

Residual Cross-Modal Fusion Networks for Audio-Visual Navigation

Yi Wang Yinfeng Yu*

School of Computer Science and Technology, Xinjiang University, Urumqi, China
Joint International Research Laboratory of Silk Road Multilingual Cognitive

Bin Ren

School of Mechatronic Engineering and Automation
Shanghai University, Shanghai, China

Abstract

Audio-visual embodied navigation aims to enable an agent to autonomously localize and reach a sound source in unseen 3D environments by leveraging auditory cues. The key challenge of this task lies in effectively modeling the interaction between heterogeneous features during multimodal fusion, so as to avoid single-modality dominance or information degradation, particularly in cross-domain scenarios. To address this, we propose a Cross-Modal Residual Fusion Network, which introduces bidirectional residual interactions between audio and visual streams to achieve complementary modeling and fine-grained alignment, while maintaining the independence of their representations. Unlike conventional methods that rely on simple concatenation or attention gating, CRFN explicitly models cross-modal interactions via residual connections and incorporates stabilization techniques to improve convergence and robustness. Experiments on the Replica and Matterport3D datasets demonstrate that CRFN significantly outperforms state-of-the-art fusion baselines and achieves stronger cross-domain generalization. Notably, our experiments also reveal that agents exhibit differentiated modality dependence across different datasets. The discovery of this phenomenon provides a new perspective for understanding the cross-modal collaboration mechanism of embodied agents.

Keywords: Audio-Visual Navigation, Multimodal Fusion, Residual Network, Embodied Intelligence

1. Introduction

With the rapid development of embodied intelligence [11, 13, 26, 35, 34], enabling agents to accomplish

*Yinfeng Yu is the corresponding author (Email: yuyinfeng@xju.edu.cn)

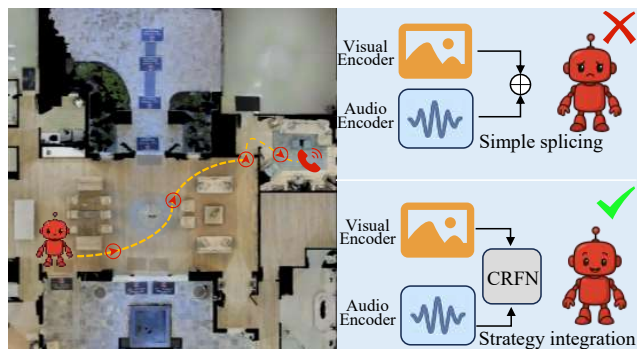


Figure 1. Humans can effortlessly integrate visual and auditory information, yet agents often suffer from performance degradation due to modal imbalance or domain discrepancy.

complex tasks in real or simulated environments has become a critical research direction in the field of artificial intelligence [15, 25]. Among related tasks, audio-visual navigation (AVN) is a key task for embodied intelligence, which requires agents to autonomously locate and navigate to target sound sources in unknown environments by leveraging both visual information and auditory cues. Such tasks hold broad value in application scenarios, including service robotics, human-computer interaction, and security monitoring, thus attracting increasing attention.

In a bustling, noisy shopping mall, you faintly hear someone calling your name. Your ears first pick up the sound’s general direction; you then turn your head and scan the crowd with your gaze. Only when you catch sight of a familiar figure do you confirm the caller’s position. Amid this chaotic setting, hearing narrows down the search range for you, while vision finalizes the identification and locking in of the target. The two work in seamless harmony, requiring hardly any deliberate effort on your part.

This kind of sensory synergy, which seems effortless to humans, poses a significant challenge for embodied agents: while these agents can acquire both visual and au-

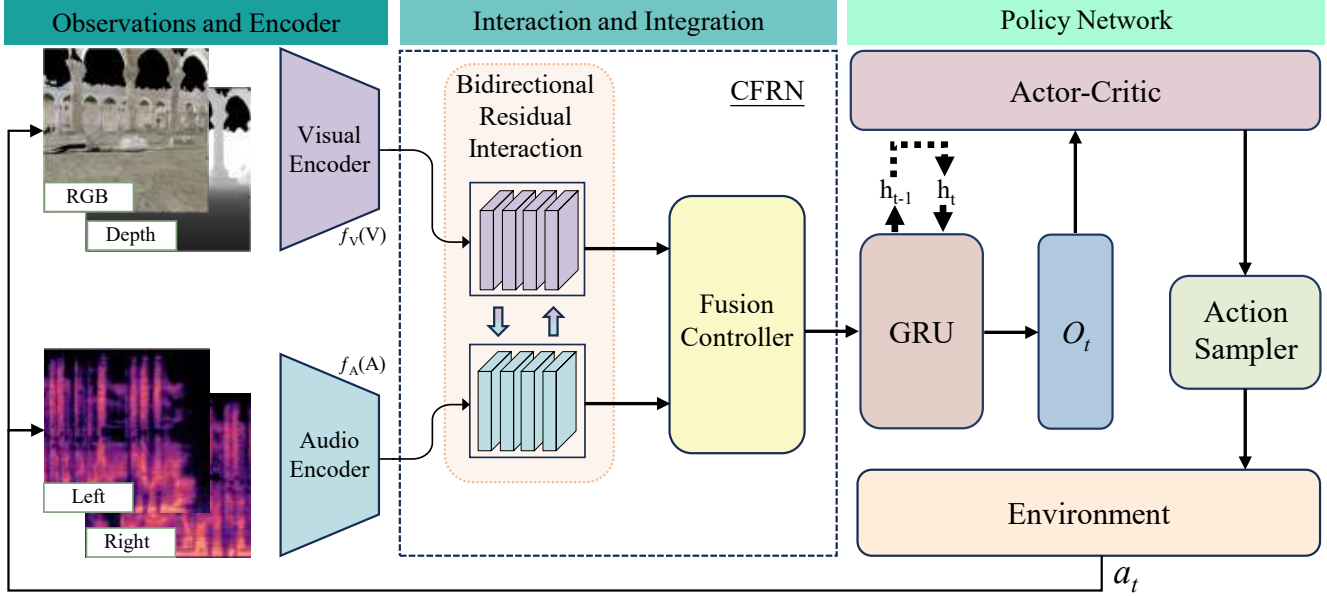


Figure 2. **Illustration of the basic architecture.** The agent’s navigation process operates in three stages: (1) Observation and Encoder, where the agent processes visual (RGB or Depth) and auditory (spectrogram) inputs via respective encoders to extract features; (2) Interaction and Integration, where our proposed CRFN module performs bidirectional residual updates to refine features and adaptively balances them using a Fusion Controller; and (3) Policy Network, where the fused representation is fed into a GRU-based Actor-Critic model to capture temporal dependencies and predict the final navigation action a_t .

ditary information, the two modalities often operate in isolation, lacking genuine complementarity and interaction—this leads to a substantial decline in navigation performance in complex environments. On the other hand, an interesting phenomenon was also observed in the experiments of this paper: the fusion effect exhibits significant differences across different environments. In synthetic high-fidelity scenes (e.g., Replica [23]), the visual modality often becomes dominant due to its clearer texture and geometric information, gradually weakening the audio modality. In contrast, in complex real-world scenes (e.g., Matterport3D [6]), a single modality can hardly complete the navigation task independently, and cross-modal complementarity instead becomes the key to performance improvement. This dynamic shift between modality dominance and cross-modal complementarity reveals the instability of existing methods in cross-environment and cross-domain generalization, and further underscores the necessity of designing a robust fusion mechanism.

To address these challenges, this paper introduces a novel Cross-Modal Residual Fusion Network (CRFN). The proposed method is built upon two core design principles: bidirectional residual interaction and a fusion control mechanism. Specifically, reciprocal residual pathways are established between audio and visual features to enable mutual refinement and complementarity while preserving the independence of each modality. In addition, a lightweight fusion controller is incorporated to adaptively regulate modal-

ity contributions during interaction, with an output normalization constraint applied to suppress single-modality dominance and stabilize training. This modular design not only enhances the robustness and generalization ability of the fusion process but also provides stronger interpretability: the evolution of residual coefficients offers an intuitive view into the dynamic collaboration between modalities.

In summary, our main contributions are as follows:

- *A novel cross-modal fusion framework.* We propose the Cross-Modal Residual Fusion Network (CRFN), which introduces bidirectional residual interactions and a lightweight fusion controller to explicitly model reciprocal refinement between audio and visual modalities while preserving their independence.
- *Improved robustness and generalization.* By combining small-value initialization of residual scaling factors with output normalization, CRFN suppresses modality imbalance, stabilizes training, and consistently outperforms concatenation and attention-based baselines on both Replica [23] and Matterport3D [6] benchmarks.
- *New empirical insights into modality dependence.* Through the analysis of residual coefficients, we uncover a novel phenomenon: in synthetic high-fidelity environments, the navigation policy tends to be dominated by the visual modality, whereas in real-world complex environments, cross-modal complementarity

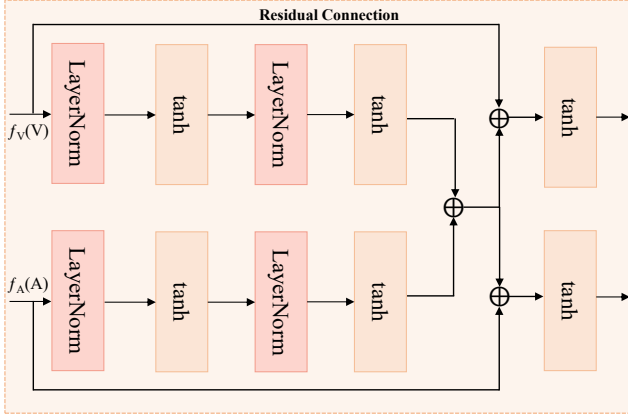


Figure 3. **Architecture of the Cross-Modal Feature Fusion Module.** This module enables visual and audio features to refine and complement each other through bidirectional residual paths. Features of each modality are updated and mutually influenced in the bidirectional interaction, ensuring a balanced exchange of information.

becomes essential. This provides a new perspective for understanding modality collaboration in embodied audio-visual navigation.

2. Related Work

Embodied intelligence aims to enable agents to learn and accomplish complex tasks through interaction with the environment. Navigation is one of its core capabilities, and early research has primarily focused on visual navigation, such as PointGoal navigation [2] and ObjectGoal navigation [4]. In these tasks, goals are typically specified by visual or textual information, such as coordinates or object categories. While the aforementioned works have achieved significant progress in visual navigation, they limit the agent’s perception to a single visual modality, which is far from the way humans and animals perceive the real world using multi-sensory synergy. In real physical environments, auditory perception provides a crucial information channel that complements vision. Sound signals possess the property of omni-directionality, allowing them to penetrate visual occlusions and provide key directional cues for targets outside the agent’s field of view. To endow embodied agents with such more powerful and biologically intuitive perceptual and action capabilities, AVN has been proposed.

The formal definition of the AVN task and the establishment of large-scale benchmarks mark the starting point for the development of this field. Chen et al. have made pioneering contributions in this regard: based on the Habitat [19, 24] platform, they constructed SoundSpaces [9] on the Replica [23] and Matterport3D [6] datasets, which is the first large-scale and high-fidelity audio-visual simulation

environment and provides a unified experimental platform and testing benchmark for subsequent research. This work has established the core problems and evaluation criteria for the AVN field; however, its fusion method is relatively elementary, making it difficult to address modal imbalance in complex environments. In parallel, Gan et al. explored a decomposed modular framework, attempting to decouple the perception, localization, and planning components to enhance interpretability and scalability [14]. These two types of methods represent the two early mainstream approaches—end-to-end learning and modular planning—respectively, and together they have established the research baseline and development direction for the AVN field.

Early AVN frameworks exhibited low efficiency when handling long-range navigation tasks [30]. To address this, AV-WaN [10] introduced an “acoustic memory map” for the first time, which structurally records sound information perceived by the agent during movement. This work innovatively proposed a hierarchical reinforcement learning framework: the high-level policy dynamically predicts mid-to-long-term “waypoints” based on multi-modal perception, while the low-level planner executes specific paths to reach these waypoints. Subsequently, SAVi [8] further pointed out the limitations of existing tasks—namely, the assumption that sound sources emit continuously and lack semantic relevance. Accordingly, they proposed a more challenging semantic AVN task, where sounds are sporadic and transient, and semantically consistent with the category of objects producing the sounds (e.g., the flushing sound of a toilet).

As fundamental issues have been addressed, researchers in the AVN field have shifted their focus toward more realistic and complex real-world scenarios. SAAVN [32, 31, 16] took the lead in investigating the navigation robustness in acoustically complex environments: it designed an adversarial environment that incorporates a “sound attacker,” which actively moves and adjusts both volume and sound category to interfere with the navigating agent. Later, a benchmark for dynamic AVN was proposed, marking the first time an agent was required to track a moving sound source [28]. Almost simultaneously, FSAAVN [29] also focused on the moving sound source tracking task and pointed out that the simple feature concatenation-based fusion methods used in previous works might overlook contextual information, and this was also the first time researchers paid attention to the issue of audio-visual cross-modal fusion since the inception of the AVN task. Subsequently, ORAN [12] leveraged Cross-Task Curriculum Policy Distillation to transfer the knowledge of a pre-trained “expert” policy (trained on the point-goal navigation task) to the AVN task. CAVEN [18] explored the role of human-computer interaction in AVN, introducing human interaction, large language models, and a new benchmark.

RILA [27] made the first attempt to solve the AVN task in a zero-shot manner, completely eliminating the need for reinforcement learning training in the environment. Shi et al. contributed BeDAVIN, a large-scale audio dataset, as well as ENMUS³-an, an architecture specifically designed for multi-sound-source scenarios [20].

Regrettably, although AVN is inherently a cross-modal task [33], the focus of existing research has largely been on innovations in navigation strategies and memory mechanisms. Its core audio-visual fusion module is to a large extent treated as a standardized component, and its inherent synergy potential has not been fully exploited.

3. Proposed Method

3.1. Framework Overview

As shown in the figure, the agent receives multimodal observations from the 3D environment, including visual images and audio signals. The visual modality consists of either depth maps or RGB images (one of which is used in the task setting), while the audio modality is represented as binaural spectrograms. First, visual and audio data are processed by dedicated encoders to extract modality-specific features, which are then fed into the proposed CRFN network. The primary role of CRFN is to establish bidirectional residual interactions and maintain a dynamic balance between modalities within a shared feature space, thereby producing robust fused representations. The fused features are subsequently passed into a recurrent unit (GRU) to capture historical information and temporal dependencies [21]. Finally, an Actor-Critic policy head leverages the temporal representation to predict action distributions and state values, driving the agent to accomplish goal-directed navigation tasks.

To achieve robust and stable cross-modal fusion, the CRFN proposed in this paper consists of two main modules: a Modal Interaction Module and a Fusion Control Module. The Modal Interaction Module explicitly models the mutual interaction between visual and audio features through bidirectional residual paths, ensuring fine-grained modal complementarity and representation alignment. The Fusion Control Module, on the other hand, adaptively adjusts the contribution ratio of each modality via learnable residual scaling factors and normalization constraints, effectively suppressing the dominance of a single modality and stabilizing the training process. The former is responsible for establishing efficient information exchange channels, while the latter acts as a “traffic controller” to dynamically regulate the intensity of such exchange, thereby achieving robust navigation performance in variable scenarios. CRFN adopts a concise and lightweight architecture, avoiding redundant network layers and complex structures to ensure efficient computation and training. With this de-

Algorithm 1: Fusion Controller in Cross-Modal Residual Fusion Network (CRFN)

Input: Visual feature \mathbf{v}_t , Audio feature \mathbf{a}_t ,

Interaction vector $\mathbf{h}_{interact}$

Output: Updated features $\hat{\mathbf{v}}_t, \hat{\mathbf{a}}_t$

Learnable parameters: β_v, β_a

Step 1: Normalize modal features

$\mathbf{v}_{norm} \leftarrow \text{LayerNorm}(\mathbf{v}_t)$

$\mathbf{a}_{norm} \leftarrow \text{LayerNorm}(\mathbf{a}_t)$

Step 2: Compute residual updates

$\mathbf{v}_{res} \leftarrow \mathbf{v}_{norm} + \beta_v \cdot \mathbf{h}_{interact}$

$\mathbf{a}_{res} \leftarrow \mathbf{a}_{norm} + \beta_a \cdot \mathbf{h}_{interact}$

Step 3: Apply activation to obtain final features

$\hat{\mathbf{v}}_t \leftarrow \tanh(\mathbf{v}_{res})$

$\hat{\mathbf{a}}_t \leftarrow \tanh(\mathbf{a}_{res})$

return $\hat{\mathbf{v}}_t, \hat{\mathbf{a}}_t$

sign, the model can efficiently achieve cross-modal information fusion and deliver excellent performance in audio-visual navigation tasks. Sections 3.2 and 3.3 below will delve into the design details and implementation methods of these two modules, respectively.

3.2. Bidirectional Residual Interaction

This module aims to construct a symmetric information pathway, enabling the feature representations of the two modalities to complement and refine each other. Specifically, at each time step t , the visual features $v_t(f_v(V))$ and audio features $a_t(f_a(A))$ obtained through encoders first pass through their respective non-linear transformation networks $U_v(\cdot)$ and $U_a(\cdot)$ to be mapped to a shared representation space that facilitates interaction. Subsequently, we average the transformed features to generate an interaction vector $h_{interact}$ that integrates information from both modalities:

$$h_{interact} = \frac{1}{2}(U_v(v_t) + U_a(a_t)), \quad (1)$$

This interaction vector $h_{interact}$ captures the core associations between audio and visual modalities and serves as a shared residual signal, which will be used to simultaneously update the features of both original modalities. This symmetric residual design ensures bidirectional flow of information between the two modalities.

3.3. Fusion Controller

The core of the fusion control module lies in adaptively adjusting modal contributions and ensuring the stability of the training process. This function is not implemented via a complex dynamic network, but rather accomplished by two

Table 1. Performance comparison with other methods under the Depth setting. SPL, SR, SNA are percentages.

Method	Replica						Matterport3D					
	Heard			Unheard sound			Heard			Unheard sound		
	SPL↑	SR↑	SNA↑	SPL↑	SR↑	SNA↑	SPL↑	SR↑	SNA↑	SPL↑	SR↑	SNA↑
Random [10]	4.9	18.5	1.8	4.9	18.5	1.8	2.1	9.1	0.8	2.1	9.1	0.8
Direction Follower [10]	54.7	72.0	41.1	11.1	17.2	8.4	32.3	41.2	23.8	13.9	18.0	10.7
Frontier W [10]	44.0	63.9	35.2	6.5	14.8	5.1	30.6	42.8	22.2	10.9	16.4	8.1
Supervised W [10]	59.1	88.1	48.5	14.1	43.1	10.1	21.0	36.2	16.2	4.1	8.8	2.9
Gan et al. [14]	57.6	83.1	47.9	7.5	15.7	5.7	22.8	37.9	17.1	5.0	10.2	3.6
SoundSpaces [9]	74.4	91.4	48.1	34.7	50.9	16.7	54.3	67.7	31.3	21.9	33.5	10.4
CRFN (Ours)	76.7	93.1	47.3	41.6	55.7	22.5	57.3	70.3	33.2	27.7	40.1	13.5

independent, learnable scalar parameters—residual scaling factors β_v and β_a . These two parameters control the update intensity of the interaction vector $\mathbf{h}_{interact}$ on visual and audio features, respectively.

We substitute these two factors into the residual update path to obtain the updated features β_v and β_a :

$$\hat{\mathbf{v}}_t = \text{act}(\text{LN}(\mathbf{v}_t) + \beta_v \cdot \mathbf{h}_{interact}), \quad (2)$$

$$\hat{\mathbf{a}}_t = \text{act}(\text{LN}(\mathbf{a}_t) + \beta_a \cdot \mathbf{h}_{interact}), \quad (3)$$

Among them, $\text{LN}(\cdot)$ refers to layer normalization applied to the residual path, which is used to stabilize the input scale of each modality; $\text{act}(\cdot)$ denotes the Tanh activation function. At the start of training, β_v and β_a are initialized to a small value. This initialization encourages the model to start learning from a “weakly coupled” state, which is more conducive for the network to independently discover the appropriate modal fusion intensity based on task requirements.

The overall computational procedure is summarized in Algorithm 1, providing a concise representation of how the residual scaling factors β_v and β_a operate within the CRFN architecture.

4. Experiments

4.1. Datasets and Implementation Detail

Our experiments were conducted in a comprehensive simulated environment, which is built on the Habitat simulator [19] and SoundSpaces [9] acoustic platform, and leverages two complementary public 3D datasets: Matterport3D [6] and Replica [23]. Among them, Matterport3D [6] is a large-scale, diverse dataset of real-world scans, containing 85 complex indoor environments. It serves as an ideal platform for evaluating the generalization

ability of agents, but its models contain geometric noise inherent to real-world scanning. In contrast, Replica [23] is a small yet high-quality dataset, comprising 18 high-fidelity synthetic indoor scenes. Renowned for its clean and realistic 3D meshes, Replica [23] is highly suitable for validating the effectiveness of models in high-fidelity environments. Using the geometric and material information from these two datasets, the SoundSpaces [9] platform generates physically realistic sound sources by convolving 102 non-repetitive natural sounds with binaural Room Impulse Responses (RIRs) corresponding to specific directions.

Implementation Details. The model parameter configuration is as follows: the sampling rate of Room Impulse Response (RIR) is 44,100 Hz for the Replica [23] dataset and 16,000 Hz for the Matterport3D [6] dataset. In all experiments, the resolution of both visual and auditory observations is 128×128 . We adopt the Proximal Policy Optimization (PPO) algorithm for training, and its hyperparameters remain consistent across the two datasets: 4 training epochs, a clip parameter of 0.1, a value loss coefficient of 0.5, and a discount factor (γ) set to 0.99. For experiments on the Replica [23] dataset, the learning rate is set to 2.5×10^{-4} , while for Matterport3D [6], the learning rate is 2.0×10^{-4} . During training, the number of steps per episode is limited to 500, and the model undergoes 40,000 update iterations on the Replica dataset and 60,000 update iterations on the Matterport3D [6] dataset.

In our experiments, there are two distinct sound source conditions: (1) Heard: the target sound source is a telephone ringtone, which is used consistently across the training, validation, and test sets; (2) Unheard: the 102 sound sources are divided into three non-overlapping groups, where 78 sound sources are used for training scenarios, 11 for validation scenarios, and the remaining 18 for test scenarios. Every scenario in the test process is new to the model—none of them were encountered in prior training or validation stages.

4.2. Evaluation Metrics

We follow the standard protocol [9] and adopt three metrics to compare the navigation performance of different methods: (1) Success weighted by Path Length (SPL) 6: After an agent successfully navigates to the target sound source, this metric measures how close the agent’s actual travel path is to the shortest feasible path [1]; (2) Success Rate (SR) 4: The proportion of test samples in which the agent successfully reaches the target location; (3) Success-Navigation Accuracy (SNA) 5: This metric evaluates the agent’s ability to not only successfully reach the target during navigation but also maintain orientation toward the correct direction.

$$SR = \frac{1}{N} \sum_{i=1}^N S_i, \quad (4)$$

$$SNA = \frac{1}{N} \sum_{i=1}^N S_i \cdot \frac{l_i}{a_i} [10], \quad (5)$$

$$SPL = \frac{1}{N} \sum_{i=1}^N S_i \cdot \frac{l_i}{\max(p_i, l_i)} [1]. \quad (6)$$

In the formulas, N is the total number of test episodes; $S_i \in \{0, 1\}$ specifies whether the i -th episode is successful; l_i is the shortest feasible path length for episode i ; p_i is the actual path length traveled by the agent; and a_i is the number of actions executed in episode i . The term a_i includes inefficient operations such as in-place rotations, so SNA penalizes excessive redundant movements.

4.3. Baselines

We compare our method with the following approaches:

- **Random Agent:** this is the simplest task in navigation, where the agent randomly selects actions (turn left, turn right, move forward) at each time step to search for the target.
- **Direction Follower:** this method adopts a hierarchical navigation framework, where the high-level policy predicts the Direction of Arrival of the sound source relying solely on auditory signals, and sets a waypoint at a fixed distance (K meters) along this direction. After the agent reaches the current waypoint, it repeats this process to gradually approach the target.
- **Frontier Waypoints:** the hierarchical baseline selects the next waypoint at the intersection of the predicted DoA and the current exploration frontier. Frontier-based waypoints are standard in visual navigation [5, 22, 7], making this a representative baseline of standard practice.
- **Supervised Waypoints:** the hierarchical baseline takes the RGB image and audio spectrogram as inputs and, using supervised (non-end-to-end) training, predicts waypoints within the field of view. Its design follows the supervised waypoint predictor of Bansal et al [3].
- **Gan et al. [14]** : one AVN method developed based on the AI2-THOR platform [17]. As a representative work in the early stage of the AVN field, its core advantage lies in the first systematic implementation of collaborative modeling between visual spatial memory and auditory sound source localization, providing a benchmark paradigm for subsequent related research.
- **SoundSpaces [9]:** the first publicly available audio-visual embodied navigation simulation platform. Our work is based on their code.

4.4. Quantitative Experimental Results

Table 1 presents a comparison of the proposed method (CRFN) with representative baselines on the Replica [23] and Matterport3D [6] datasets under the Depth setting. As shown, CRFN consistently achieves the best performance across all evaluation metrics, demonstrating its effectiveness in cross-modal feature integration and robust audio-visual navigation. On the Replica dataset, CRFN surpasses all prior methods in both Heard and Unheard settings. Specifically, in the Heard scenario, the SPL improves from 74.4 to 76.7, and SR rises to 93.1, indicating that the residual cross-modal fusion effectively enhances audio-visual synergy. In the Unheard scenario, the improvement becomes more pronounced—SPL reaches 41.6, about 7.6 percentage points higher than SoundSpaces [9]—showing stronger generalization and robustness to unseen sound sources. Likewise, on the more complex Matterport3D dataset, CRFN achieves the highest scores in both settings, with SPL = 57.3 and SR = 70.3 in the Heard case, significantly outperforming existing approaches. These results confirm that CRFN better balances the complementary relationship between visual and auditory signals in realistic environments, leading to improved multimodal perception, decision-making, and generalization in unseen scenes.

4.5. Qualitative Experimental Results

Fig. 4 and 5 provide a visual comparison of the top-down navigation trajectories between our CRFN model and the SoundSpaces baseline across a series of typical scenes. Specifically, Fig. 4 showcases two representative scenes from the Replica dataset, while Fig. 5 presents four complex scenes from the Matterport3D [6] dataset. From these trajectories, it can be clearly observed that the navigation paths generated by CRFN are significantly superior to those




Figure 4. Navigation trajectories on the top-down map in the Replica scenes. Agent paths transition from dark to light blue temporally, while green indicates the shortest geodesic path.

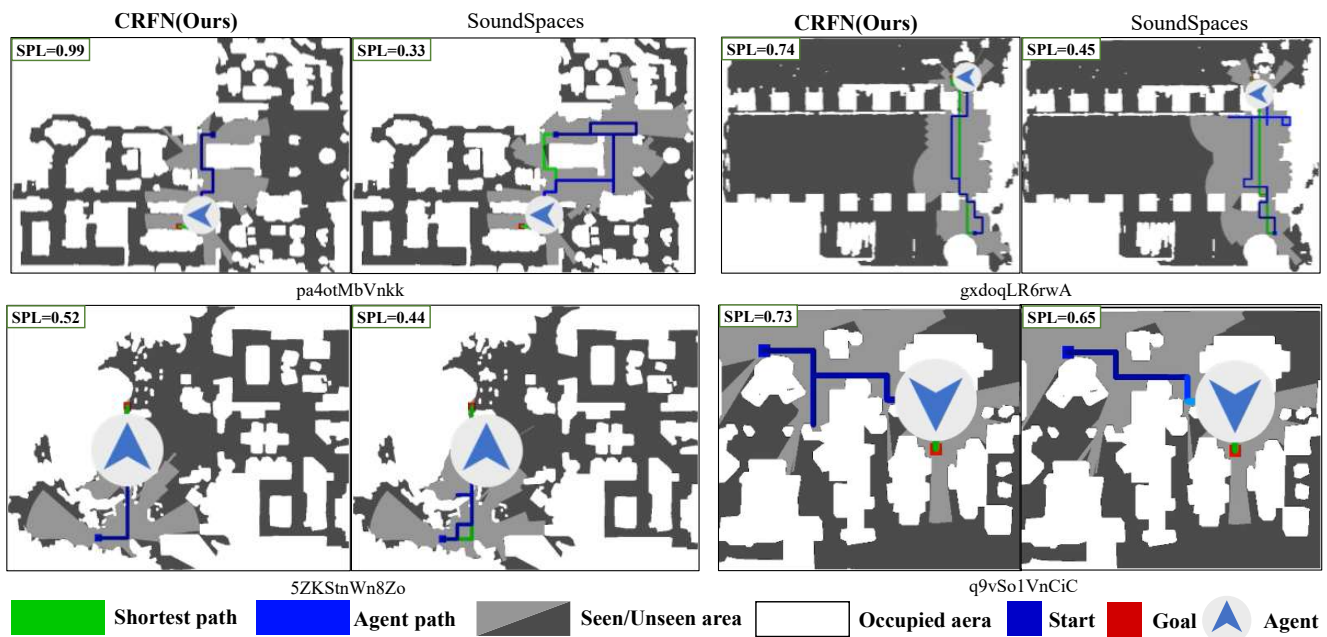


Figure 5. Navigation trajectories on the top-down map in the Matterport3D scenes. Agent paths transition from dark to light blue temporally, while green indicates the shortest geodesic path.

of the baseline, whether in the structurally simpler scenes of Replica (e.g., *apartment1* and *apartment4* in Fig. 4) or in the more complex layouts of Matterport3D [6]. CRFN’s trajectories adhere more closely to the shortest path (green), exhibiting less invalid exploration and backtracking, which is also corroborated by its higher SPL scores. Particularly in the task with scene ID *pa4otMbVnkk*, CRFN almost perfectly replicates the shortest path (SPL=0.99), whereas

the SoundSpaces [9] baseline, unable to establish a stable audio-visual correspondence, wanders extensively and inefficiently within the environment, resulting in extremely low path efficiency (SPL=0.33). This indicates that our model can efficiently fuse cross-modal information to quickly lock onto the target’s bearing even in acoustically simple, open spaces, thereby suppressing aimless exploration. Furthermore, in scenarios like *apartment4*, where the sound source




Figure 6. **Variation of Modal Weights with Training Progress.** In the figure, the curves represent the evolutionary trends of the visual update weight (β_v) and auditory update weight (β_a) on the Replica (left) and Matterport3D (right) datasets, respectively.

Table 2. Ablation study on the Fusion Controller of CRFN on the Replica dataset.

Method	Telephone			Unheard		
	SPL \uparrow	SR \uparrow	SNA \uparrow	SPL \uparrow	SR \uparrow	SNA \uparrow
w/o FC	69.3	84.9	44.4	34.0	43.5	19.2
CRFN (Ours)	76.7	93.1	47.3	41.6	55.7	22.5

Table 3. Ablation study on the Fusion Controller of CRFN on the Matterport3D dataset.

Method	Telephone			Unheard		
	SPL \uparrow	SR \uparrow	SNA \uparrow	SPL \uparrow	SR \uparrow	SNA \uparrow
w/o FC	53.1	65.9	32.2	28.9	37.7	14.7
CRFN (Ours)	57.3	70.3	33.2	27.7	40.1	13.5

may be partially occluded by walls, CRFN is still able to progressively correct its course and successfully complete the navigation, while the baseline model exhibits greater uncertainty, leading to lower path efficiency.

4.6. Modal Dependence Analysis

To further reveal the internal learning dynamics of our model across different environments, we visualize the temporal evolution of the update weights for the visual (β_v) and auditory (β_a) modalities during training. As shown in Fig. 6, we observe a notable phenomenon: rather than fusing multimodal information in a static manner, the model adaptively adjusts its modality dependency according to the structural complexity and perceptual statistics of the environment, exhibiting a dynamic balancing strategy between modalities.

Specifically, on the Replica [23] dataset, the model progressively reduces its reliance on the auditory modality (β_a) in the later stages of training, eventually converging to a vi-

sually dominated navigation strategy. This behavior can be attributed to the relatively small scale and regular structure of Replica [23] environments, where the agent can easily memorize spatial layouts through visual cues. Once the visual signal alone becomes sufficient for localization and planning, the model naturally suppresses the contribution of the temporally uncertain audio input, leading to a more efficient unimodal decision process.

In contrast, on the Matterport3D [6] dataset, the model maintains a balanced dependence on both modalities throughout training. Matterport3D [6] contains large-scale, cluttered real-world scenes with complex geometry, occlusions, and reverberations, where a single modality provides insufficient information for reliable navigation. In this case, the audio modality offers complementary global directional cues that can penetrate obstacles, compensating for the limitations of visual perception. This results in a more robust cross-modal collaboration, highlighting the model’s inherent ability to adaptively modulate modality weights across diverse environments.

4.7. Ablation Studies

To evaluate the effectiveness of the Fusion Controller (FC) in our CRFN, we conducted ablation experiments on both the Replica [23] and Matterport3D [6] datasets. As shown in Tables 2 and 3, removing the Fusion Controller (w/o FC) leads to a clear performance drop across all metrics. On the Replica [23] dataset, the SPL decreases from 76.7 to 69.3 and SR from 93.1 to 84.9, indicating that the absence of adaptive fusion regulation causes imbalance between modalities and degrades navigation efficiency in structured synthetic scenes. On the Matterport3D [6] dataset, performance also declines, though to a lesser extent, suggesting that in complex real-world environments the model can still partially rely on natural audio-visual

Table 4. Effect of the residual scaling initialization factor β_{init} on navigation performance. Results show that a moderate initialization ($\beta_{\text{init}} = 0.2$) yields the best balance between training stability and accuracy on both Replica and Matterport3D datasets.

β_{init}	Replica			Matterport3D		
	SPL \uparrow	SR \uparrow	SNA \uparrow	SPL \uparrow	SR \uparrow	SNA \uparrow
0.1	76.4	91.6	46.4	53.2	64.4	31.3
0.2	76.7	93.1	47.3	57.3	70.3	33.2
0.3	75.2	92.6	46.1	53.3	65.5	32.0

complementarity. Overall, these results demonstrate that the Fusion Controller plays a crucial role in stabilizing cross-modal interactions, mitigating modality dominance, and enhancing generalization across diverse environments.

In this experiment, we investigate the effect of the initial residual scaling factor β_{init} on the performance of the proposed model. The residual scaling factor controls the initial strength of the cross-modal residual connections in CRFN, which is crucial for balancing the contributions from both audio and visual modalities during training. We test three different values of β_{init} , namely 0.1, 0.2, and 0.3, across two popular datasets, Replica and Matterport3D. The results, summarized in Table 4, show that a moderate value of $\beta_{\text{init}} = 0.2$ leads to the best overall performance, achieving higher SPL, SR, and SNA across both datasets. This suggests that initializing the residual scaling factor with a moderate value helps ensure a stable learning process, avoiding overdominance of any single modality while maintaining effective cross-modal fusion.

5. Conclusion

This paper presents CRFN, a cross-modal residual fusion network designed for audio-visual navigation tasks. CRFN achieves fine-grained alignment and complementary modeling between modalities through a bidirectional residual interaction mechanism, while a fusion controller dynamically adjusts the contribution of each modality, effectively suppressing imbalance and feature degradation. Experimental results demonstrate that CRFN achieves superior navigation performance and stable cross-domain generalization in diverse and complex environments, validating the effectiveness and robustness of the proposed fusion mechanism.

Moreover, by analyzing the dynamic evolution of residual coefficients, we uncover an interesting adaptive modality dependence phenomenon: agents tend to rely more on visual cues in high-fidelity synthetic environments, whereas cross-modal complementarity becomes crucial in complex real-world scenes. This finding provides new insights into how embodied agents coordinate multimodal perception and decision-making.

Acknowledgements

This research was financially supported by the National Natural Science Foundation of China (Grant No. 62463029).

References

- [1] 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 preprint arXiv:1807.06757*, 2018. 6
- [2] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3674–3683, 2018. 3
- [3] S. Bansal, V. Tolani, S. Gupta, J. Malik, and C. Tomlin. Combining optimal control and learning for visual navigation in novel environments. In *Conference on Robot Learning*, pages 420–429. PMLR, 2020. 6
- [4] D. Batra, A. Gokaslan, A. Kembhavi, O. Maksymets, R. Mottaghi, M. Savva, A. Toshev, and E. Wijmans. Objectnav revisited: On evaluation of embodied agents navigating to objects. *arXiv preprint arXiv:2006.13171*, 2020. 3
- [5] J. A. Caley, N. R. Lawrance, and G. A. Hollinger. Deep learning of structured environments for robot search. *Autonomous Robots*, 43(7):1695–1714, 2019. 6
- [6] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niebner, M. Savva, S. Song, A. Zeng, and Y. Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *7th IEEE International Conference on 3D Vision, 3DV 2017*, pages 667–676. Institute of Electrical and Electronics Engineers Inc., 2018. 2, 3, 5, 6, 7, 8
- [7] D. S. Chaplot, R. Salakhutdinov, A. Gupta, and S. Gupta. Neural topological slam for visual navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12875–12884, 2020. 6
- [8] C. Chen, Z. Al-Halah, and K. Grauman. Semantic audio-visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15516–15525, 2021. 3
- [9] C. Chen, U. Jain, C. Schissler, S. V. A. Gari, Z. Al-Halah, V. K. Ithapu, P. Robinson, and K. Grauman. Soundspaces: Audio-visual navigation in 3d environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17–36. IEEE/CVF, 2020. 3, 5, 6, 7
- [10] C. Chen, S. Majumder, Z. Al-Halah, R. Gao, S. K. Ramakrishnan, and K. Grauman. Learning to set waypoints for audio-visual navigation. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7 2021*. OpenReview.net. 3, 5, 6
- [11] J. Chen, B. Lin, X. Liu, L. Ma, X. Liang, and K.-Y. K. Wong. Affordances-oriented planning using foundation models for continuous vision-language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23568–23576, 2025. 1

- [12] J. Chen, W. Wang, S. Liu, H. Li, and Y. Yang. Omnidirectional information gathering for knowledge transfer-based audio-visual navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10993–11003, 2023. [3](#)
- [13] Q. Chen, D. Pitawela, C. Zhao, G. Zhou, H.-T. Chen, and Q. Wu. Webvln: Vision-and-language navigation on websites. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 1165–1173, 2024. [1](#)
- [14] C. Gan, Y. Zhang, J. Wu, B. Gong, and J. B. Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9701–9707. IEEE, 2020. [3](#), [5](#), [6](#)
- [15] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2616–2625, 2017. [1](#)
- [16] J.-H. Kim, S.-W. Lee, D. Kwak, M.-O. Heo, J. Kim, J.-W. Ha, and B.-T. Zhang. Multimodal residual learning for visual qa. *Advances in neural information processing systems*, 29, 2016. [3](#)
- [17] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, M. Deitke, K. Ehsani, D. Gordon, Y. Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017. [6](#)
- [18] X. Liu, S. Paul, M. Chatterjee, and A. Cherian. Caven: An embodied conversational agent for efficient audio-visual navigation in noisy environments. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 3765–3773, 2024. [3](#)
- [19] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9339–9347, 2019. [3](#), [5](#)
- [20] Z. Shi, L. Zhang, L. Li, and Y. Shen. Towards audio-visual navigation in noisy environments: A large-scale benchmark dataset and an architecture considering multiple sound-sources. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 14673–14680, 2025. [4](#)
- [21] N. Srivastava and R. R. Salakhutdinov. Multimodal learning with deep boltzmann machines. *Advances in neural information processing systems*, 25, 2012. [4](#)
- [22] G. J. Stein, C. Bradley, and N. Roy. Learning over subgoals for efficient navigation of structured, unknown environments. In *Conference on robot learning*, pages 213–222. PMLR, 2018. [6](#)
- [23] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. [2](#), [3](#), [5](#), [6](#), [8](#)
- [24] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. S. Chaplot, O. Maksymets, et al. Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in neural information processing systems*, 34:251–266, 2021. [3](#)
- [25] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. *arXiv preprint arXiv:1911.00357*, 2019. [1](#)
- [26] Y. Wu, P. Zhang, M. Gu, J. Zheng, and X. Bai. Embodied navigation with multi-modal information: A survey from tasks to methodology. *Information Fusion*, 112:102532, 2024. [1](#)
- [27] Z. Yang, J. Liu, P. Chen, A. Cherian, T. K. Marks, J. Le Roux, and C. Gan. Rila: Reflective and imaginative language agent for zero-shot semantic audio-visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16251–16261, 2024. [4](#)
- [28] A. Younes, D. Honerkamp, T. Welschehold, and A. Valada. Catch me if you hear me: Audio-visual navigation in complex unmapped environments with moving sounds. *IEEE Robotics and Automation Letters*, 8(2):928–935, 2023. [3](#)
- [29] Y. Yu, L. Cao, F. Sun, X. Liu, and L. Wang. Pay self-attention to audio-visual navigation. In *33rd British Machine Vision Conference 2022, BMVC 2022*, page 46, London, UK, November 2022. BMVA Press. [3](#)
- [30] Y. Yu, L. Cao, F. Sun, C. Yang, H. Lai, and W. Huang. Echo-enhanced embodied visual navigation. *Neural Computation*, 35(5):958–976, 04 2023. [3](#)
- [31] Y. Yu, C. Chen, L. Cao, F. Yang, and F. Sun. Measuring acoustics with collaborative multiple agents. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-23)*, Macao, China, August 2023. [3](#)
- [32] Y. Yu, W. Huang, F. Sun, C. Chen, Y. Wang, and X. Liu. Sound adversarial audio-visual navigation. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. [3](#)
- [33] Y. Yu, Z. Jia, F. Shi, M. Zhu, W. Wang, and X. Li. Weavenet: End-to-end audiovisual sentiment analysis. In *International Conference on Cognitive Systems and Signal Processing*, pages 3–16, 2021. [4](#)
- [34] Y. Yu and D. Yang. Dope: Dual object perception-enhancement network for vision-and-language navigation. In *Proceedings of the 2025 International Conference on Multimedia Retrieval*, pages 1739–1748, 2025. [1](#)
- [35] D. Zheng, S. Huang, L. Zhao, Y. Zhong, and L. Wang. Towards learning a generalist model for embodied navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13624–13634, 2024. [1](#)