
Achieving Cross Modal Generalization with Multimodal Unified Representation Appendix

Anonymous Author(s)
Affiliation
Address
email

1 A Details of Downstream Tasks

Table 1: The illustration of all downstream tasks.

Task	Pretrained Modality	Downstream Dataset	Generalization Direction
Cross-Modal	Audio-Visual	AVE	$A \rightarrow V$
Event Classification	Audio-Visual-Text		$V \rightarrow A$
Cross-Modal	Audio-Visual	AVVP	$A \rightarrow V$
Event Localization	Audio-Visual-Text		$V \rightarrow A$
Cross-Modal	Audio-Text	AVS-S4	$A \rightarrow T$
Video Segmentation	Audio-Visual-Text		$T \rightarrow A$
Cross Modal & Dataset	Audio-Visual	AVE & AVVP	$A(AVE) \rightarrow V(AVVP)$
Event Localization	Audio-Visual-Text		$V(AVE) \rightarrow A(AVVP)$

2 B Implementation and Training Details

3 B.1 Pretraining Details

4 Following previous work [1], we extract pool5 feature maps from sampled 16 RGB video frames by
5 VGG-19 network for each 1s visual segment, then we use global average pooling over the 16 frames
6 to generate visual feature maps as $7 \times 7 \times 512$ -D. Also, we use the VGG-like model pre-trained
7 on AudioSet to extract 128-D audio features for each 1s audio segment. For text representations,
8 we first design several prompt templates for each video label of the VGGSound-AVEL [2] dataset,
9 and transform it into a descriptive sentence, see Section. B.3 for more details. Then we use Bert [3]
10 as our text encoder and can get 768-D feature for each word. For all modalities, we extract their
11 semantic features dimension as 256-D. We use three convolution networks to get visual-specific
12 feature as $T \times 2048$ and use an average pooling and Linear layer to compress it to apply mutual
13 information minimization with the visual-semantic feature. Since audio and text do not contain
14 spatial information, we use a Linear layer to get audio-specific and text-specific representations
15 as $T \times 256$, respectively. For Cross-CPC, we set the default prediction steps as 2. For the CLUB
16 applied in each modality, we divide its training process into first forward (mainly used to optimize
17 the approximation network q_θ) and second forward (mainly used for MI minimization), during each
18 epoch, the first forward will be updated 5 times and the second forward will be updated only once.
19 We set the learning rate as 0.0004, the γ in MM-EMA as 0.99, and the batch size as 64. For the
20 inactivated code reset mechanism, we set the N_{re} as 200.

21 For three modalities of unified representation pre-training, most details are the same as previous
22 stage. We need to apply Cross-CPC between every two modalities, i.e., visual-text, visual-audio, and
23 audio-text. And the default prediction step is set as 1. All the above pre-training experiments are
24 conducted on one NVIDIA A100 GPU.

25 B.2 Downstream Tasks

26 **Cross-Modal Event Classification:** The AVE [1] dataset contains a total of 28 different event types,
27 and the length of each audio and video is 10 seconds. Taking video-to-audio generalization (V2A)
28 as an example: during downstream training, we use the Encoder obtained in the abovementioned
29 pre-training phase to encode video input into the unified discrete space, the 10s second video can
30 get discrete vectors of length 10. Then we use two Linear layers as the Decoder to map the discrete
31 vectors into a 28-D feature space, and use the softmax function to find the event with the highest
32 prediction probability and calculate the loss with Ground Truth. During the whole process, the
33 parameters of the Encoder are frozen, only the Decoder will be updated. We set the learning rate
34 as $2.5 \times e^{-4}$, and the batch size as 256. After training, we directly replace the video with audio
35 as input, and use the obtained Encoder and Decoder to predict events in the audio. The process of
36 audio-to-video generalization (A2V) is similar.

37 **Cross-Modal Event Localization:** The AVVP [4] dataset contains a total of 25 different event
38 types, and the length of each audio and video is 10 seconds. Different from each video in the AVE
39 dataset that only contains one event type, the audio and video in the AVVP dataset may contain
40 multiple different event types. For example, in the same video, the audio information contains events
41 A, B, while visual information contains events A, C, D. As with the previous task, we also use the
42 pre-trained Encoder to map the video input to a unified discrete space, and then we use two layers of
43 Linear as the Decoder to map the discrete vectors to a 25-dimensional space, activate with Sigmoid
44 function and combine with Ground Truth to calculate the loss. Other settings are the same as the
45 cross-modal event classification.

46 **Cross-Modal Video Segmentation:** The AVS-S4 [5] dataset contains 4932 five-seconds videos
47 over 23 categories, including humans, animals, vehicles, and musical instruments. As the same with
48 VGGSound-AVEL dataset, we transform the text label corresponding to each video into a descriptive
49 sentence. Taking audio-to-text generalization (A2T) as an example: during downstream training, we
50 use the Encoder obtained in Audio-Text or Audio-Visual-Text pre-training phase to encode audio
51 input into the unified discrete space. Then we use the architecture proposed in AVS [5] as our visual
52 encoder, audio-video interaction module, and video decoder. During training, the audio encoder
53 is frozen while others are trainable. We use the Adam optimizer with a learning rate of $1e^{-4}$ and
54 the batch size is set to 4. After training, we directly replace the audio with text as input to test the
55 performance of text-based video segmentation. The process of text-to-audio generalization (T2A) is
56 similar.

57 **Cross Modal&Dataset Event Localization:** To further prove that our model is widely applicable
58 in various downstream tasks, we also test the cross-modal generalization ability in cross-dataset
59 scenarios. There are 12 common event categories shared in AVE and AVVP datasets: **dog, car,**
60 **helicopter, violin fiddle, frying food, motorcycle, acoustic guitar, banjo, baby cry, chainsaw, cat,**
61 **accordion.** During the downstream task, we use the video or audio modality in the AVE training
62 set to train the localization model in a weakly-supervised manner, and then directly test the event
63 localization performance of the opposed modality in the AVVP val set. We choose the F1 score as
64 the evaluation metric. The other training details are the same as the Cross-Modal Event Localization
65 task. All the above downstream tasks are training and evaluating on one NVIDIA A100 GPU.

66 B.3 Text Prompts

67 We design different prompt templates for different event categories in VGGSound-AVEL and AVS-S4
68 datasets. For example, for event **race car**, we design the following prompts: *The roar of a high-speed*
69 *race car engine.;* *The race car is running on the road and making a loud engine sound.;* *There is*
70 *a high-speed race car running on the road.* For event **playing violin**, we can design the following
71 prompts: *The sweet and melodious sound of a violin being played.;* *The player is playing the violin.;*
72 *Someone is playing beautiful music on the violin.* For different kinds of events, we will design
73 several unique templates for each event type according to their characteristics. Each piece of data
74 will randomly choose a prompt from these templates as the audio-visual description.

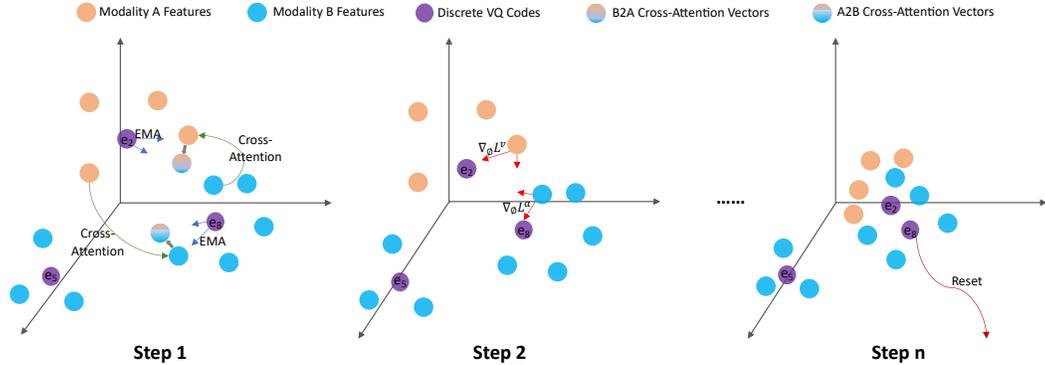


Figure 1: The overview of our proposed MM-EMA. Step 1, we use the cross-attention to obtain composite features and use EMA to update the latent VQ codes. Step 2, we use the modified commitment loss to update the representation of modality A and B. After several steps, these different modalities will be aggregated together, resulting in some redundant discrete code which will be reset.

Table 2: Ablation studies of audio-visual-text pre-training on three downstream tasks.

Method	VGGsounds-AVEL 40K							
	AVE		AVVP		AVE→AVVP		AVS-S4 (mIoU)	
	V→A	A→V	V→A	A→V	V→A	A→V	A→T	T→A
Encoder Frozen	54.1	55.0	63.4	71.0	53.0	52.4	78.0	77.7
Encoder Fine-tuned	44.5	54.9	54.0	50.8	44.5	50.0	77.8	77.1

75 C Details of Multi-Modal EMA

76 During MM-EMA, we use the cross-attention vectors \mathbf{r}_i^b and \mathbf{r}_i^a as the bridge to help the two modalities
77 move closer to each other. At the initial training, the codes mapped from different modalities are not
78 yet overlapped, thus the MM-EMA algorithm is similar to the original EMA equation:

$$N_i^{(t)} = \gamma N_i^{(t-1)} + (1-\gamma)n_i^{(t)} \quad \mathbf{o}_i^{(t)} = \gamma \mathbf{o}_i^{(t-1)} + (1-\gamma) \sum_{j=1}^{n_i^{(t)}} \left[\frac{\mathbf{z}_{i,j}^{a(t)} + \mathbf{r}_{i,j}^{b(t)}}{2} \right] \quad \mathbf{e}_i^{(t)} = \mathbf{o}_i^{(t)} / N_i^{(t)}, \quad (1)$$

79 with the help of the \mathbf{r}_i^b , \mathbf{r}_i^a and the modified commitment loss, the distance between semantic features
80 \mathbf{z}_i^a and \mathbf{z}_i^b is gradually decreasing. With the training progress, more and more latent discrete codes
81 can be mapped from two modalities, thus the MM-EMA algorithm can be written as:

$$N_i^{(t)} = \gamma N_i^{(t-1)} + (1-\gamma)[n_i^{a(t)} + n_i^{b(t)}] \quad \mathbf{e}_i^{(t)} = \mathbf{o}_i^{(t)} / N_i^{(t)} \quad (2)$$

$$\mathbf{o}_i^{(t)} = \gamma \mathbf{o}_i^{(t-1)} + (1-\gamma) \left[\sum_{j=1}^{n_i^{a(t)}} \frac{\mathbf{z}_{i,j}^{a(t)} + \mathbf{r}_{i,j}^{b(t)}}{2} + \sum_{j=1}^{n_i^{b(t)}} \frac{\mathbf{z}_{i,j}^{b(t)} + \mathbf{r}_{i,j}^{a(t)}}{2} \right]$$

82 Finally, the information with the same semantics is converged into the same latent discrete code. This
83 process will generate many redundant codes that will never be utilized, as shown in Fig. 1, thus, it is
84 obliged to use the code reset mechanism to reset these quantized vectors.

85 D More Ablation Studies

86 D.1 The Effect of Whether Froze Encoder in Downstream Tasks

87 To evaluate whether the parameters of the pre-trained encoder should be frozen in downstream tasks,
88 we conduct a series of experiments, as shown in Table. 2. The results demonstrate a significant
89 decline in performance when the encoder is fine-tuned during these tasks. We argue the reason is
90 that when the encoder parameters are not frozen, gradient updates from the downstream tasks may
91 lead to updates in the encoder corresponding to the modality during training, thereby disrupting the
92 alignment relationship between different modalities in the latent space.

Table 3: Ablation studies of the number of Cross-CPC prediction steps on three downstream tasks.

Prediction Steps	VGGsounds-AVEL 40K					
	AVE		AVVP		AVE→AVVP	
	V→A	A→V	V→A	A→V	V→A	A→V
step = 1	54.1	55.0	63.4	71.0	53.0	52.4
step = 2	51.8	54.4	58.2	72.2	50.1	48.2
step = 3	51.4	53.5	59.7	70.3	51.8	49.2
step = 4	52.7	51.8	50.0	62.9	47.4	40.1

Table 4: Ablation studies of Cross-CPC and InfoNCE on two downstream tasks.

Method	VGGsounds-AVEL 40K			
	AVE		AVVP	
	V→A	A→V	V→A	A→V
Cross-CPC	54.1	55.0	63.4	71.0
InfoNCE	54.2	55.9	46.5	39.7

93 D.2 The Effect of the Prediction Steps in Cross-CPC

94 In this paper, we set the default prediction step in the audio-visual pre-training experiments as 2, while
 95 setting it in the audio-visual-text and audio-text pre-training experiments as 1. To further illustrate the
 96 impact of the number of prediction steps on the effectiveness of unified representation, we conduct
 97 a series of experiments on audio-visual-text pre-training, as shown in Table. 3. Our findings reveal
 98 that when the prediction steps are less than 4, the model’s performance in all downstream tasks
 99 exhibits only minor fluctuations. However, a significant decline in performance is observed when the
 100 prediction step reaches 4. This indicates that an excessive number of prediction steps can hinder the
 101 model’s ability to learn fine-grained multimodal alignment relationships, as predicting the distant
 102 future often proves to be more challenging.

103 D.3 Replace Cross-CPC with InfoNCE

104 To further evaluate the effectiveness of our proposed Cross-CPC, we replace it with a traditional
 105 InfoNCE loss. We first compress the sequence audio, visual, and text information into single vectors
 106 and then directly apply contrastive learning among these three modalities. As we can see in Table. 4,
 107 when the Cross-CPC is replaced with InfoNCE, the model achieves comparable performance in
 108 cross-modal event classification tasks, but exhibits inferior performance in event localization tasks.
 109 The results demonstrate that in the case of coarse-grained alignment, our proposed Cross-CPC has
 110 similar effects to InfoNCE. However, in comparison to InfoNCE, our method attains fine-grained
 111 alignment of distinct modality information, which further demonstrates the effectiveness of our
 112 approach in multi-modal sequence unified representation.

113 E More Qualitative Analyses

114 **Visualization of Latent Discrete Codes:** In order to illustrate the process of our method aligning
 115 different modalities in the latent space more intuitively, we visualize the discrete latent codes obtained
 116 from different training epochs and compare our method with the baseline model. As shown in Fig. 2,
 117 we can see that there exist many inactivated codes in the baseline model in the early stage, with
 118 the audio and visual modalities converging separately rather than aligning. In contrast, our method
 119 effectively aligns these distinct modalities. The figure also showcases the efficacy of our inactivated
 120 code reset mechanism. In epoch 1, our approach leaves only a few codes unactivated, which are
 121 promptly reactivated in subsequent epochs. Meanwhile, the baseline model continues to exhibit a
 122 large number of inactivated discrete codes.

123 **Visualization of Segmentation Results:** Also we present more visualization results of A2T and
 124 T2A generalization of our model on AVS-S4 dataset. As we can see in Fig. 3 and 4, our model can
 125 accurately localize the area where the sound is produced even though the test modality has never
 126 been seen before.

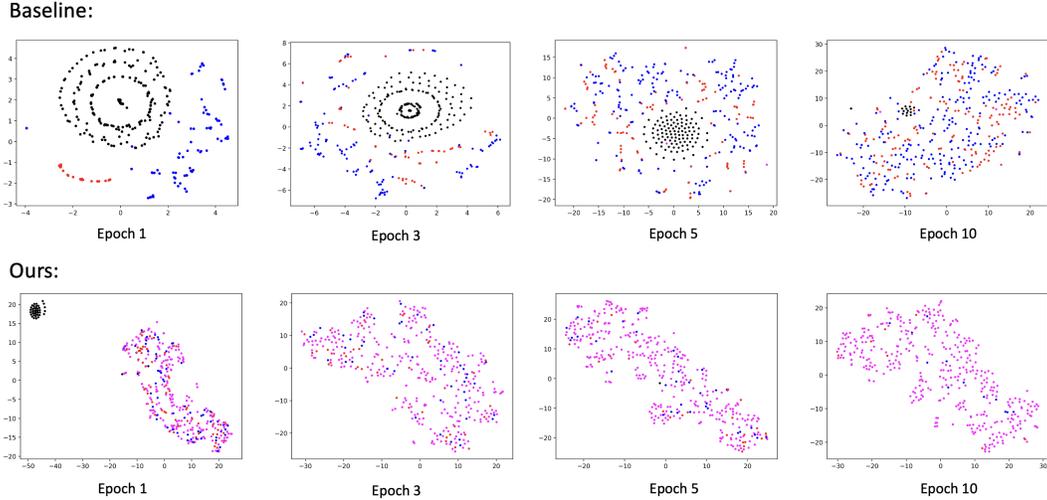


Figure 2: The visualization of latent codes of our model and the baseline model after training different epochs. Red represents audio-only code mapping, blue represents visual-only code mapping, purple represents audio-visual co-mapping, and black means that the code has not been activated.

127 **Visualization of Fine-grained Prediction Score:** To further illustrate that our model has fine-grained
 128 cross-modal generalization capabilities, we visualize the prediction results of the model on the AVVP
 129 dataset, as shown in Fig. 5. For different temporal segments, our model predict different probabilities
 130 for the occurrence of different events, and can accurately predict the types of events that appear in
 131 unknown modalities and the moments when these events occur.

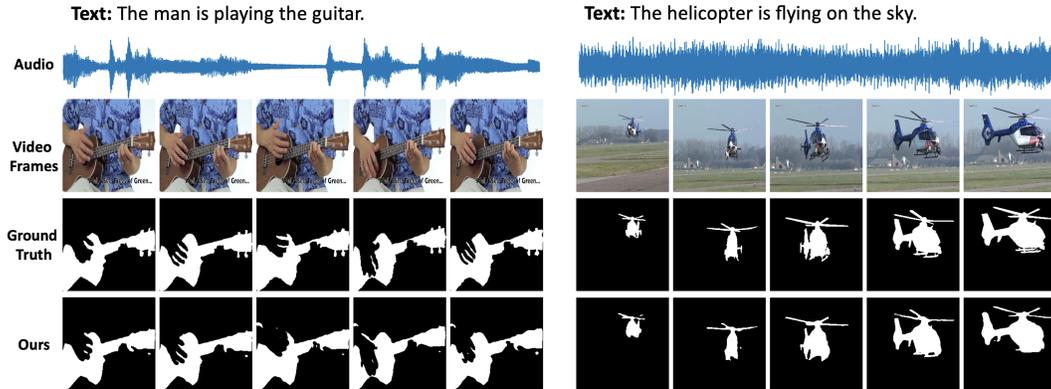


Figure 3: More visualization results of audio-to-text generalization on AVS-S4 dataset.

132 F Discussion

133 During the pre-training phase, our goal is to map information with the same semantics from different
 134 modalities together. To achieve this, we first utilize DCID to extract semantic information from
 135 various modalities, avoiding the introduction of modal-specific information that may affect alignment
 136 results. We then employ Cross-CPC for contrastive learning prediction, aligning different modalities
 137 at a fine-grained level. Subsequently, MM-EMA can assist in mapping these segments with identical
 138 information to the same discrete code. Once the training is complete, we can obtain the multi-modal
 139 unified representation. Then in the downstream training, we can assign additional knowledge for the
 140 known modality, such as event label, and relation with visual objects. Therefore, when the model
 141 is dealing with unknown modalities, by mapping it into a unified space, the corresponding labels
 142 learned from the known modality can be obtained, thus achieving cross-modal generalization and
 143 knowledge transferring.

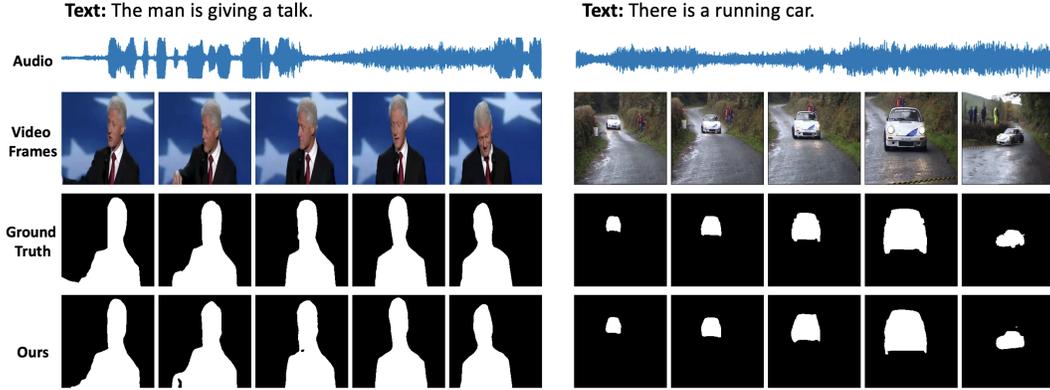
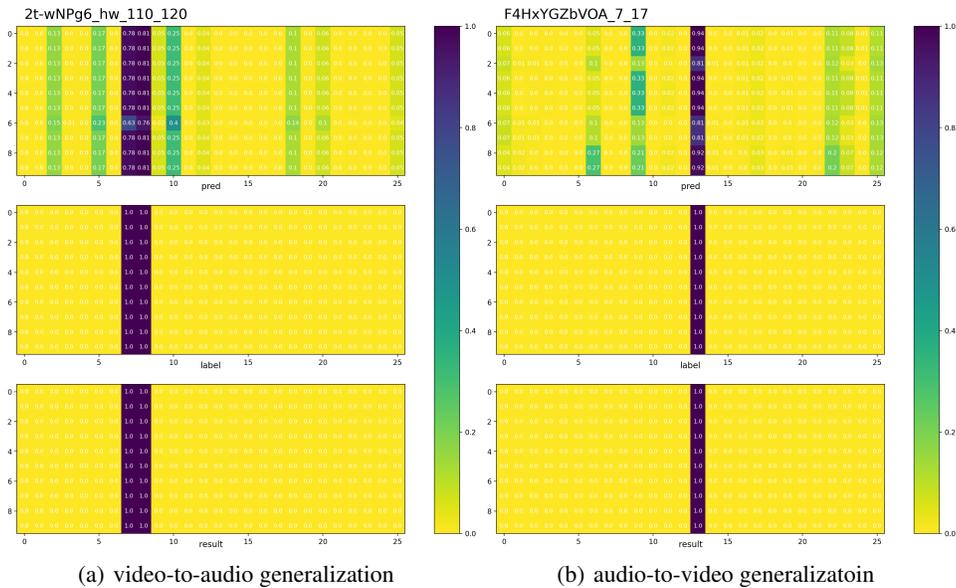


Figure 4: More visualization results of text-to-audio generalization on AVS-S4 dataset.



(a) video-to-audio generalization

(b) audio-to-video generalization

Figure 5: Visualization of the prediction scores on cross-modal event localization tasks. In each subfigure, the vertical axis represents temporal dimension, and the horizontal axis represents event types. The top layer is the possible score of each event at each segment predicted by the model. The middle layer is the ground truth label. The bottom layer represents the results of multiplying the prediction scores and ground truth.

144 References

- 145 [1] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event
 146 localization in unconstrained videos. In *Proceedings of the European Conference on Computer
 147 Vision (ECCV)*, pages 247–263, 2018.
- 148 [2] Jinxing Zhou, Dan Guo, and Meng Wang. Contrastive positive sample propagation along the
 149 audio-visual event line. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- 150 [3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
 151 bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- 152 [4] Yapeng Tian, Dingzeyu Li, and Chenliang Xu. Unified multisensory perception: Weakly-
 153 supervised audio-visual video parsing. In *Computer Vision—ECCV 2020: 16th European Con-
 154 ference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 436–454. Springer,
 155 2020.

- 156 [5] Jinxing Zhou, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo,
157 Lingpeng Kong, Meng Wang, and Yiran Zhong. Audio–visual segmentation. In *Computer Vision–*
158 *ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings,*
159 *Part XXXVII*, pages 386–403. Springer, 2022.