Measuring Acoustics with Collaborative Multiple Agents

Yinfeng Yu[1,6], Changan Chen[2], Lele Cao[1,3], Fangkai Yang[4], Fuchun Sun[1,5,*]

[1]Tsinghua University [2]UT Austin [3]EQT [4]Microsoft Research
[5] THU-Bosch JCML Center [6]Xinjiang University

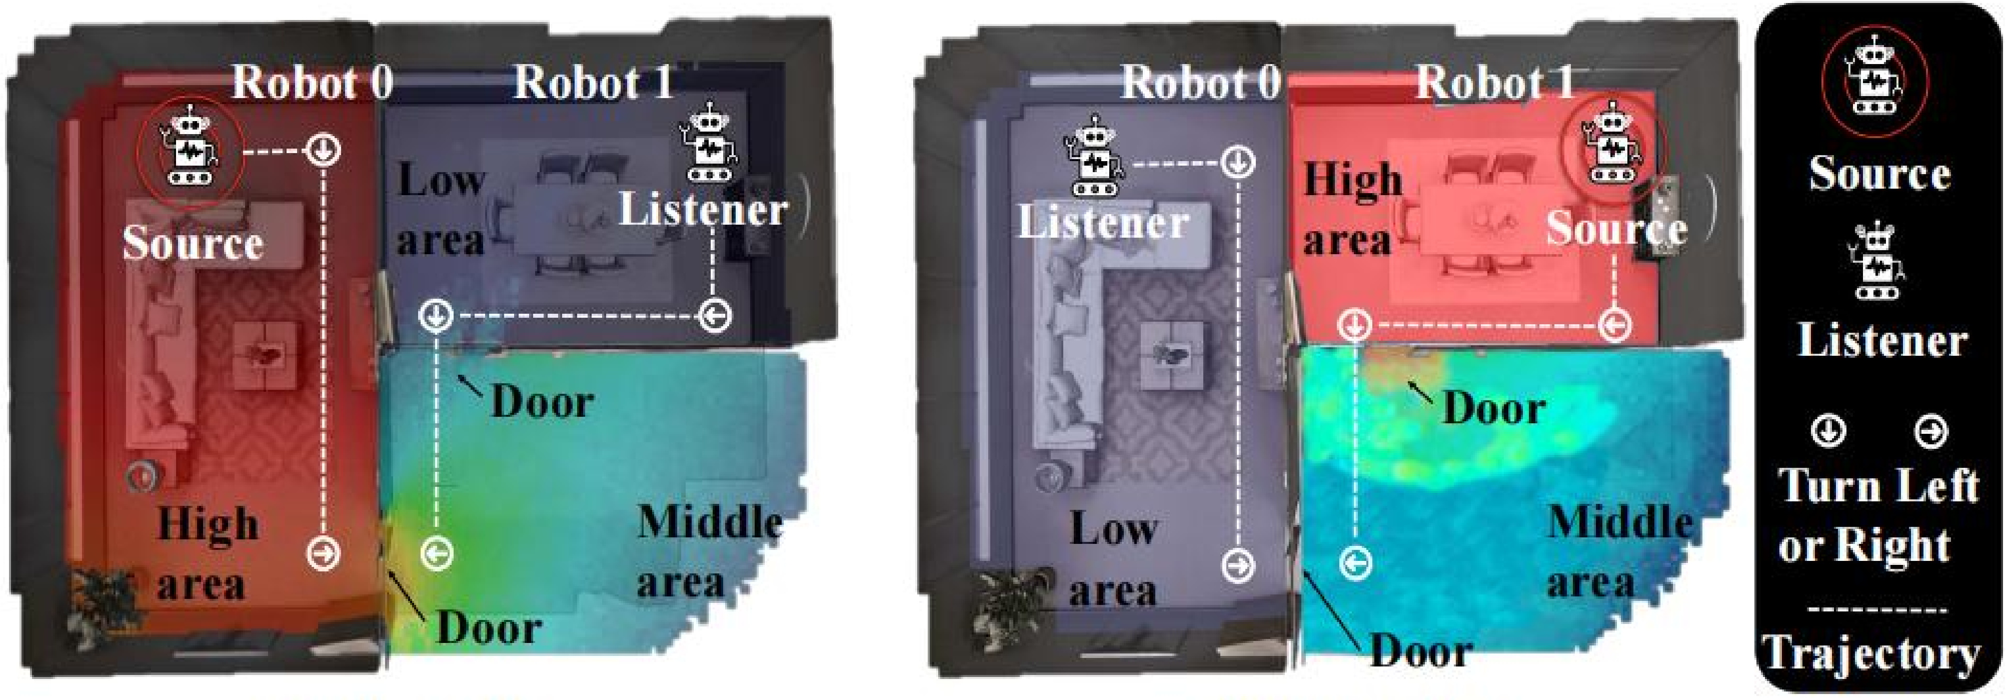

(a) Phase One (b) Phase Two

This work aims to learn to measure environment acoustics with two collaborative robots.

∗ Corresponding author: Fuchun Sun.

Content

- Motivation
- Challenges
- Methodology
- Experiments
- Conclusions

Motivation

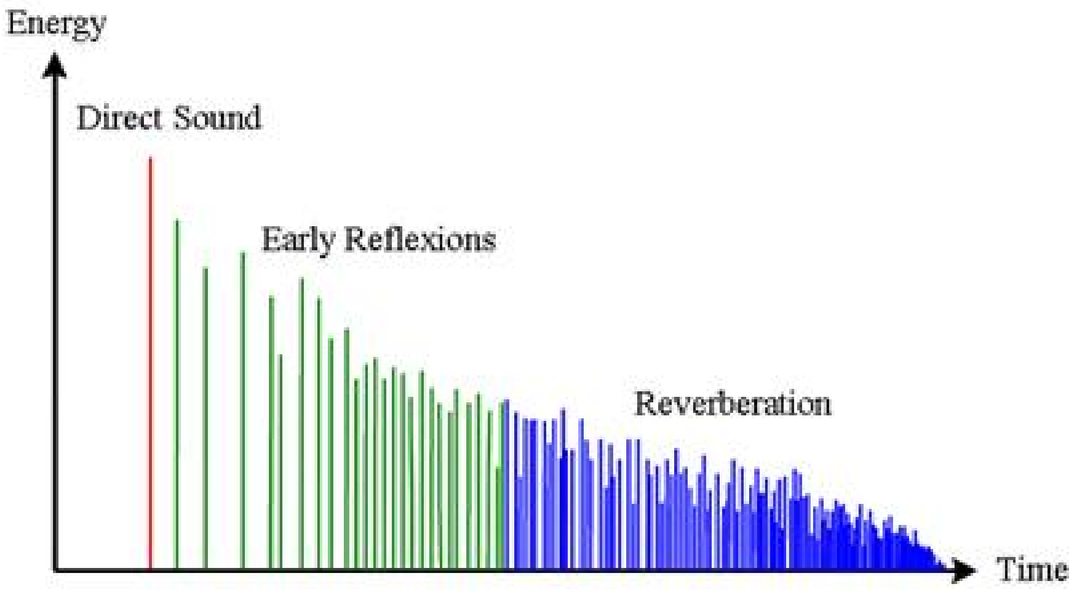

Schematic energetic room impulse response[1]

RIR is short for room impulse response

- Measuring is time-consuming due to the large number of samples to traverse
- FAST-RIR [Ratnarajah et al., 2022b] relies on handcrafted features,
- MESH2IR [Ratnarajah et al., 2022a] uses the scene's mesh and source/listener locations to generate RIR while ignoring the measurement cost
- Neural Acoustic Field [Luo et al., 2022] tries to learn the parameters of the acoustic field, but its training time and model storage cost grow linearly with the number of environments [Majumder et al., 2022].
- Some work [Singh et al., 2021] suggests that storing the original RIR data of the sampled points is optional, and only the acoustic field parameters must be stored.

[1] https://acoustics.org/pressroom/httpdocs/152nd/behler.html

Challenges

Challenge 1:
Given a limited number of action steps, it is challenging to model the acoustic field.

Challenge 2:
What are the evaluation metrics for collaborative RIR measurements? There are no such metrics existed before!

Contributions

- We propose a new setting for planning RIR measuring under finite time steps and a solution to measure the RIR with two-agent cooperation in low resource situations;
- We design a novel reward function for the multi-agent decomposition to encourage coverage of environment acoustics;
- We design evaluation metrics for the collaborative measurement of RIR, and we experimentally verify the effectiveness of our model.

Task

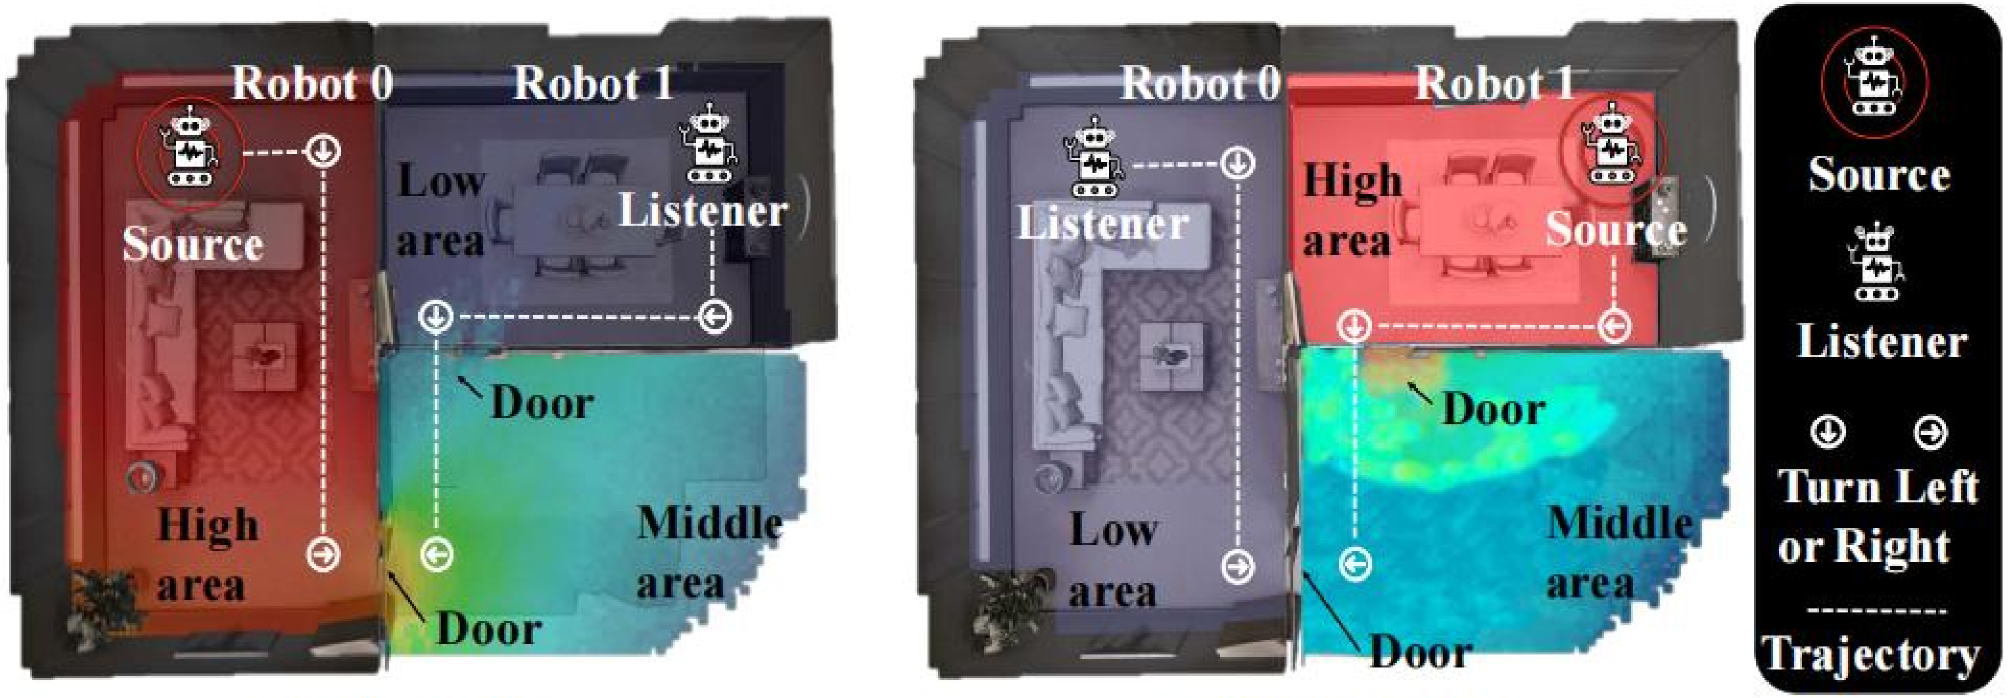

(a) Phase One (b) Phase Two

Learn to measure environment acoustics with two collaborative robots.
The background color indicates sound intensity (``High'', ``Middle'' and ``Low'' areas).
Each step (one step per second) embodies three steps: 1) robot 0 emits a sound, and robot 1 receives the sound; 2) robot 1 emits the sound, and robot 0 receives the sound; 3) two robots make a movement following their learned policies.
This process repeats until reaching the maximum number of time steps.

Methodology

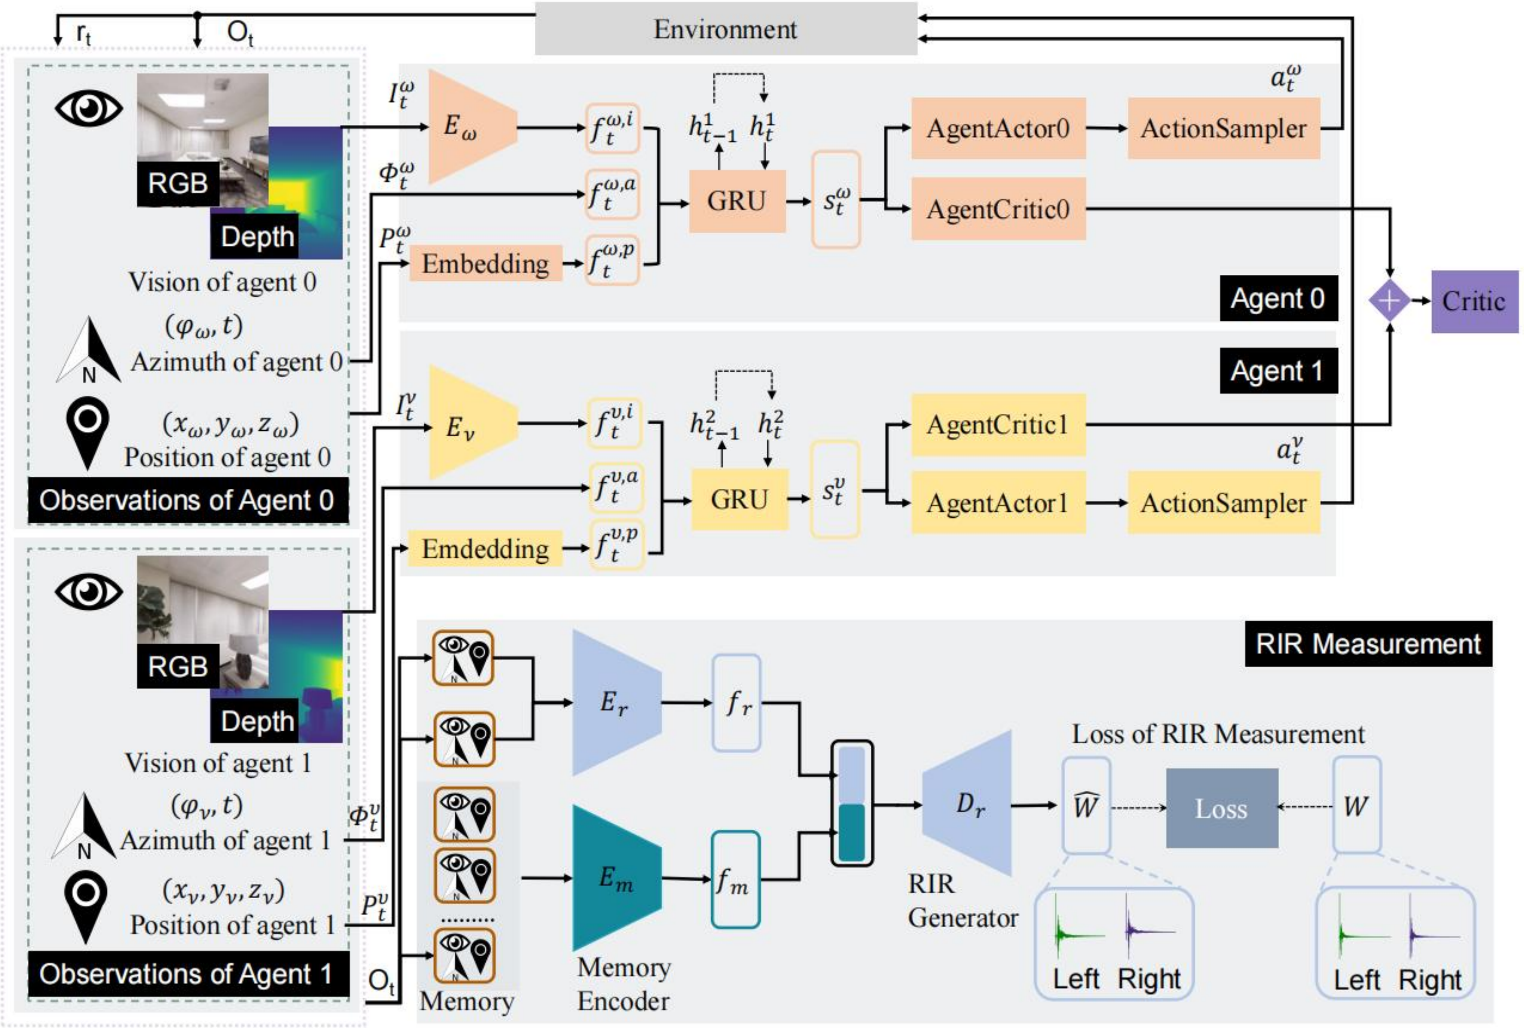

The MACMA architecture: the agent 0 and the agent 1 first learn to encode observations as s_t^ω and s_t^ν respectively using encoder E_ω and E_ν, which are fed to actor-critic networks to predict the next action a_t^ω and a_t^ν.

The RIR measurement learns how to predict RIR $\widehat{W}_t$ guided by ground truth RIR W_t.

IJCAI/2023 MACAO

Experiments (1) : metrics

- **CR**:coverage rate, $CR = \frac{N_v}{N_e}$
- ***PE:*** prediction error, $PE = \Delta(W_t, \widehat{W_t})$
- **WCR**:weighted coverage rate,
$$WCR = (1.0 - \lambda) * CR + \lambda * (1.0 - PES)$$
$$PES = \frac{2.0}{1+e^{-PE}} - 1.0$$
- **RTE**: RT60 Error,
- **SiSDR**: scale-invariant signal-to-distortion ratio,

$$SiSDR = 10log_{10}\frac{||X_T||^2}{||X_E||^2}$$

where

- N_v is the total number of unique nodes that two agents have visited together
- N_e is the total number of all individual nodes in the current episode
- W_t is the ground truth RIR
- $\widehat{W_t}$ is the predicted RIR
- Δ is the function of the STFT distance (Eq. 8)
- PES is Scaled Prediction Error
- $0 \leq \lambda \leq 1.0$ is a hyperparameter
- RTE describes the difference between the ground truth RT60 value and the predicted one
- X_T is the ground truth vector
- X_E is the error vector

Experiments (2) : dataset

We used the SoundSpaces dataset, including:

- Replica
- Matterport3D

Experiments (3) : baselines

We compare our model to the following baselines and existing works:

- **Random**: uniformly samples one of three actions and executes Stop when it reaches maximum steps.
- **Nearest neighbor:** predict from closest experience
- **Occupancy**[1]: orient to occupy more area, making the area of Γ_{ABCD} in Figure 3 larger.
- **Curiosity**[2]: strives to visit nodes that have not been visited already in the current episode.

[1] Ramakrishnan S K, Al-Halah Z, Grauman K. Occupancy anticipation for efficient exploration and navigation, ECCV 2020
[2] Pathak D, Agrawal P, Efros A A, et al. Curiosity-driven exploration by self-supervised prediction, ICML 2017

Experiments (4) : Quantitative comparison of the two datasets.

Model	Replica					Matterport3D				
	WCR (↑)	PE (↓)	CR (↑)	RTE (↓)	SiSDR (↑)	WCR (↑)	PE (↓)	CR (↑)	RTE (↓)	SiSDR (↑)
Random	0.3103	5.4925	0.3439	14.7427	20.3534	0.2036	5.5552	0.2254	23.5281	12.3042
Nearest neighbor	0.3444	5.4533	0.3817	14.0269	22.0135	0.2099	5.3342	0.2321	28.8765	15.2351
Occupancy	0.4464	3.7224	0.4907	12.5532	23.0666	0.2225	4.5327	0.2449	20.3399	18.3848
Curiosity	0.4327	3.4883	0.4742	**10.9565**	**23.8669**	0.2111	4.4255	0.2319	29.5572	20.0031
MACMA (Ours)	**0.6977**	**3.2509**	**0.7669**	13.8896	23.6501	**0.3030**	**4.0113**	**0.3327**	**15.9338**	**21.3187**

Table 1: The results of quantitative comparison between our proposed method (MACMA) and baselines.

As a result, we can conclude that MACMA quantitatively outperforms baselines over both datasets.

Experiments (5) : Qualitative comparison on exploration capability.

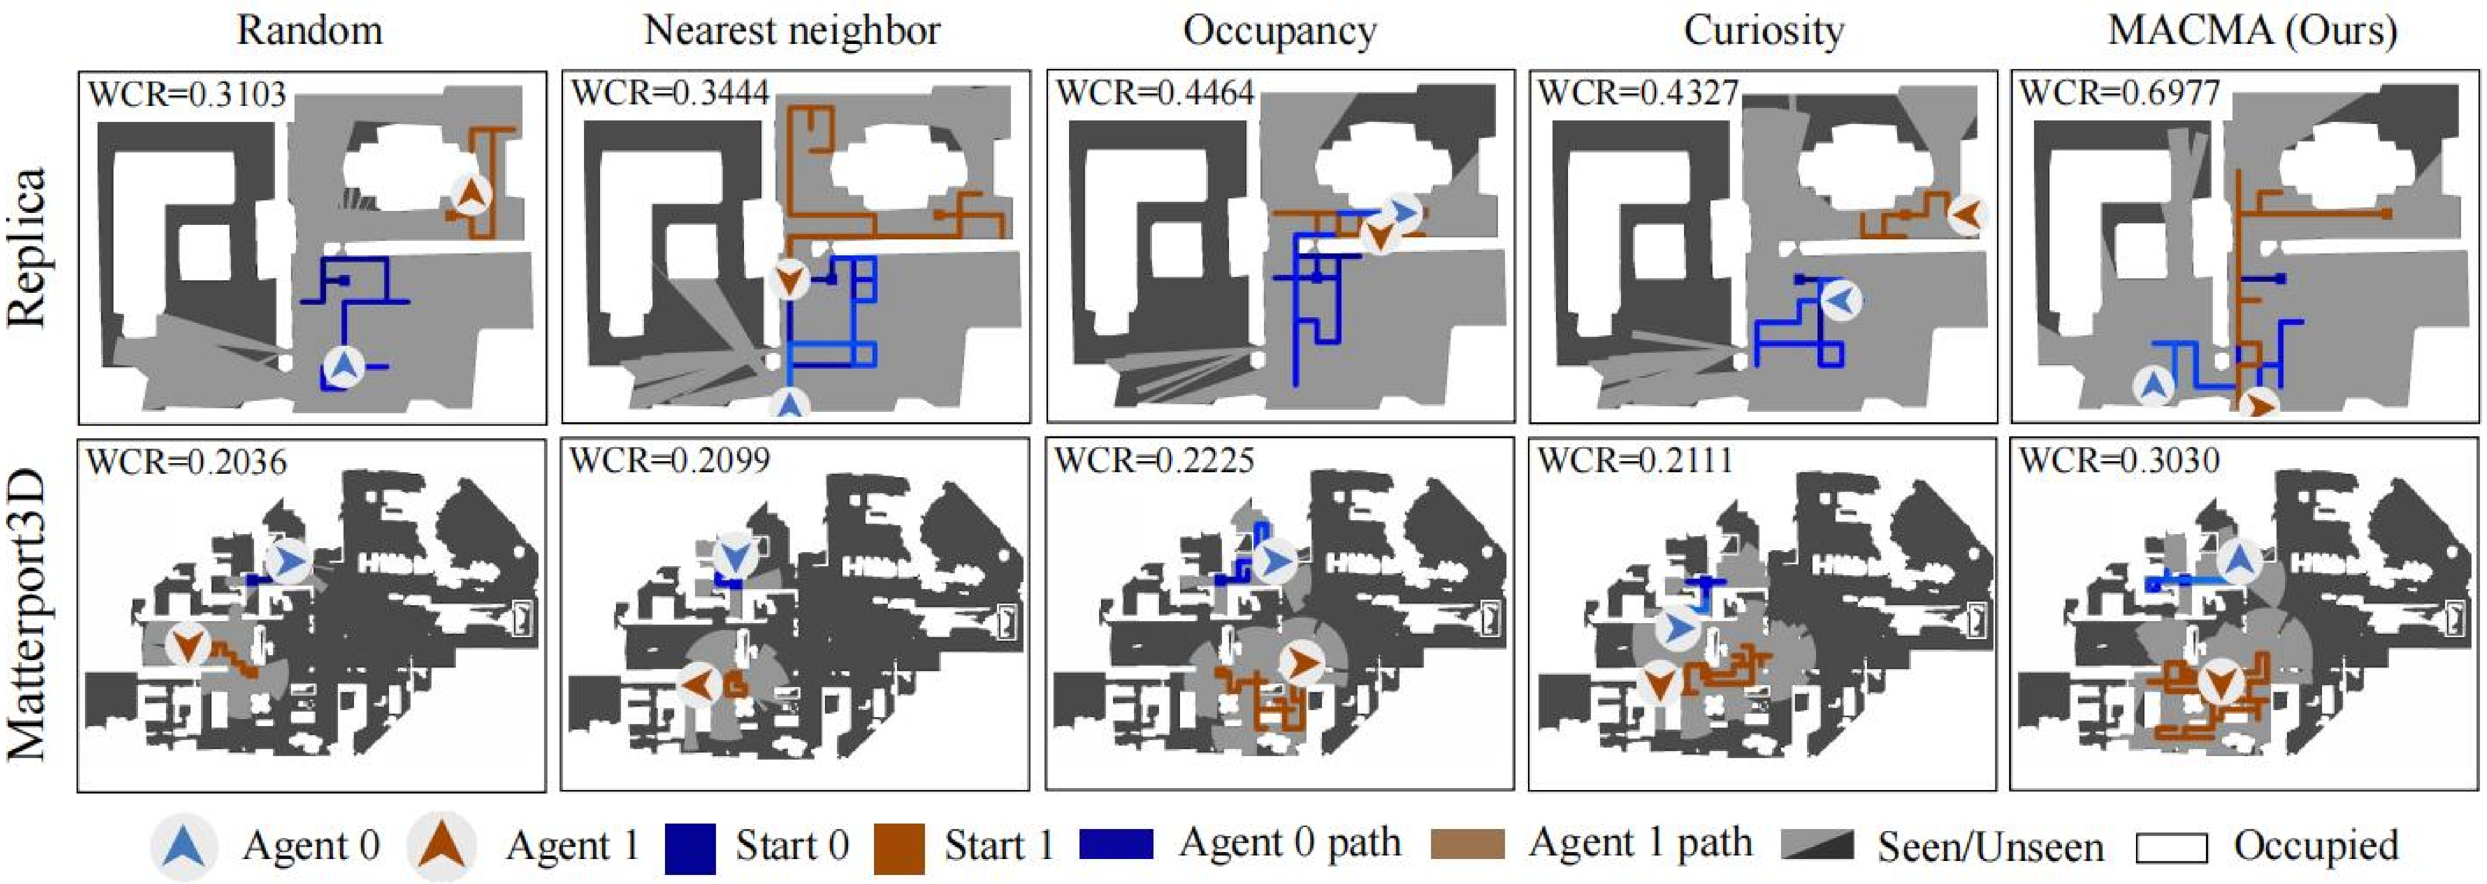

Visualization of the navigation trajectories by the end of a particular episode from Replica (top row) and Matterport3D (bottom row) dataset. Higher WCR values and bigger "seen" areas (colored in light-grey) indicate better performances.

Experiments (6) : Qualitative comparison on RIR prediction.

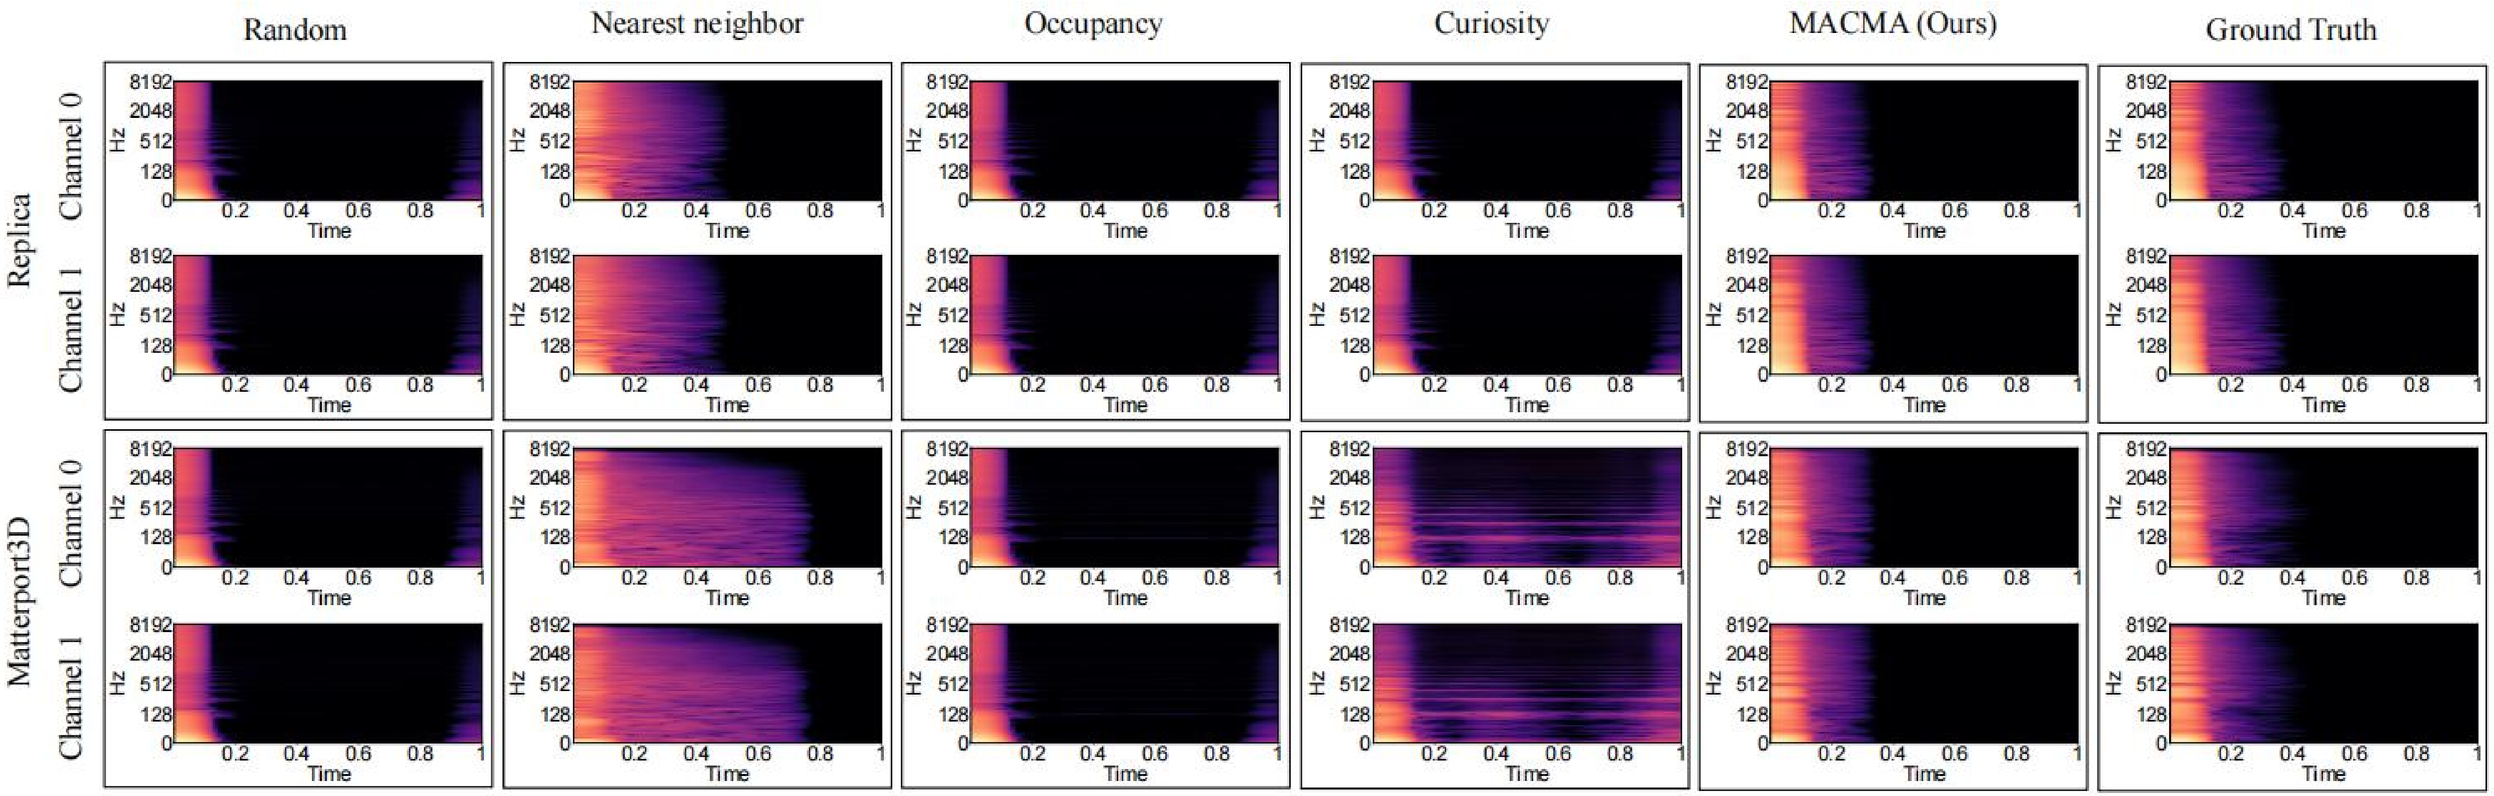

Qualitative comparison of RIR prediction (Binaural RIR with channel 0 and channel 1) in spectrogram from Replica (top row) and Matterport3D (bottom row) dataset. Every column is the result of one model except last one. The last column is the ground truth of RIR.

Experiments (7) : Ablation on modality.

Vision	WCR (↑)	PE (↓)	CR (↑)	RTE (↓)	SiSDR (↑)
Blind	0.5020	3.4966	0.5512	14.2049	23.0903
RGB	0.5930	3.8204	0.6541	15.5897	**23.7713**
Depth	0.5068	3.4927	0.5566	29.6905	23.5089
RGBD	**0.6977**	**3.2509**	**0.7669**	**13.8896**	23.6501

Table 2: Ablation on modality.

As shown in Table 2, RGBD (vision with RGB images and Depth input) seems to be the best choice.

Experiments (8)： MACMA with Reward Assignment.

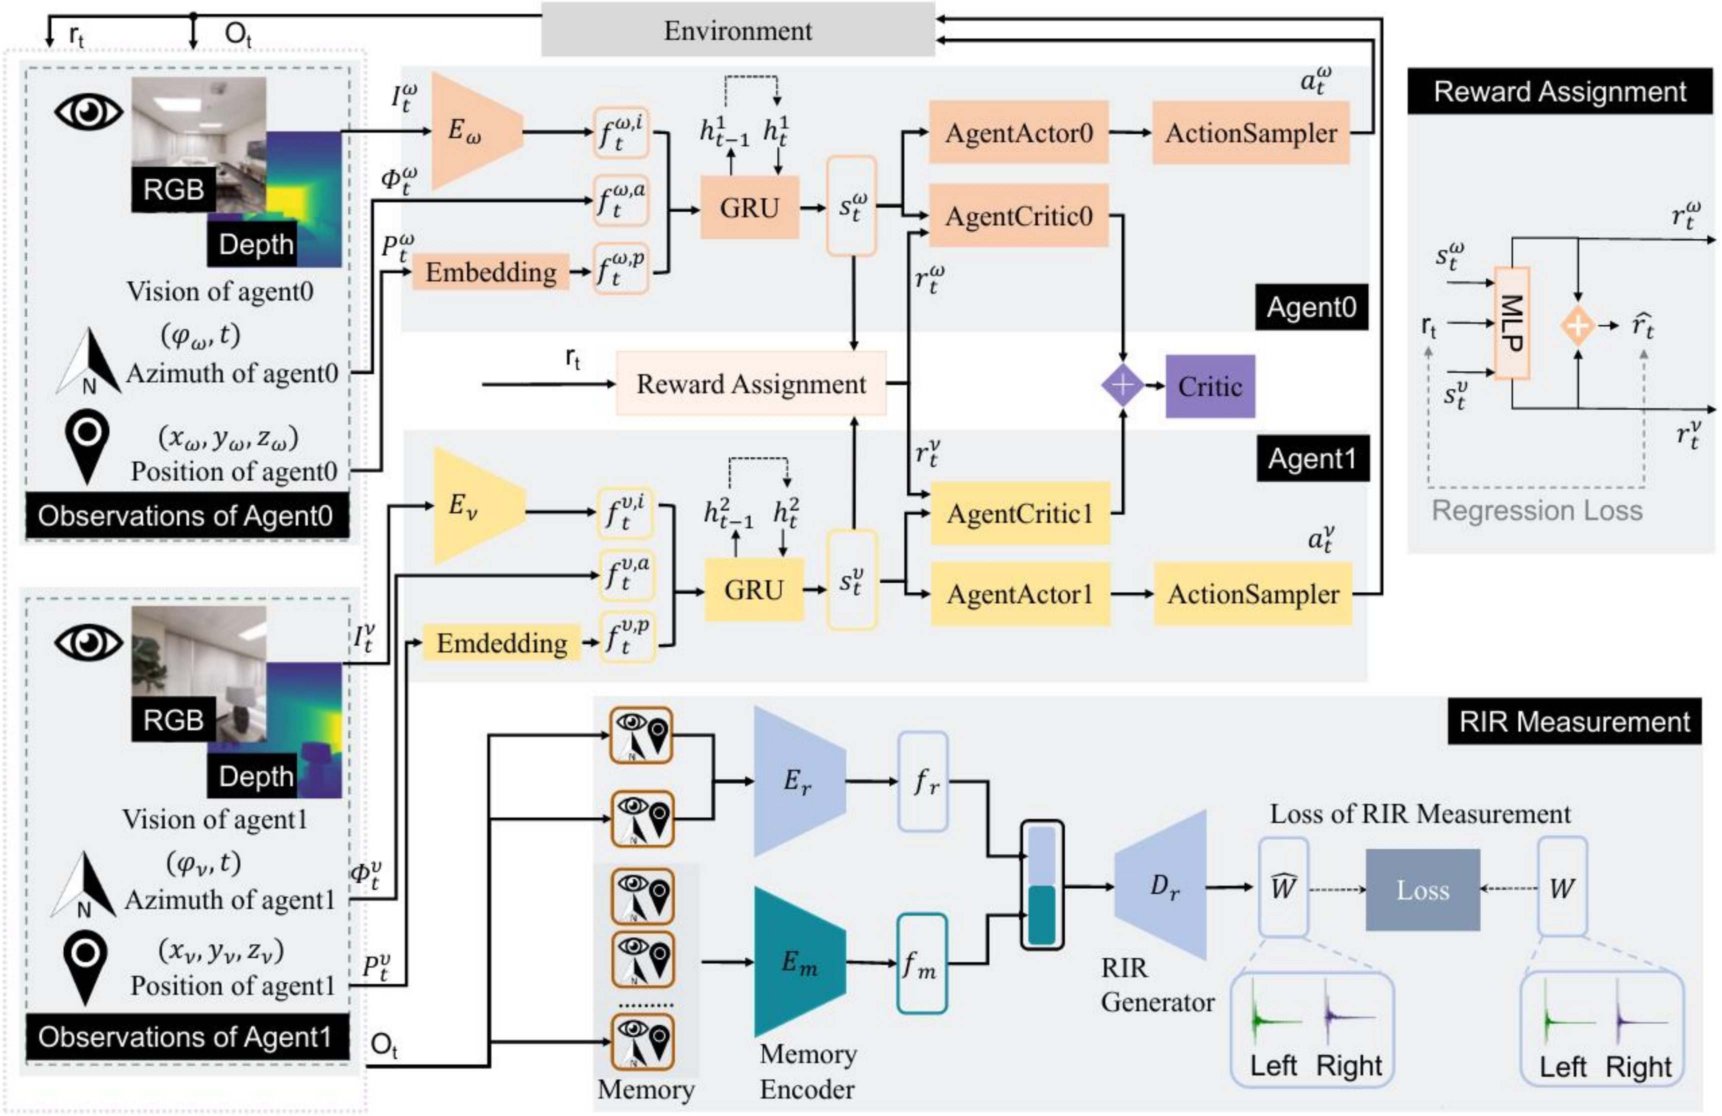

The MACMARA architecture: add **Reward Assignment to MACMA,** the reward assignment learns how to assign reward r_t given by the environment supervised by r_t.

Experiments (9): Experimental results of MACMARA.

ρ	WCR (↑)	PE (↓)	CR (↑)	RTE (↓)	SiSDR (↑)
-1.0♡	**0.6977**	3.2509	**0.7669**	13.8896	23.6501
0.0♣	0.6614	3.5727	0.7288	12.6987	23.1183
0.2	0.5287	3.3398	0.5798	14.0822	23.6715
0.4	0.5040	3.1860	0.5512	15.3278	23.6485
0.6	0.5384	3.3099	0.5904	**12.4184**	23.2581
0.8	0.4510	**2.9755**	0.4903	14.5194	**23.8725**
1.0	0.5471	3.2316	0.5995	12.9246	23.8201

♡: Without reward assignment module, and every agent get all the environment reward ($r_t^\omega = r_t^\nu = r_t$).
♣: With reward assignment module, and every agent get all the environment reward ($r_t^\omega = r_t^\nu = 0.5r_t$).

Experiments (10)： more discussions

- The dynamic relationship between modality importance and action selection
- The dynamic relationship between modal importance and RIR measurement accuracy
- Studies on memory size κ
- Studies on the reward component
- ...

For more details, please refer to the origin paper.

Conclusion

- This paper proposes a novel task where two collaborative agents learn to measure room impulse responses of an environment by moving and emitting/receiving signals in the environment within a given time budget.

- To tackle this task, we design a collaborative navigation and exploration policy.

- Our approach outperforms several other baselines on the environment's coverage and prediction error.

For more details, please refer to the origin paper.

Measuring Acoustics with Collaborative Multiple Agents

Yinfeng Yu[1,6], Changan Chen[2], Lele Cao[1,3], Fangkai Yang[4], Fuchun Sun[1,5,*]

[1]Tsinghua University [2]UT Austin [3]EQT [4]Microsoft Research
[5] THU-Bosch JCML Center [6]Xinjiang University

The project and code can be viewed at the following website:

https://yyf17.github.io/MACMA/index.html

Thanks for your attention!