

KDC-MAE: Knowledge Distilled Contrastive Mask Auto-Encoder

Maheswar Bora, Saurabh Atreya, Aritra Mukherjee, Abhijit Das
Machine Intelligence Group, Department of CS&IS,
Birla Institute of Technology and Science, Pilani – Hyderabad Campus
Secunderabad, Telangana 500078

abhijit.das@hyderabad.bits-pilani.ac.in

Abstract

In this work, we attempted to extend the thought and showcase a way forward for the Self-supervised Learning (SSL) learning paradigm by combining contrastive learning, self-distillation (knowledge distillation) and masked data modelling, the three major SSL frameworks, to learn a joint and coordinated representation. The proposed technique of SSL learns by the collaborative power of different learning objectives of SSL. Hence to jointly learn the different SSL objectives we proposed a new SSL architecture KDC-MAE, a complementary masking strategy to learn the modular correspondence, and a weighted way to combine them coordinately. Experimental results conclude that the contrastive masking correspondence along with the KD learning objective has lent a hand to performing better learning for multiple modalities over multiple tasks.

1. Introduction

A natural tendency of human beings is to learn about the environment by themselves or be self-taught". Since our childhood, we start learning via supervised training. Slowly we start analyzing and learning things by ourselves or start getting "self-taught". There are several aspects by which we get self-taught and a combination of them builds our life.

Even the learning representation community adopted the self-taught [40] as SSL. Learning without training typically involves recalling the data, grouping similar things and unrelated separately for learning the abstract knowledge from data. Common SSL learning objectives are *pre-text task* [30], *contrastive learning* [23], and *teacher-student methods* [3, 44, 54]. For pretext learning on images one popular task was colorization [22, 28, 29] and jigsaw puzzle solving [5, 25, 36, 37, 41, 43, 57]. In the past, the jigsaw puzzle rearrangement has been used both on spatial and temporal domains for learning better video representations [51]. Various works [45, 48, 50, 55] used video reconstruction as learning objectives by disentangling motion

and content, using scene dynamics and transformers. Super-resolution is gaining popularity as a learning objective for scale-independent generalized learning [58].

Contrastive learning is a widely used learning objective. A comprehensive survey [12] shows that approaches can be classified into three categories. First, Binary classification of "good pair" and "bad pair" using Binary Cross Entropy (BCE) loss [26, 60]. Second, SSL can be done with discriminators, which is the fuzzy version of the binary classification [46, 52]. Third, the most popular approach in SSL is to use Noise-Contrastive Estimation (NCE) loss [18] to pretrain the model by pulling similar data pairs closer and pushing dissimilar data pairs further. Though originally devised for Natural language Processing (NLP) tasks [34], the concept is now widely used in video [19, 27, 33, 39, 61], images [7, 10, 35] and audio [32].

In recent years since Masked Autoencoders (MAE) was proposed for vision tasks [20], the domain has been heavily explored by researchers resulting in a huge amount of research outcomes in recent years. Due to huge interest in this area, comprehensive surveys [62, 63, 66] were made in recent times. Subsequently, video-MAE [13, 47, 56], Audio-visual MAE [15] and CAV-MAE [16] are proposed on this direction of research. In recent time-motion-guided masking [21, 59] and dual masking [53] have been explored in this regard. It is observed that the selection of a masking policy has a profound impact on the learnability of models. Knowledge distillation also plays an important role in self-distillation between instances with different masking policies [9, 64]. It is also observed that knowledge distillation can provide a significant boost to SSL [6, 17, 38, 65].

The aforementioned learning objectives of SSL are well explored individually and found to perform nonuniformly for different scenarios when they are employed individually. It is expected that they will work better if they learn jointly to find mutual correspondence. A recent work [16] attempted to unify the concepts of masked autoencoder for the reconstruction of video frames and contrastive learning between video and audio tokens to learn better multimodal

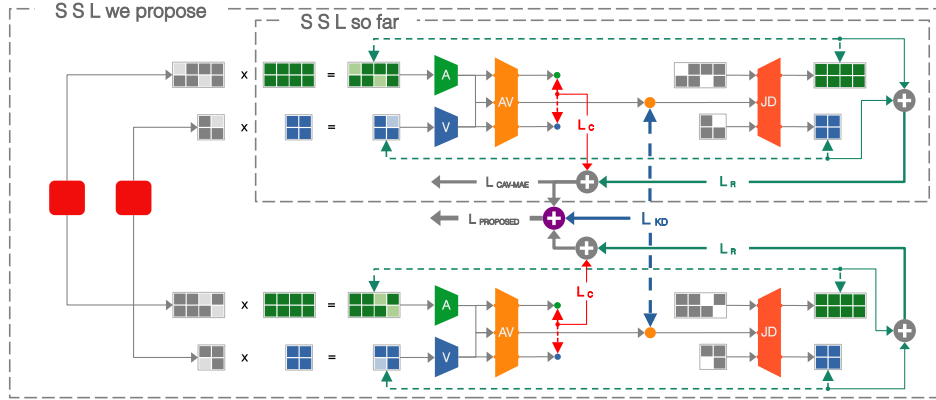


Figure 1. Proposed improvement on existing SSL by complementary mask and self-distillation. The model uses shared weights with two masked versions of the same audio-video pair passed through the encoder, generating separate joint latent embeddings. KL divergence loss aligns these embeddings, followed by a joint decoder that splits audio and video, with contrastive loss applied to the latent embeddings.

representations.

The reconstruction task of CAV-MAE forces its representation to encode the input information in the fusion and the constructive task helps to find the explicit audio-visual correspondence objective but lacks in finding the modular correspondence. This motivates us to propose Knowledge Distilled Contrastive Mask Auto-Encoder (KDC-MAE), an optimized manner to use modular knowledge distillation along with masked data modelling, and inter-modal contrastive learning to learn joint multimodal correspondence of audio and video for unified learning representation. We proposed to use dual complementary masks as input to dual head mask-autoencoder in a weight-sharing fashion to add modular correspondence. Further, to foster the modular correspondence, Kullback–Leibler divergence (KL) of the representation of the dual mask of each modality is computed in the encoding space and propagated as a loss (See Fig. 1).

From our experiments, we conclude that these learning objectives can boost learning in an optimal combination. We demonstrate our findings through various ablation studies at the pretraining level. To support our claim we utilized the pre-trained models in further several downstream tasks over different modalities and their combinations, on which we obtained better results. In summary, our contributions are:

- To propose that three paradigms of SSL namely contrastive learning, masked data modelling, and knowledge distillation objectives are not complementary for unified joint modelling.
- Introduction of complementary masking strategy for finding the modular correspondence of the modalities.
- Introduction of static regularizers as weights for equating the various losses of the different tasks of SSL

for joint and collaborative learning to generalise the contextual information extraction capability of the encoder.

2. Proposed Methodology

In our proposed work we focused on optimization of the ensemble of the SSL, few questions arose for the training approach: (a) Investigating whether reconstruction from masked input, contrastive learning and knowledge distillation is complementary; (b) If they are not complementary, in what way can they be best combined? (b) How can we incorporate mutual learning as a form of self distillation [65], with different unmasked tokens as input to weight shared instances of the encoder, to train it better? (d) How can we make the SSL learning task tougher by experimenting with more masking strategies in order to make the encoder learn better?

2.1. Preliminaries

We follow a similar technique for pre-processing and tokenization as in MAE [11] which was further adopted by CAV-MAE [16]. Hence, we take on audio-visual data by considering a frame of the 10sec video clip as the visual input and the Mel Spectrogram of the entire audio as an image, as the audio input. The encoder architecture for both modalities is the same and the decoder is common. The heart of the encoder is based on transformer layers [49]. Both the audio encoder(AE) and video encoder(VE) are composed of 11 transformer layers and the joint audio-video encoder(AVE) contains 1 layer. The total 12 layer encoder is based on the ViT base model [11]. Thus its input is also 16×16 patches, termed as unmasked tokens. The video input frames are scaled-center-cropped 224×224 thus resulting in 196 patches. The audio spec-

trogram is of size 128×1024 (features \times time) and thus it results in $512 \times 16 \times 16$ patches. The masking ratio was kept at 75% and thus 49 video patches and 128 audio patches were the input for the audio and the video encoder. After passing through the 11 layers of both the embeddings are concatenated and passed through the single layer of the audio-video encoder. Three sets of embeddings, *i.e.* the video, audio and audio video are passed in three consecutive passes to keep the time complexity in check. The audio and video embeddings of the i th sample (a_i, v_i) are subjected to meanpool operation to compute (c_i^v, c_i^a) for further calculation of contrastive loss \mathcal{L}_c . The concatenated embeddings (\mathbf{x}') are passed through the joint decoder (JD) which has 16 attention heads and has 8 layers. Both audio and video embeddings pass through common 7 layers of JD and the last modality-specific layer for final patch reconstruction. The masked tokens are passed as trainable masked tokens, 16×16 patches attached to the computation graph, during the process. Each reconstructed token gets into its proper place due to modality-specific sinusoidal position embedding. The idea of using zero-valued patches is best explained in [4]. The reconstruction loss \mathcal{L}_r is computed between the masked patches (75% of the video frame and the spectrogram). The equations for \mathcal{L}_c and \mathcal{L}_r are as follows:

$$\mathcal{L}_c = -\frac{1}{N} \sum_{i=1}^N \log \left[\frac{\exp(s_{i,i}/\tau)}{\sum_{k \neq i} \exp(s_{i,k}/\tau) + \exp(s_{i,i}/\tau)} \right] \quad (1)$$

where $s_{i,j} = \|c_i^v\|^T \|c_j^a\|$ and τ is a temperature parameter.

$$\hat{\mathbf{a}}_i, \hat{\mathbf{v}}_i = D(\mathbf{x}' + [\mathbf{E}'_{\mathbf{a}}, \mathbf{E}'_{\mathbf{v}}] + [\mathbf{E}^{\mathbf{P}'_{\mathbf{a}}}, \mathbf{E}^{\mathbf{P}'_{\mathbf{v}}}] \quad (2)$$

$$\mathcal{L}_r = \frac{1}{N} \sum_{i=1}^N \left[\frac{\sum (\hat{\mathbf{a}}_i^\mu - \text{norm}(\mathbf{a}_i^\mu))^2}{|\mathbf{a}_i^\mu|} + \frac{\sum (\hat{\mathbf{v}}_i^\mu - \text{norm}(\mathbf{v}_i^\mu))^2}{|\mathbf{v}_i^\mu|} \right] \quad (3)$$

where N is the mini-batch size; $\mathbf{a}^\mu, \mathbf{v}^\mu, \hat{\mathbf{a}}^\mu, \hat{\mathbf{v}}^\mu$ denote the original and predicted masked patches; $|\mathbf{a}^\mu|$ and $|\mathbf{v}^\mu|$ denote the number of masked audio and visual patches, $D()$ is the joint decoder (JD), $\mathbf{E}'_{\mathbf{a}}$ and $\mathbf{E}'_{\mathbf{v}}$ are modality type embedding (for $D()$ to recognize audio and video tokens), $\mathbf{E}^{\mathbf{P}'_{\mathbf{a}}}$ and $\mathbf{E}^{\mathbf{P}'_{\mathbf{v}}}$ are modality specific 2D sinusoidal embeddings, respectively. The final loss is computed as $\mathcal{L}_{\text{total}} = \mathcal{L}_r + \lambda_c \cdot \mathcal{L}_c$ where λ_c is a regulating parameter empirically chosen as 0.01.

2.2. Complementary mask: why, what and how?

There are a few important things that we observed revisiting the closely related previous works and the reason why we proposed complementary mask:

\mathcal{L}_c only does positive pair mining, *i.e.* pull together corresponding audio and video. Thus there is no way to find the

correspondence between the modality itself, which is important as we are heavily sampling both the audio and video before encoding. *This motivates us to use self-distillation between modality.*

Due to the nature of implementation, the JD is unaware of the contrastive learning effect. This measure is necessary as the decoder will be discarded for downstream tasks and thus its only concern should be how to reconstruct the embeddings to the masked input patches, the only possible annotation available for for SSL. *Thus self distillation should be done in the encoded space like constructive task.*

The masking strategy is random and thus there is no implied correspondence between the video and audio tokens, that are passing through VE and AE respectively. As it is impossible to devise any temporal correlation between random spatial patches of a single video frame and the total audio spectrogram, thus self-distillation between video and audio to find the modular correspondence directly is impossible in the present form of data representation. However, distillation between individual modalities of audio and video and the joint embeddings makes sense as the couples belong to the same modality/category. Thus to materialise the same the duel head MAE is employed with non-overlapping or overlapping complementary patches as input to each heads.

The concept of the complementary mask is simply the generation of two masks so that M_1 and M_2 are generated as sets of unmasked tokens in such a way that the intersection of them is a null set *i.e.* $M_1 \cap M_2 = \phi$. The concept is pictorially described in Fig. 2. It should be noted that n complementary random masks are only possible for a masking ratio $m_r \geq (n/N)$ where N is the total number of input patches. In our case, $n = 2$ thus $m_r \geq 0.5$ will suffice. As the masking ratio is 0.75 thus there was no need to change the masking ratio.

At first, a normal random mask is generated with masking ratio m_r . Let the set of all tokens for audio modality be $U_a = 512$. Thus if the first mask M_1^a is considered as a set then the residue tokens left for consideration in M_2^a is $U_a - M_1^a$, thus the set M_2^a is randomly chosen from that residual set so that $|M_1^a| = |M_2^a| = m_r \cdot |U_a|$ and $M_1^a \cap M_2^a = \phi$. Similar masks M_1^v and M_2^v are generated for the video modality from U_v where $|U_v| = 196$. It should be noted that there is no relation imposed between the audio and video masks while generating the complementary mask.

The motivation behind complementary patches is making the *embedding mask agnostic to find modal correspondence, i.e.* if the embeddings are treated as probability distributions, they should be closer irrespective of the input tokens, provided that they are from the same data point. After all a masked encoder is supposed to be a function $F(x)$ where x is a subset of the input X in the form of randomly

chosen patches. Thus if two subsets x_1 and x_2 are there such as $x_1 \subset X, x_2 \subset X, x_1 \cap x_2 = \phi$ are there, $F()$ should be trained in such a way that $F(x_1) \simeq F(x_2)$. As X will be reconstructed from x_1 and x_2 by a single decoder which is shallower than the encoder, it is the task of the encoder to incorporate the entire information of all patches in the unmasked patches. Thus the best encoder is one that can generate similar input to the decoder irrespective of the input subset, even if the two subsets have no overlap at all.

2.3. Learning strategy of the proposed SSL-model

Gong et al. [16] already proved different learning strategies of SSL are not complementary. Their work demonstrated that a combination of contrastive and reconstruction loss works well towards learning a generic multi-modal encoder for SSL. However, they have not explored all combinations of SSL strategies to optimize their combination. Hence, we introduced in this context self-distillation along with contrastive-mask encoding. The idea of self-distillations works like contrastive learning here but the contrast is between two forms of the same input across all modalities and their combinations, rather than just two modalities of the same input. In the reconstruction loss, the accuracy of the masked token reconstruction is checked and is trained to the patches of the masked input patches with the help of information provided in the JD, by the encoded audio-visual embeddings. Our goal is to force the embeddings to preserve better information so that it can provide that to the trainable masked token in the JD, in turn improving the power of the encoder.

DINO is a popular model to start from when it comes to self-distillation with no labels [6]. But at any one point in time, it is unidirectional in nature. An older approach called Deep Mutual Learning(DML) [65] appeared to be a better candidate for our consideration after initial experiments. We tuned the nature of DML by redesigning the training methodology with masked data modelling in mind. For that, we wagered our approach on a symmetric form of KL divergence. The Kullback-Leibler (KL) distance (discrete version) is defined as:

$$\mathcal{D}(p_1||p_2) = \sum_{i=0}^{n-1} (p_1(i) \log \frac{p_1(i)}{p_2(i)}) \quad (4)$$

where p_1 and p_2 are two probability distributions of n distinct samples. In terms of latent embedding the n translates to the dimension size. Though KL divergence provides a very good log-likelihood measure of two distributions, the major caveat is its asymmetric property *i.e.* $\mathcal{D}(p_1||p_2) \neq \mathcal{D}(p_2||p_1)$. Also, the inputs must be a probability distribution or $\sum_{i=0}^{n-1} p(i) = 1$ and $p(i) \geq 0 \forall i$ which may not be the case for an embedding.

The embeddings between which the KL divergence is measured are the mean pool for all audio and video to-

kens that will be subjected to pass to JD (without the positional encoding). In our case we meanpool \mathbf{x}' to χ where $n = |\chi| = 768$ as the patches have dimension 16×16 with three channels. To get the embedding for KL divergence we first find the minimum value χ_{min} in the embedding vector (χ) and then sum or subtract that to the vector, *i.e.* $\chi = \langle \chi(i)(+/-)|\chi_{min}|, \forall i \rangle$, depending on whether χ_{min} is negative or positive, respectively. This is done to make sure $\chi(i) \geq 0 \forall i$.

Then χ is converted to a probability distribution by simple linear normalization *i.e.* $p = (\chi(i) / \sum_{j=0}^{n-1} \chi(j)), \forall i$. Something as common as softmax is avoided as that is tuned for classification problems, which exaggerates the divergence of the distribution for the sake of clarity. Our purpose is not classification but a near-accurate projection of χ into a probability distribution p so that KL-divergence can be applied.

Now to eliminate the asymmetric effect of KL divergence, we use the self-distillation loss (\mathcal{L}_{kd}) formula:

$$\mathcal{L}_{kd}(p_1, p_2) = \frac{\mathcal{D}(p_1||p_2) + \mathcal{D}(p_2||p_1)}{2} \quad (5)$$

Where p_1 and p_2 are the probability distribution projection of the origin corrected mean-pooled audio-video embeddings χ_1 and χ_2 coming from the parallel streams (refer Fig. ??) with input masks M_1 and M_2 , respectively.

If we consider \mathbf{x}' to be the probability distribution in a translated latent space $\mathbf{L}' \in \mathbb{R}^n$ then a properly learnt encoder $F(x)$ should compute $\mathbf{x}'_1 = F(x_1)$ and $\mathbf{x}'_2 = F(x_2)$ so that $\mathbf{x}'_1 \simeq \mathbf{x}'_2$. For MAE, the decoder $D()$ is getting trained such as $D(\mathbf{x}'_1) = D(\mathbf{x}'_2) = X$. As $D()$ is to be discarded, only the condition $\mathbf{x}'_1 \simeq \mathbf{x}'_2$ can ensure such an outcome. Thus we chose to propagate the \mathcal{L}_{kd} only through $F()$, to make it powerful without being dependent on $D()$. As it is proven [1] that model selection can be best measured with KL divergence from the ground truth distribution, in our case, due to lack of ground truth in \mathbf{L}' , mutual ground truth is the only option left. The idea is that whatever may be the true latent representation, the encoder should be able to approximate a distribution near that, irrespective of input. MAE [20] has already proved that random masking solves this purpose, thus making the encoder learn better. Hence, complementary inputs will make this divergence, as at the encoder level there is no option for information exchange between the two sets of complementary tokens.

2.4. Combining all the losses and learning objectives

We have discussed about three losses, this is how to translate to different learning objectives. Contrastive loss \mathcal{L}_c translates to the learning objective “audio and video of the same datapoint should be closer”. Reconstruction loss \mathcal{L}_r translates to “even with masked input, the deep encoder should extract enough information to use a shallow decoder

for reconstructing the total input”. Self Distillation loss \mathcal{L}_{kd} translates to “whatever may be the mask, the latent audio-video representation of a datapoint should not vary.”

A scalar factor $\lambda_c = 0.01$ for the contrastive loss \mathcal{L}_c and $\lambda_{kd} = 10$ for self-distillation loss \mathcal{L}_{kd} was used before adding to the reconstruction loss \mathcal{L}_r . Due to the nature of implementation, during backward pass \mathcal{L}_r affects the modality-specific encoders, the joint encoder and the joint decoder but the contrastive loss \mathcal{L}_c and \mathcal{L}_{kd} only effects later two. This is due to the way the graph updates weights based on the local gradient computed in the forward pass and the downstream gradient computed in the backward pass.

The self-distillation loss \mathcal{L}_{kd} works best for two mutual teacher/student (stream) configurations. For more parallel streams, the performance degrades as proven by experiments. If more than two streams are used then \mathcal{L}_{kd} is taken as the mean of all three streams *i.e.* $\mathcal{L}_{kd}(p_1, p_2, p_3) = (\mathcal{L}_{kd}(p_1, p_2) + \mathcal{L}_{kd}(p_1, p_3) + \mathcal{L}_{kd}(p_2, p_3))/3$. In such a configuration the divergence between two streams affects the third stream, to which it has no contribution. Though the higher number of student cohorts of the parallel stream is supposed to enhance accuracy [65] the trend reverses due to too much exposure to data as 3 complementary masks expose 75% of the data, which overfits the model. Thus it is proven that a complementary masking strategy is not the holy grail when it comes to training an MAE.

3. Experimental results

We performed extensive experiments and ablation studies with different learning objectives combined in different ways. Our final configuration is a result of incremental change in various key factors and those results are presented in an ordered fashion to elaborate our path of analysis. We have used 4 datasets for SSL pretraining and modality-specific finetuning. The downstream task in finetuning has been only classification across different modalities of data.

3.1. Datasets and hardware

The major two datasets that we worked with are AudioSet [14] and VGGsound [8] for training and testing the methodology on audio-visual modality. We have used the Kinetics [24] dataset for video-only action recognition. We have not experimented separately on image modality as the video frame reconstruction is the same task.

AudioSet [14], is a collection of audio label annotated YouTube clips of 10 seconds duration. Due to the unavailability of many videos, the dataset on which *we based our experiments has 1472186 training videos and 7486 test videos (for label classification)*. The total number of classes is 527 that are not distributed in a balanced fashion. Our use of AudioSet was mainly in SSL pretraining. A balanced subset called AudioSet 20k with 8227 datapoints is used for finetuning tasks. Like the full Audioset, many videos are

missing from the original list, in the subset also.

VGGsound [8], is a also collection of audio label annotated YouTube clips of 10 seconds duration like AudioSet. For the same reason as AudioSet, we were able to experiment with 164003 training videos and 13560 test videos. The total number of classes is 300 and the samples are not balanced over classes.

Kinetics-400 [24] is also a huge collection of 10 second YouTube clips where only the video modality is used for human action recognition. We have used 230k training videos and 19k validation and testing videos.

ILSVRC2017 [42] is a subset of the ImageNet dataset consisting of 1,281,167 training images, 50,000 validation images and 100,000 test images, all belonging to one of 1000 classes.

For all the dataset numbers may vary from the officially released number due to the aforementioned reasons of working with a collection of YouTube videos. Regarding hardware, we have used multiple A100-80 and A6000-48 cards in PyTorch dataparallel mode [31].

3.2. Implementation details

For pretraining we have used a batch size of 120 and for finetuning it is 48 for VGGsound and 36 for AudioSet20K. The learning rate starts from 0.00005 and starts decreasing by half every five epochs, after the 15th epoch. For finetuning, the learning rate starts from 0.0001 and is halved every epoch starting from 2nd epoch. Only a few experiments were done on the much larger AudioSet dataset, due to its sheer size and thus delay in training. In the tables, AudioSet is abbreviated as AS, and VGGsound is abbreviated as VGG. The metric of measurement is classification accuracy in every form of experiment.

3.3. Experimental results

Table 1 shows three rows the first being the results of finetuning on different combinations of both datasets VGG sound and AS-20K on the CAV-MAE. The second one is Moving Average mutual Knowledge Distillation (MAKD) [17], where instead of having shared weights in the two streams the one with lesser loss (teacher) propagated its weight to the one with higher loss (student), for a certain iteration. This role reverses dynamically and thus the “mutual” terms come in play. This was our primary attempt towards incorporating KD as the third pillar of SSL after masked data modelling and contrastive learning. The third shows the best result obtained by our proposed ensemble technique of dual complementary mask and self-distillation-based approach. It can be observed that our ensembled method showed a jump in accuracy over the vanilla technique in the multimodal configuration and the video-only configuration in finetuning. It can also be seen

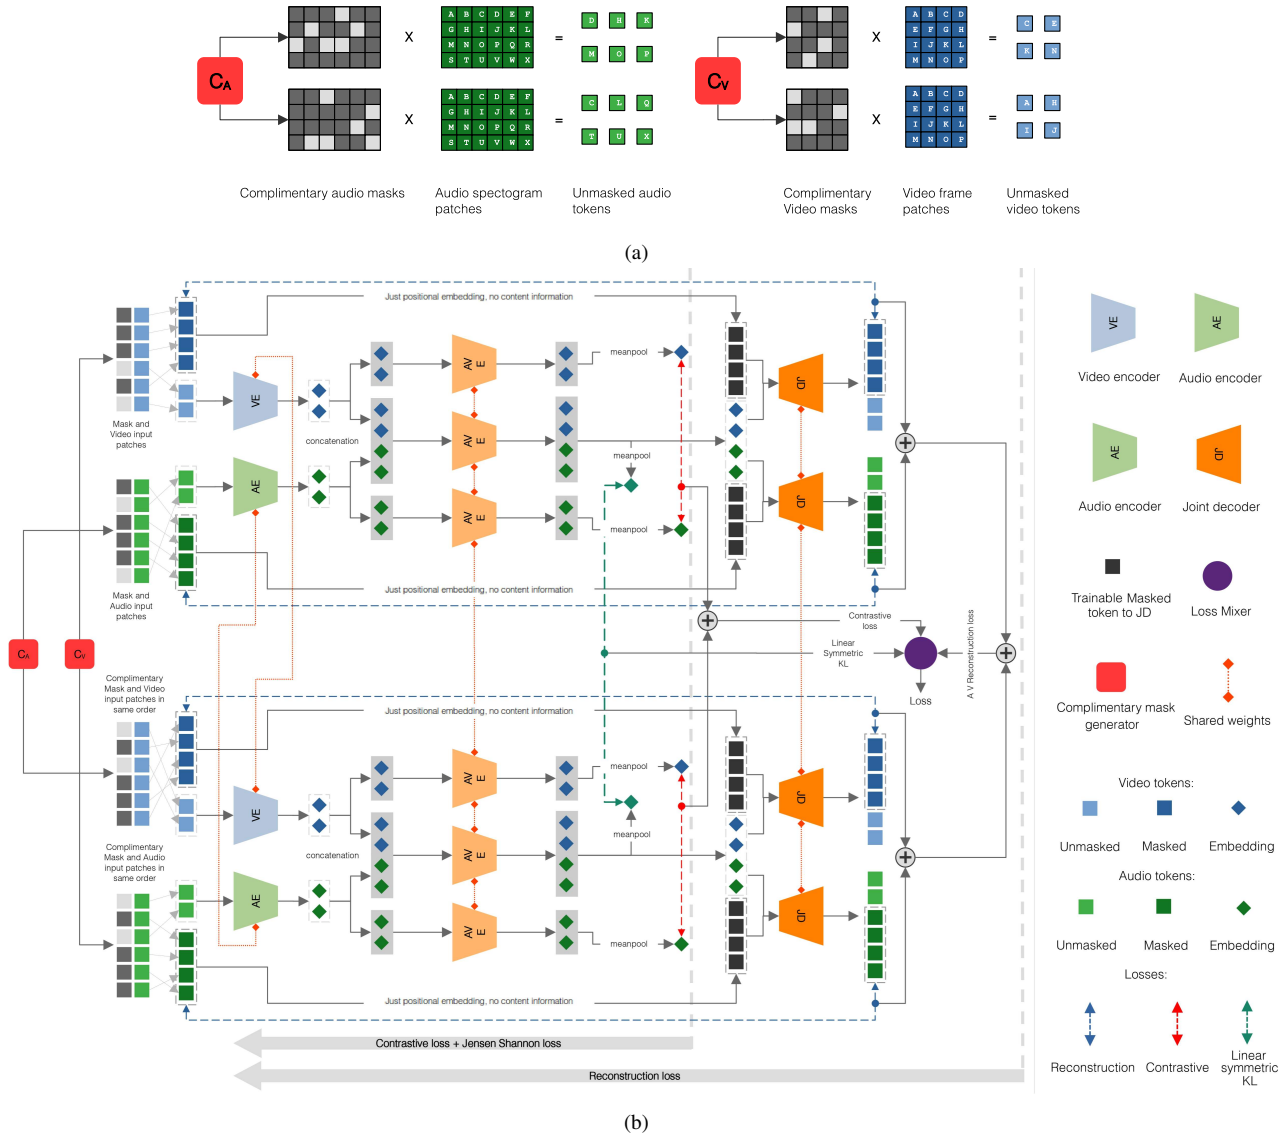


Figure 2. (a) The complementary mask generation. Note: Two video masks and audio masks are complementary but there is no relation between any two audio and video masks, (b) The proposed architecture of KDC-MAE: (Symbol index on the right)

Table 1. Comparison of finetuning accuracy (MAKD is Moving Average mutual Knowledge Distillation.)

	VGG FT AV	VGG FT A only	VGG FT V only	AS-20k FT AV	AS-20k FT A only	AS-20k FT V only
CAV-MAE	63.89	58.6	43.20	39.61	37.95	31.00
MAKD	62.44	59.2	43.37	40.80	37.70	30.86
Proposed	64.23	58.73	43.43	41.03	37.62	30.95

that for AudioSet the proposed method performed better for AV classification.

The final loss propagated with our ensemble method is either determined by dynamic weights as learnable parameters during training or by manually assigning the weights λ_c and λ_{kd} for the contrastive and self-distillation losses respectively. As per CAV-MAE authors, $\lambda_c = 0.01$ is an optimal choice as per their study so we kept it unchanged and experimented with various values of λ_{kd} , as shown in Table 2. By overlapped and non-overlapped masks, we mean n random masks M_i so that $M_i \cap M_j = \phi \forall i, j, i \neq j$ and $M_i \cap M_j = \phi \forall i, j, i \neq j$, respectively. We observe that for $\lambda_{kd} = 10$ with complementary masking, we achieve the best result among all our experiments. Note that in con-

figurations where VGGsound results were not satisfactory, the experiments were not repeated for AudioSet. In this regard, it must be clarified that we have used the abbreviation “kd” for self-distillation. Also, the result of adaptive masking strategy [2] is compared with our complementary masking strategy. In that work, the mask is adaptively generated based on the input image.

Table 2. Results of different datasets (for description refer Table 1). The first row shows two random masks with no restrictions on being complementary and $\lambda_{kd} = 10$. In the second row onwards, masks are complementary with different values for λ_{kd} .

	VGG FT AV	VGG FT A only	VGG FT V only	AS-20k FT AV	AS-20k FT A only	AS-20k FT V only
Adaptive Masking [2]	63.65	58.5	43.2	41.23	37.11	31.01
With Overlap Dual Mask $\lambda_{kd}=10$	63.85	58.77	43.68	41.37	37.72	31.59
No Overlap Dual Mask $\lambda_{kd}=10$	64.23	58.73	43.43	41.03	37.62	30.95
No Overlap Dual Mask $\lambda_{kd}=15$	63.58	58.85	43.71	41.13	36.57	31.12
Proposed No Overlap Dual Mask $\lambda_{kd}=5$	64.01	58.84	43.51	41.62	38.42	31.25

The self-distillation strategy is about measuring the divergence between the latent vectors out of the joint encoder and then propagating that as a loss back through the joint encoder, audio and video encoders. Though while training a multimodal encoder, the joint encoding matters the most, we experimented with the divergences between the outputs of the audio encoders and video encoders of both streams. Our motivation was to study the fact that whether a better joint representation learning can be enhanced by ensuring better modality-specific learning. In Table 3 we have the results with the configurations of overlapping dual mask and non-overlapping dual mask in row 1 and 2 respectively for the different dataset combinations. We also experimented with the same configurations with dynamic weights replacing λ_c and λ_{kd} . In Fig. 2, the violet round denotes the loss mixer, which is named so that the mixing can be either done with manual weights or dynamic weights. The dynamic weights are learnable parameters in the PyTorch computation graph that get updated by the global loss. Row 3 and 4 show the results with dynamic weight in Table 3, where row 4 is the configuration with dynamic weight and divergence on just the joint encoder. It can be observed that dynamic weight delivers better accuracy when distillation between all pairs of encoders (AE, VE, AVE) are considered.

As discussed in the later part of Section 2.4, the consideration of more than two parallel streams behaves counter-intuitively as proven by experimental results in Table 4. It can be observed that for three parallel streams, complementary masks perform worse than overlapping masks. The analyzed reason for this behaviour is due to too much data

Table 3. Results of different datasets (for description refer Table 1) when trained with dual masks and self distillation between the joint encoder (AVE) and also between the audio and video encoders (AE and VE) of the two streams (referred to as AA and VV), respectively. (DW means Dynamic Weight)

	VGG FT AV	VGG FT A only	VGG FT V only	AS-20k FT AV	AS-20k FT A only	AS-20k FT V only
With Overlap Dual Mask AA VV	63.98	58.79	43.62	41.42	35.60	31.44
No Overlap Dual Mask AA VV	64.05	58.78	43.73	41.01	37.18	30.94
No Overlap Dual Mask AA VV DW	64.06	59.04	43.31	41.56	37.57	31.58
No Overlap Dual Mask DW	63.82	58.90	43.83	40.98	38.28	31.87

Table 4. Ablation study to demonstrate the effect of multiple student self-distillation techniques (for description refer Table 1).

	VGG FT AV	VGG FT A only	VGG FT V only	AS-20k FT AV	AS-20k FT A only	AS-20k FT V only
Triple Mask No Overlap $\lambda_{kd} = 10$	63.79	58.60	43.53	40.78	38.27	31.74
Triple Mask With Overlap $\lambda_{kd} = 10$	64.18	59.00	43.85	41.55	38.43	30.70

Table 5. Finetune accuracy (%) on different datasets

	CAV-MAE	Dual Mask With Overlap	Dual Mask Without Overlap
AS-2M (V only)	30.26	32.59	32.27
AS-2M (A only)	41.71	42.27	42.51
AS-2M (AV)	41.85	43.55	42.73
AS-20k (V only)	28.86	30.83	31.57
AS-20k (A only)	37.21	37.92	37.01
AS-20k (AV)	37.02	41.22	41.34
ILSVRC (V only)	73.73	73.57	73.76
Kinetics-400 (V only)	65.74	66.19	65.78
Kinetics-400 (AV)	71.55	71.53	71.69

being exposed, especially for the video modality, which overfits the model gradually. For the same reason overlapping masks perform better as the amount of exposed data is bound to be lesser than complementary masks if we take uniform random sampling for keeping unmasked tokens. One interesting observation is that for audioset the triple-stream technique fares better with audio-only finetuning too. The reason is that a larger dataset mitigates overfitting and due to the comprehensive nature of the audio data (spectrogram is for the entire clip, not a single frame sample from a clip of 300 frames) the model gets trained better (contradicting videos cannot nullify the knowledge of audio modality). In this regard, it should also be mentioned

Table 6. Results on retrieval, inpainting and visual sound source localization tasks on the VGG-Sound dataset

Model	Retrieval						Inpainting		Localization
	Audio \rightarrow Visual			Visual \rightarrow Audio			Loss A	Loss V	Avg Cos Sim
	R @ 1	R @ 5	R @ 10	R @ 1	R @ 5	R @ 10			
CAV-MAE	0.0580	0.1599	0.2142	0.0662	0.1917	0.2537	2.1869	2.9981	0.2884
With Overlap	0.0594	0.1629	0.2245	0.0402	0.1367	0.1994	2.2948	3.0026	0.2807
No Overlap	0.0506	0.1431	0.1944	0.0416	0.1304	0.1734	1.9985	2.8756	0.2798

Table 7. Results of different masking strategies on VGG-Sound

Model	Classification Accuracy (%)			
	Time	Frequency	TF	Uniform
CAV-MAE	63.58	63.77	63.51	63.89
With Overlap	64.91	64.90	64.21	63.98
No Overlap	64.83	64.68	64.01	64.05

that the audio-only finetuning fares much better than video-only finetuning as the datasets are annotated for audio and not video. The same can also be observed from the results in Table 5, which reports the finetuning accuracy on more datasets like ILSRVC, Kinetics-400 and Audioset 20k subset. Again, it can be seen that classification accuracy only improves when using the proposed dual mask method in all datasets.

Another aspect that has been explored is the evaluation of downstream tasks such as retrieval, localization and inpainting as shown in Table 6. It can be seen that the proposed model performs better in the case of audio-to-visual retrieval and the case of inpainting and is quite competitive for localisation tasks. Yet another important observation can be made by comparing the performance improvement with structured masking vs uniform masking as described in Table 7. It can be concluded that uniform unstructured masking provides the best result when compared to time, frequency or time-frequency(TF) based methods. More results on downstream tasks, masking and visualization are reported in the supplementary material.

In order to demonstrate the efficacy of using self-distillation along with the complementary masking strategy for SSL pretraining, we have fine-tuned various model configurations on the Kinetics dataset. We pre-trained all the encoders (*i.e.* AE, VE and AVE) on Kinetics and we finetuned for human action recognition downstream task. The model was initiated with AudioSet pre-trained weights as the starting point. The results are **71.90%**, **71.78%** and **71.69%** for no overlap, with overlap and CAVMAE. We also tested on ILSVRC, with the accuracy of **73.73%**, **74.53%** and **74.76%** for CAVMAE, overlap and nonoverlap respectively. These results prove that pretraining with a task-specific dataset makes the encoder learn well (which is obvious!) and self-distillation in some form (complementary masked or overlapping masked input pairs) makes the

encoder learn better, unequivocally. It can be observed that the overlapped dual mask performed better than the complementary mask as for Kinetics the annotation is based on video rather than audio, thus overexposure of video modality caused it to overfit.

From the experiments conducted, we can summarize our findings about our proposed model in the following points: complementary masks have performed well or generated competitive results depending on the downstream task, especially on multimodal datasets. As the video sampling is not as uniform as audio, complementary masks don't scale well when the downstream annotation is based on video instead. However, for audio-based annotation in downstream tasks, complementary masking scales up the results elegantly. Moreover, the induction of knowledge distillation with KL-divergence on the complementary brings in the modular correspondence. This makes the encoder struggle to learn mask agnostic modelling and the model starts to learn fine modality-specific information. Further, along with the mask encoder that forces encoding the input information in the fusion and the constructive task that helps to find the explicit audio-visual correspondence, the knowledge-distillation helps to find the modular correspondence as results provide better modelling.

4. Conclusion

In the quest for finding an optimal learning strategy for SSL on large datasets, we ended up exploring many untrodden nooks and corners of knowledge distillation and masked data modelling. Masked autoencoders first leveraged the idea of masked data modelling on image modality with reconstruction from incomplete data as the sole learning objective. Later CAV-MAE elevated that for training a task-agnostic multimodal learner by incorporating intermodal contrastive learning as an extra learning objective. In our proposed learning strategy we have demonstrated that adding knowledge distillation, driven by divergence between embeddings from two inputs subjected to complementary masks, elevated the result still further. We have compared our results with the state-of-the-art methodology for the task, CAV-MAE and found that the proposed KDC-MAE achieved better results. This work opened the door for many paths to explore and future versions for more extensive studies.

References

- [1] D Anderson and K Burnham. Model selection and multi-model inference. *Second. NY: Springer-Verlag*, 63(2020):10, 2004. 4
- [2] Wele Gedara Chaminda Bandara, Naman Patel, Ali Ghomami, Mehdi Nikkhah, Motilal Agrawal, and Vishal M Patel. Adamae: Adaptive masking for efficient spatiotemporal learning with masked autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14507–14517, 2023. 7
- [3] Silvia Bucci, Antonio D’Innocente, Yujun Liao, Fabio M Carlucci, Barbara Caputo, and Tatiana Tommasi. Self-supervised learning across domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5516–5528, 2021. 1
- [4] Shuhao Cao, Peng Xu, and David A Clifton. How to understand masked autoencoders. *arXiv preprint arXiv:2202.03670*, 2022. 3
- [5] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019. 1
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 1, 4
- [7] Krishna Chaitanya, Ertunc Erdil, Neerav Karani, and Ender Konukoglu. Contrastive learning of global and local features for medical image segmentation with limited annotations. *Advances in neural information processing systems*, 33:12546–12558, 2020. 1
- [8] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggssound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 721–725. IEEE, 2020. 5
- [9] Yabo Chen, Yuchen Liu, Dongsheng Jiang, Xiaopeng Zhang, Wenrui Dai, Hongkai Xiong, and Qi Tian. Sdae: Self-distilled masked autoencoder. In *European Conference on Computer Vision*, pages 108–124. Springer, 2022. 1
- [10] Julien Denize, Jaonary Rabarisoa, Astrid Orcesi, Romain Hérault, and Stéphane Canu. Similarity contrastive estimation for self-supervised soft contrastive learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2706–2716, 2023. 1
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [12] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3299–3309, 2021. 1
- [13] Christoph Feichtenhofer, Yanghao Li, Kaiming He, et al. Masked autoencoders as spatiotemporal learners. *Advances in neural information processing systems*, 35:35946–35958, 2022. 1
- [14] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 776–780. IEEE, 2017. 5
- [15] Mariana-Iuliana Georgescu, Eduardo Fonseca, Radu Tudor Ionescu, Mario Lucic, Cordelia Schmid, and Anurag Arnab. Audiovisual masked autoencoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16144–16154, 2023. 1
- [16] Yuan Gong, Andrew Rouditchenko, Alexander H Liu, David Harwath, Leonid Karlinsky, Hilde Kuehne, and James Glass. Contrastive audio-visual masked autoencoder. *arXiv preprint arXiv:2210.07839*, 2022. 1, 2, 4
- [17] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 1, 5
- [18] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. 1
- [19] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *Advances in Neural Information Processing Systems*, 33:5679–5690, 2020. 1
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 1, 4
- [21] Bingkun Huang, Zhiyu Zhao, Guozhen Zhang, Yu Qiao, and Limin Wang. Mgmoe: Motion guided masking for video masked autoencoding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13493–13504, 2023. 1
- [22] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Let there be color! joint end-to-end learning of global and local image priors for automatic image colorization with simultaneous classification. *ACM Transactions on Graphics (ToG)*, 35(4):1–11, 2016. 1
- [23] Ashish Jainwal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2020. 1
- [24] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola,

- Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5
- [25] Dahun Kim, Donghyeon Cho, Donggeun Yoo, and In So Kweon. Learning image representations by completing damaged jigsaw puzzles. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 793–802. IEEE, 2018. 1
- [26] Joshua Knights, Ben Harwood, Daniel Ward, Anthony Vanderkop, Olivia Mackenzie-Ross, and Peyman Moghadam. Temporally coherent embeddings for self-supervised video representation learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 8914–8921. IEEE, 2021. 1
- [27] Giorgos Kordopatis-Zilos, Giorgos Toliás, Christos Tzelepis, Ioannis Kompatsiaris, Ioannis Patras, and Symeon Papadopoulos. Self-supervised video similarity learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4755–4765, 2023. 1
- [28] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 577–593. Springer, 2016. 1
- [29] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6874–6883, 2017. 1
- [30] Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34:309–323, 2021. 1
- [31] Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, et al. Pytorch distributed: Experiences on accelerating data parallel training. *arXiv preprint arXiv:2006.15704*, 2020. 5
- [32] Shuo Liu, Adria Mallol-Ragolta, Emilia Parada-Cabaleiro, Kun Qian, Xin Jing, Alexander Kathan, Bin Hu, and Bjorn W Schuller. Audio self-supervised learning: A survey. *Patterns*, 3(12), 2022. 1
- [33] Guillaume Lorré, Jaonary Rabarisoa, Astrid Orcesi, Samia Ainouz, and Stéphane Canu. Temporal contrastive pretraining for video action recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 662–670, 2020. 1
- [34] Andriy Mnih and Koray Kavukcuoglu. Learning word embeddings efficiently with noise-contrastive estimation. *Advances in neural information processing systems*, 26, 2013. 1
- [35] Duy MH Nguyen, Hoang Nguyen, Truong TN Mai, Tri Cao, Binh T Nguyen, Nhat Ho, Paul Swoboda, Shadi Albarqouni, Pengtao Xie, and Daniel Sonntag. Joint self-supervised image-volume representation learning with intra-inter contrastive clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 14426–14435, 2023. 1
- [36] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016. 1
- [37] Mehdi Noroozi, Ananth Vinjimoor, Paolo Favaro, and Hamed Pirsiavash. Boosting self-supervised learning via knowledge transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9359–9367, 2018. 1
- [38] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1
- [39] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huiheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6964–6974, 2021. 1
- [40] Rajat Raina, Alexis Battle, Honglak Lee, Benjamin Packer, and Andrew Y Ng. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th international conference on Machine learning*, pages 759–766, 2007. 1
- [41] Bin Ren, Yahui Liu, Yue Song, Wei Bi, Rita Cucchiara, Nicu Sebe, and Wei Wang. Masked jigsaw puzzle: A versatile position embedding for vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20382–20391, 2023. 1
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 5
- [43] Mohammadreza Salehi, Ainaz Eftekhari, Niusha Sadjadi, Mohammad Hossein Rohban, and Hamid R Rabiee. Puzzle-ae: Novelty detection in images through solving puzzles. *arXiv preprint arXiv:2008.12959*, 2020. 1
- [44] Madeline C Schiappa, Yogesh S Rawat, and Mubarak Shah. Self-supervised learning for videos: A survey. *ACM Computing Surveys*, 55(13s):1–37, 2023. 1
- [45] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pages 843–852. PMLR, 2015. 1
- [46] Li Tao, Xueting Wang, and Toshihiko Yamasaki. Self-supervised video representation learning using inter-intra contrastive framework. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2193–2201, 2020. 1
- [47] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022. 1
- [48] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for

- video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. **1**
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **2**
- [50] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. *Advances in neural information processing systems*, 29, 2016. **1**
- [51] Guodong Wang, Yunhong Wang, Jie Qin, Dongming Zhang, Xiuguo Bao, and Di Huang. Video anomaly detection by solving decoupled spatio-temporal jigsaw puzzles. In *European Conference on Computer Vision*, pages 494–511. Springer, 2022. **1**
- [52] Jinpeng Wang, Yuting Gao, Ke Li, Jianguo Hu, Xinyang Jiang, Xiaowei Guo, Rongrong Ji, and Xing Sun. Enhancing unsupervised video representation learning by decoupling the scene and the motion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10129–10137, 2021. **1**
- [53] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae v2: Scaling video masked autoencoders with dual masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14549–14560, 2023. **1**
- [54] Lin Wang and Kuk-Jin Yoon. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE transactions on pattern analysis and machine intelligence*, 44(6):3048–3068, 2021. **1**
- [55] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luowei Zhou, and Lu Yuan. Bevt: Bert pretraining of video transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14733–14743, 2022. **1**
- [56] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14668–14678, 2022. **1**
- [57] Chen Wei, Lingxi Xie, Xutong Ren, Yingda Xia, Chi Su, Jiaying Liu, Qi Tian, and Alan L Yuille. Iterative reorganization with weak spatial constraints: Solving arbitrary jigsaw puzzles for unsupervised representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1910–1919, 2019. **1**
- [58] Bin Xia, Jingwen He, Yulun Zhang, Yitong Wang, Yapeng Tian, Wenming Yang, and Luc Van Gool. Structured sparsity learning for efficient video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22638–22647, 2023. **1**
- [59] Haosen Yang, Deng Huang, Bin Wen, Jiannan Wu, Hongxun Yao, Yi Jiang, Xiatian Zhu, and Zehuan Yuan. Self-supervised video representation learning with motion-aware masked autoencoders. *arXiv preprint arXiv:2210.04154*, 2022. **1**
- [60] Ting Yao, Yiheng Zhang, Zhaofan Qiu, Yingwei Pan, and Tao Mei. Seco: Exploring sequence supervision for unsupervised representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10656–10664, 2021. **1**
- [61] Liangzhe Yuan, Rui Qian, Yin Cui, Boqing Gong, Florian Schroff, Ming-Hsuan Yang, Hartwig Adam, and Ting Liu. Contextualized spatio-temporal contrastive learning with self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13977–13986, 2022. **1**
- [62] Chaoning Zhang, Chenshuang Zhang, Junha Song, John Seon Keun Yi, and In So Kweon. A survey on masked autoencoder for visual self-supervised learning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 6805–6813, 2023. **1**
- [63] Chaoning Zhang, Chenshuang Zhang, Junha Song, John Seon Keun Yi, Kang Zhang, and In So Kweon. A survey on masked autoencoder for self-supervised learning in vision and beyond. *arXiv preprint arXiv:2208.00173*, 2022. **1**
- [64] Qi Zhang, Yifei Wang, and Yisen Wang. How mask matters: Towards theoretical understandings of masked autoencoders. *Advances in Neural Information Processing Systems*, 35:27127–27139, 2022. **1**
- [65] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4320–4328, 2018. **1, 2, 4, 5**
- [66] Zexian Zhou and Xiaojing Liu. Masked autoencoders in computer vision: A comprehensive survey. *IEEE Access*, 2023. **1**