exhibiting strong robustness to short-duration sounds and long-distance navigation. Ablation studies further validate the effectiveness of our proposed method. In future work, we plan to extend this framework to more complex auditory scenarios involving multiple or dynamic goals.

7. Acknowledgements

This work was supported in part by the Zhongguancun Academy (Grant No. 20240306), the National Natural Science Foundation of China (Grant Nos. 62471340, 62376266, 62376140, U23A20315), the Special Fund for Taishan Scholar Project of Shandong Province, and the Tsinghua University-Toyota Research Center.

References

- [1] Sharath Adavanne, Archontis Politis, and Tuomas Virtanen. Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network. In *Proc. EU-SIPCO*, pages 1462–1466, 2018. 2
- [2] Sharath Adavanne, Archontis Politis, Joonas Nikunen, and Tuomas Virtanen. Sound event localization and detection of overlapping sources using convolutional recurrent neural networks. *IEEE J. Sel. Topics Signal Process.*, 13(1):34–48, 2019. 2
- [3] Dong An, Hanqing Wang, Wenguan Wang, Zun Wang, Yan Huang, Keji He, and Liang Wang. ETPNav: Evolving topological planning for vision-language navigation in continuous environments. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 5130–5145, 2024. 3
- [4] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018. 6
- [5] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3674–3683, 2018. 1, 3
- [6] Davide Berghi, Peipei Wu, Jinzheng Zhao, Wenwu Wang, and Philip JB Jackson. Fusion of audio and visual embeddings for sound event localization and detection. In *Proc. IEEE ICASSP*, pages 8816–8820, 2024. 2
- [7] Emre Cakir, Giambattista Parascandolo, Toni Heittola, Heikki Huttunen, and Tuomas Virtanen. Convolutional recurrent neural networks for polyphonic sound event detection. *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 25(6):1291–1303, 2017. 2
- [8] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niebner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. In *Proc. Int. Conf. 3D Vis.*, pages 667–676, 2017. 2, 3
- [9] Changan Chen, Unnat Jain, Carl Schissler, Sebastia Vicenc Amengual Gari, Ziad Al-Halah, Vamsi Krishna Ithapu, Philip Robinson, and Kristen Grauman. SoundSpaces: Audio-visual navigation in 3D environments. In *Proc. Eur. Conf. Comput. Vis.*, pages 17–36, 2020. 1, 3, 6
- [10] Changan Chen, Ziad Al-Halah, and Kristen Grauman. Semantic audio-visual navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15516–15525, 2021. 1, 3, 4, 6
- [11] Changan Chen, Sagnik Majumder, Ziad Al-Halah, Ruohan Gao, Santhosh Kumar Ramakrishnan, and Kristen Grauman. Learning to set waypoints for audio-visual navigation. In *Proc. Int. Conf. Learn. Represent.*, pages 4861–4876, 2021. 6
- [12] Changan Chen, Carl Schissler, Sanchit Garg, Philip Kobernik, Alexander Clegg, Paul Calamia, Dhruv Batra, Philip Robinson, and Kristen Grauman. SoundSpaces 2.0: A simulation platform for visual-acoustic learning. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 8896–8911, 2022. 3
- [13] Changan Chen, Jordi Ramos, Anshul Tomar, and Kristen Grauman. Sim2real transfer for audio-visual navigation with frequency-adaptive acoustic field prediction. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 8595–8602, 2024. 3
- [14] Kehan Chen, Dong An, Yan Huang, Rongtao Xu, Yifei Su, Yonggen Ling, Ian Reid, and Liang Wang. Constraint-aware zero-shot vision-language navigation in continuous environments. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(11):10441–10456, 2025. 3
- [15] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 5834–5847, 2021. 1
- [16] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Think global, act local: Dual-scale graph transformer for vision-and-language navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16537–16547, 2022. 1
- [17] David Diaz-Guerra, Antonio Miguel, and Jose R Beltran. gpuRIR: A python library for room impulse response simulation with GPU acceleration. *Multimedia Tools Applic.*, 80(4):5653–5671, 2021. 3
- [18] Kuan Fang, Alexander Toshev, Li Fei-Fei, and Silvio Savarese. Scene memory transformer for embodied agents in long-horizon tasks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 538–547, 2019. 4, 6
- [19] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-follower models for vision-and-language navigation. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 3314–3325, 2018. 1
- [20] Chuang Gan, Yiwei Zhang, Jiajun Wu, Boqing Gong, and Joshua B Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. In *Proc. IEEE Int. Conf. Robotics Autom.*, pages 9701–9707, 2020. 1
- [21] Guillermo García-Barrios, Daniel Aleksander Krause, Archontis Politis, Annamaria Mesaros, Juana M Gutiérrez-Arriola, and Rubén Fraile. Binaural source localization using deep learning and head rotation information. In *Proc. EUSIPCO*, pages 36–40, 2022. 5
- [22] Georgios Georgakis, Karl Schmeckpeper, Karan Wanchoo, Soham Dan, Eleni Miltsakaki, Dan Roth, and Kostas Daniilidis. Cross-modal map learning for vision and language

- navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15460–15470, 2022. 1
- [23] Pierre-Amaury Grumiaux, Srđan Kitić, Laurent Girin, and Alexandre Guérin. A survey of sound source localization with deep learning methods. *J. Acoust. Soc. Am.*, 152(1): 107–151, 2022. 2
- [24] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2616–2625, 2017. 1
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. 4
- [26] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez-Opazo, and Stephen Gould. VLN-BERT: A recurrent vision-and-language bert for navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1643–1653, 2021. 1
- [27] Yicong Hong, Zun Wang, Qi Wu, and Stephen Gould. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15439–15449, 2022. 3
- [28] Jinbo Hu, Yin Cao, Ming Wu, Fang Kang, Feiran Yang, Wenwu Wang, Mark D Plumbley, and Jun Yang. PSELD-Nets: Pre-trained neural networks on a large-scale synthetic dataset for sound event localization and detection. *IEEE Trans. Audio, Speech, Lang. Process.*, 33:2845–2860, 2025. 2
- [29] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *Proc. Int. Conf. Learn. Represent.*, pages 1–15, 2015. 5, 6
- [30] Haru Kondoh and Asako Kanezaki. Multi-goal audio-visual navigation using sound direction map. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 5219–5226, 2023. 3
- [31] Jacob Krantz and Stefan Lee. Sim-2-Sim transfer for vision-and-language navigation in continuous environments. In *Proc. Eur. Conf. Comput. Vis.*, pages 588–603, 2022. 3
- [32] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the Nav-Graph: Vision-and-language navigation in continuous environments. In *Proc. Eur. Conf. Comput. Vis.*, pages 104–120, 2020. 3, 6
- [33] Jacob Krantz, Aaron Gokaslan, Dhruv Batra, Stefan Lee, and Oleksandr Maksymets. Waypoint models for instruction-guided navigation in continuous environments. In *Proc. Int. Conf. Comput. Vis.*, pages 15162–15171, 2021. 3
- [34] Daniel Aleksander Krause, Guillermo García-Barrios, Archontis Politis, and Annamaria Mesaros. Binaural sound source distance estimation and localization for a moving listener. *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 32: 996–1011, 2023. 4, 5
- [35] Daniel Aleksander Krause, Archontis Politis, and Annamaria Mesaros. Sound event detection and localization with distance estimation. In *Proc. EUSIPCO*, pages 286–290, 2024. 5, 8
- [36] Xiulong Liu, Sudipta Paul, Moitrey Chatterjee, and Anoop Cherian. CAVEN: An embodied conversational agent for efficient audio-visual navigation in noisy environments. In *Proc. AAAI*, pages 3765–3773, 2024. 3
- [37] Alexander Pashevich, Cordelia Schmid, and Chen Sun. Episodic transformer for vision-and-language navigation. In *Proc. Int. Conf. Comput. Vis.*, pages 15942–15952, 2021. 5
- [38] Sudipta Paul, Amit Roy-Chowdhury, and Anoop Cherian. AVLEN: Audio-visual-language embodied navigation in 3D environments. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 6236–6249, 2022. 3
- [39] Laureline Perotin, Romain Serizel, Emmanuel Vincent, and Alexandre Guerin. CRNN-based multiple DOA estimation using acoustic intensity features for Ambisonics recordings. *IEEE J. Sel. Topics Signal Process.*, 13(1):22–33, 2019. 2
- [40] Archontis Politis, Annamaria Mesaros, Sharath Adavanne, Toni Heittola, and Tuomas Virtanen. Overview and evaluation of sound event localization and detection in DCASE 2019. *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 29: 684–698, 2021. 8
- [41] 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. In *Proc. Int. Conf. Comput. Vis.*, pages 9339–9347, 2019. 3
- [42] Robin Scheibler, Eric Bezzam, and Ivan Dokmanić. Py-roomacoustics: A python package for audio room simulation and array processing algorithms. In *Proc. IEEE ICASSP*, pages 351–355, 2018. 3
- [43] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 6
- [44] Zhanbo Shi, Lin Zhang, Linfei Li, and Ying Shen. Towards audio-visual navigation in noisy environments: A large-scale benchmark dataset and an architecture considering multiple sound-sources. In *Proc. AAAI*, pages 14673–14680, 2025. 3, 4
- [45] Kazuki Shimada, Yuichiro Koyama, Naoya Takahashi, Shusuke Takahashi, and Yuki Mitsufuji. ACCDOA: Activity-coupled Cartesian direction of arrival representation for sound event localization and detection. In *Proc. IEEE ICASSP*, pages 915–919, 2021. 2, 5
- [46] Kazuki Shimada, Yuichiro Koyama, Shusuke Takahashi, Naoya Takahashi, Emiru Tsunoo, and Yuki Mitsufuji. Multi-ACCDOA: Localizing and detecting overlapping sounds from the same class with auxiliary duplicating permutation invariant training. In *Proc. IEEE ICASSP*, pages 316–320, 2022. 2
- [47] Kazuki Shimada, Archontis Politis, Parthasaarathy Sudarsanam, Daniel A. Krause, Kengo Uchida, Sharath Adavanne, Aapo Hakala, Yuichiro Koyama, Naoya Takahashi, Shusuke Takahashi, Tuomas Virtanen, and Yuki Mitsufuji. STARSS23: An audio-visual dataset of spatial recordings of real scenes with spatiotemporal annotations of sound events. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 72931–72957, 2023. 2
- [48] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, et al. Habitat 2.0: Training home assistants to rearrange their habitat. In *Proc. Adv. Neural Inform. Process. Syst.*, pages 251–266, 2021. 3

- [49] Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6629–6638, 2019. 1
- [50] Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *Proc. Int. Conf. Learn. Represent.*, pages 5648–5668, 2020. 5
- [51] Yi Wu, Yuxin Wu, Aviv Tamar, Stuart Russell, Georgia Gkioxari, and Yuandong Tian. Bayesian relational memory for semantic visual navigation. In *Proc. Int. Conf. Comput. Vis.*, pages 2769–2779, 2019. 1
- [52] Xianjun Xia, Roberto Togneri, Ferdous Sohel, Yuanjun Zhao, and Defeng Huang. Sound event detection using multiple optimized kernels. *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 28:1745–1754, 2020. 2
- [53] Zeyuan Yang, Jiageng Liu, Peihao Chen, Anoop Cherian, Tim K Marks, Jonathan Le Roux, and Chuang Gan. RILA: Reflective and imaginative language agent for zero-shot semantic audio-visual navigation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16251–16261, 2024. 3
- [54] Abdelrahman Younes, Daniel Honerkamp, Tim Welschehold, and Abhinav Valada. Catch me if you hear me: Audio-visual navigation in complex unmapped environments with moving sounds. *IEEE Robot. Autom. Lett.*, 8(2):928–935, 2023. 3
- [55] Yinfeng Yu, Wenbing Huang, Fuchun Sun, Changan Chen, Yikai Wang, and Xiaohong Liu. Sound adversarial audio-visual navigation. In *Proc. Int. Conf. Learn. Represent.*, pages 26129–26152, 2022. 3
- [56] Lu Yue, Dongliang Zhou, Liang Xie, Feitian Zhang, Ye Yan, and Erwei Yin. Safe-VLN: Collision avoidance for vision-and-language navigation of autonomous robots operating in continuous environments. *IEEE Robot. Autom. Lett.*, 9(6):4918–4925, 2024. 3
- [57] Yuke Zhu, Daniel Gordon, Eric Kolve, Dieter Fox, Li Fei-Fei, Abhinav Gupta, Roozbeh Mottaghi, and Ali Farhadi. Visual semantic planning using deep successor representations. In *Proc. Int. Conf. Comput. Vis.*, pages 483–492, 2017. 1
- [58] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *Proc. IEEE Int. Conf. Robotics Autom.*, pages 3357–3364, 2017. 1