

Noise Estimation Using Density Estimation for Self-Supervised Multimodal Learning

Elad Amrani^{1,2}, Rami Ben-Ari¹, Daniel Rotman¹, Alex Bronstein²

¹IBM Research AI

²Technion

Abstract

One of the key factors of enabling machine learning models to comprehend and solve real-world tasks is to leverage multimodal data. Unfortunately, annotation of multimodal data is challenging and expensive. Recently, self-supervised multimodal methods that combine vision and language were proposed to learn multimodal representations without annotation. However, these methods often choose to ignore the presence of high levels of noise and thus yield sub-optimal results. In this work, we show that the problem of noise estimation for multimodal data can be reduced to a multimodal density estimation task. Using multimodal density estimation, we propose a noise estimation building block for multimodal representation learning that is based strictly on the inherent correlation between different modalities. We demonstrate how our noise estimation can be broadly integrated and achieves comparable results to state-of-the-art performance on five different benchmark datasets for two challenging multimodal tasks: Video Question Answering and Text-To-Video Retrieval. Furthermore, we provide a theoretical probabilistic error bound substantiating our empirical results and analyze failure cases. Code: <https://github.com/elad-amrani/ssml>.

1 Introduction

Multimodal learning is a well established methodology for tackling complex and challenging artificial intelligence tasks such as Visual Question Answering (Antol et al. 2015; Jang et al. 2017; Gao et al. 2018; Xu et al. 2017; Fan et al. 2019) and Text-to-Video Retrieval (Liu et al. 2019; Mithun et al. 2018; Yu, Kim, and Kim 2018; Miech, Laptev, and Sivic 2018; Song and Soleymani 2019). The motivation for glean- ing information from multiple correlated data sources comes from how we as humans perceive the world and learn from ex- perience. Using the correlation between speech and vision, a person is able to recognize objects by their names while learn- ing the visual characteristics. Additionally, concepts can be learned separately and a combination can be comprehended automatically, e.g., ‘running’ and ‘beach’ vs. ‘running on the beach’.

Manual annotation of large-scale datasets and specifically multimodal datasets is challenging and expensive. This dif- ficulty results in a shortage which limits the progress of

supervised machine learning and has become the key devel- opment bottleneck. Recently, to combat costs and effort of annotation, self-supervised machine learning (Zhang, Isola, and Efros 2016; Noroozi and Favaro 2016; Pathak et al. 2016; Misra, Zitnick, and Hebert 2016; Wei et al. 2018; Vondrick, Pirsiavash, and Torralba 2016; Srivastava, Mansimov, and Salakhudinov 2015) presents new ways to better utilize the abundant unlabeled data on the web. However, most self- supervised systems aim to learn from a single data modality, which limits their applicability.

In contrast to the above, (Miech et al. 2019, 2020; Amrani et al. 2019; Moriya et al. 2019; Sun et al. 2019b,a) recently showed that unlabeled instructional videos could be used as training data for a self-supervised multimodal learning system due to the high correlation between the spoken word and the ongoing visuals. Unfortunately, such systems are forced to deal with high noise levels and thus yield sub- optimal results as we show in this paper.

In this paper, we propose a novel noise robust multimodal representation learning building block for self-supervised learning. We utilize the inherent correlation between different modalities for efficient multimodal learning in the presence of extreme levels of noise. Specifically, we show that noise estimation can be reduced to a density estimation problem. We define a multimodal similarity function and show that based on this function, noise is correlated with sparsity and vice versa.

Ultimately, we integrate our proposed building block into an embedding model and learn superior joint video-text rep- resentations that achieve comparable state-of-the-art perfor- mance on five datasets: MSRVT (Xu et al. 2016), LSMDC (Rohrbach et al. 2015), MSVD (Chen and Dolan 2011), MSRVT-QA (Xu et al. 2017) and MSVD-QA (Xu et al. 2017); for two different tasks: Video Question Answering and Text to Video Retrieval. Additionally, we provide a theo- retical probabilistic error bound substantiating our empirical results and analyze failure cases.

Contributions. The key contributions of this paper are four fold:

1. We show that the problem of noise estimation for mul- timodal data can be efficiently reduced to a multimodal density estimation task.
2. We propose a novel building block for noise-robust mul- timodal representation learning and demonstrate its integra-

tion into the max margin ranking loss function.

3. We demonstrate comparable state-of-the-art performance on five datasets for two different challenging multimodal tasks by utilizing our approach for self-supervised multimodal learning with the HowTo100M dataset (Miech et al. 2019).
4. We substantiate our empirical results with a theoretical analysis of the proposed method that includes a probabilistic error bound.

2 Related Work

Self-Supervised Learning. Self-supervised learning methods strive to learn informative data representations by defining and solving a pretext task. In these tasks, pseudo labels can be generated automatically and compact data representations must be learned in order to solve these tasks. Many pretext tasks were proposed in recent years: colorizing grayscale images (Zhang, Isola, and Efros 2016), image jigsaw puzzle (Noroozi and Favaro 2016), image inpainting (Pathak et al. 2016), video frame order verification (Misra, Zitnick, and Hebert 2016), video future prediction (Vondrick, Pirsaviash, and Torralba 2016; Srivastava, Mansimov, and Salakhutdinov 2015), audio-visual correspondence (Korbar, Tran, and Torresani 2018; Arandjelovic and Zisserman 2017), speech-visual correspondence (Miech et al. 2019, 2020; Amrani et al. 2019; Moriya et al. 2019; Sun et al. 2019b,a), etc. In this work, we focus on speech-visual correspondence in unlabeled instructional videos, where speech is converted to text using an automatic speech recognition system. Speech-visual correspondence is considered a difficult pretext task due to extremely noisy pseudo labels, yet it can be a highly advantageous task since it provides semantic information of visual features in the form of natural text. Such valuable information can be utilized to solve many challenging multimodal downstream tasks as we show in Section 6.

Multimodal Representation Learning. The word *modality* refers to a particular form of sensory perception, such as the visual and auditory modalities. A machine learning task or dataset is said to be *multimodal* when it includes a number of modalities. Multimodal representation learning frameworks can be divided into three types: (a) joint representation which aims to learn a shared semantic subspace (Salakhutdinov and Hinton 2009; Srivastava and Salakhutdinov 2012; Antol et al. 2015); (b) an encoder-decoder framework which aims to translate from one modality into another and keep their semantics consistent (Mao et al. 2014; Rohrbach, Rohrbach, and Schiele 2015; Venugopalan et al. 2015; Reed et al. 2016); and (c) coordinated representation which aims to learn separated yet coordinated representations for each modality under some constraints (Weston, Bengio, and Usunier 2010; Frome et al. 2013; Socher et al. 2014; Wang, Li, and Lazebnik 2016; Vendrov et al. 2016; Tian, Krishnan, and Isola 2019). In this work, we focus on coordinated representations that enforce similarity among them. Our goal is to enforce the multimodal representations of similar ‘concepts’ to be close to each other. E.g., a video of a man running on the beach should be close in representation to the textual representation of ‘a man running

on the beach’ as opposed to ‘a man cooking in the kitchen’. The multimodal representation described above is highly valuable for solving multimodal machine learning tasks. If a machine learning model learns to link between the visuals and text of specific concepts it should be able, for example, to answer natural language questions about visual content, or do cross-modal retrieval more easily (Section 5.2).

Density Estimation. The aim of density estimation is to estimate the probability density function underlying the data, which is assumed to be *i.i.d.* Existing density estimation algorithms can be divided into two categories: (a) parametric or semi-parametric approaches such as Gaussian Mixture models (McLachlan and Krishnan 2007; Wang and Wang 2015) and probabilistic graphical models (Koller and Friedman 2009); and (b) non-parametric approaches such as histograms, Splines (Stone 1994), neural network-based density estimation (Papamakarios, Pavlakou, and Murray 2017) and Kernel Density Estimation (Silverman 2018; Terrell and Scott 1992). For an extended review on density estimation for high-dimensional data see (Wang and Scott 2019). In this work, we utilize *multimodal* k -Nearest Neighbor density estimation, which is a special case of Kernel Density Estimation. With it, we form a novel noise-robust multimodal representation learning model.

Learning with Noisy Data. Learning with noisy data can be divided into two approaches: (a) formulating explicit or implicit noise models to characterize the distribution of noisy and true labels using neural networks (Goldberger and Ben-Reuven 2017; Jiang et al. 2018; Sukhbaatar et al. 2015), graphical models (Xiao et al. 2015; Li et al. 2017), etc. and (b) using correction methods. E.g., relabeling the data during training (Reed et al. 2015), jointly optimizing the model’s parameters and estimating true labels (Tanaka et al. 2018), using noise-tolerant loss function (Ghosh, Kumar, and Sastry 2017; Van Rooyen, Menon, and Williamson 2015) or noise tolerant training algorithms (Li et al. 2019). However, these methods do not deal with multimodal association label (Definition 1) and often require a small set of data with clean labels to be available. In this work, we propose a true label estimation method for multimodal data that does not require availability of clean labels. We base our estimation on the correlation between modalities alone.

3 Method

3.1 Motivation

In multimodal data, a sample is said to be noisy when two or more modalities do not share the same semantic meaning. For example, a video-text pair that is associated with each other, yet the text is not related to the ongoing visuals. Existing multimodal embedding models are susceptible to such noisy data, i.e., the model is likely to adjust itself to the noise in the data and thus yield sub-optimal results. This scenario is very common in the case of self-supervised multimodal learning and even when learning from unlabeled instructional videos. Although in these instructional videos there is some correlation between caption (speech transcription) and vision,

unfortunately often a person is talking about something that is not present visually. For example, in the HowTo100M dataset (Miech et al. 2019), the authors manually inspected 400 randomly sampled clip-caption pairs and found that in about half there was not a single object or action mentioned in the caption that was also visually present in the video clip. To deal with noise, we suggest to utilize the inherent correlation between the different modalities that is based on the Definition and Assumption below. See Fig. 1 for a visualization and a detailed explanation.

Definition 1. A *correctly (wrongly) associated pair* is a clip-caption pair (v, c) that share (do not share) the same semantic meaning or concept, i.e., the caption c describes (does not describe) the ongoing visuals v .

Assumption 1 (Mixture Model). The distributions of the videos and captions can be represented using a general mixture model of T components in the corresponding modality. Denoting by $a, b \in \{1, \dots, T\}$, respectively, the concept to which the video v and the caption c belong, we can write $v|a \sim p_v(v|a)$ and $c|b \sim p_c(c|b)$.

If Assumption 1 holds, then *correctly associated* pairs form dense clusters in both modalities that contain pairs that are also associated with each other (see Fig. 1a). Thus, by defining a multimodal similarity function (i.e., a similarity measure between pairs), we can formulate the task of finding *correctly associated* pairs simply as a multimodal density estimation task. In this formulation, pairs in dense areas will be more likely to be *correctly associated*, while pairs in sparse areas will be more likely to be *wrongly associated* (see Fig. 1b).

3.2 Notation and Problem Formulation

Let $\{(v_i, c_i) \in \mathbb{R}^{d_v} \times \mathbb{R}^{d_c}\}_{i=1}^M$ denote the set of clip-caption pairs, where for each i , the video clip v_i is associated with the caption sentence c_i , and M denotes the size of the dataset. Let $p_i \in \{0, 1\}$ denote a binary indicator for whether the pair (v_i, c_i) is *correctly associated* ($p_i = 1$) or *wrongly associated* ($p_i = 0$). Let $f_v : \mathbb{R}^{d_v} \rightarrow \mathbb{R}^d$ and $f_c : \mathbb{R}^{d_c} \rightarrow \mathbb{R}^d$ denote the embedding functions of the videos and the captions, respectively, into a common representation space. The task of *noise robust* multimodal representation learning aims to map all of the data modalities to a single embedding space such that for all v_i that is *correctly associated* with c_i , $f_v(v_i) \approx f_c(c_i)$ in the sense of some similarity function.

3.3 Noise Estimation Using Multimodal Density Estimation

For the ease of notation, we will denote the pair as $z_i = (v_i, c_i)$. Let us define a similarity function between pairs, $S : \mathbb{R}^{d_v+d_c} \times \mathbb{R}^{d_v+d_c} \rightarrow \mathbb{R}$.

$$S(z_i, z_j) \triangleq \min \left\{ \frac{s(v_i, v_j) - \bar{\mu}_v}{\bar{\sigma}_v}, \frac{s(c_i, c_j) - \bar{\mu}_c}{\bar{\sigma}_c} \right\}, \quad (1)$$

where s can be, for example, the cosine similarity function $s(x, y) = \frac{x^\top y}{\|x\| \|y\|}$; $\bar{\mu}_v, \bar{\mu}_c$ and $\bar{\sigma}_v, \bar{\sigma}_c$ are the sample means and standard deviations of each modality, i.e., the similarity values of each modality are normalized before taking the

minimum. Using (1), a pair z_i is close to z_j only if v_i is close to v_j and c_i is close to c_j as well.

We denote by \hat{p}_i the estimated probability of z_i being *correctly associated*, and compute it using its local k -NN density estimation normalized such that $\hat{p}_i \in [0, 1]$:

$$\hat{p}_i \triangleq \frac{\bar{S}_i - \min(\{\bar{S}_i\}_{i=1}^M)}{\max(\{\bar{S}_i\}_{i=1}^M) - \min(\{\bar{S}_i\}_{i=1}^M)}, \quad (2)$$

where,

$$\bar{S}_i = \frac{1}{K} \sum_{k=1}^K S(z_i, z_{i_k}), \quad i \in [M], \quad (3)$$

z_{i_k} is the k -th nearest neighbor of z_i and S is the similarity function defined in (1).

3.4 Soft Max Margin Ranking Loss

Integrating our noise estimation component from above into a max margin ranking loss function (Wang et al. 2014; Schroff, Kalenichenko, and Philbin 2015) is straightforward. We weight each pair z_i with its fixed estimated probability \hat{p}_i of being *correctly associated*. We call it *Soft Max Margin Ranking*:

$$L_{soft-rank} = \sum_{i \in P} \left(\hat{p}_i \sum_{j \in N_i} \max\{0, s_{ij} - s_{ii} + \delta\} + \max\{0, s_{ji} - s_{ii} + \delta\} \right), \quad (4)$$

where, P is the set of noisy associated (positive) pairs, N_i is the set of negative pairs for clip-caption pair (v_i, c_i) , \hat{p}_i is defined in (2), s_{ij} is the similarity score between the embedding of the clip-caption pair $(f_v(v_i), f_c(c_j))$, and δ is the margin. The first term in the equation above is for matching a video with a negative caption and the second term is for matching a caption with a negative video.

We note that a different integration approach is to discard samples with \hat{p}_i below a certain threshold. However, for real-data experiments we found the performance to be substantially worse.

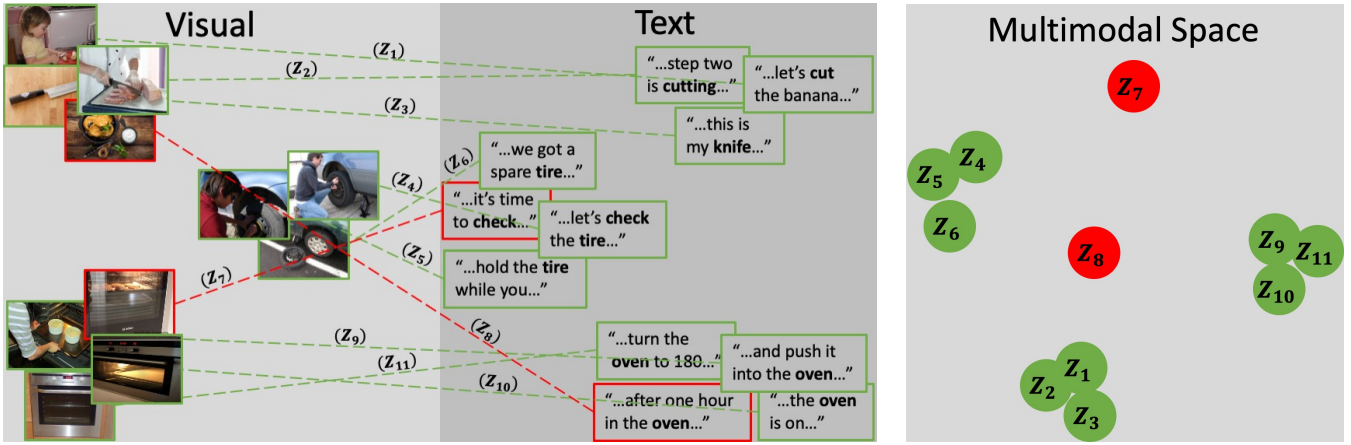
4 Theoretical Analysis

We present a theoretical probabilistic error upper bound of our noise estimation approach. For simplicity we assume the data is distributed under a Gaussian Mixture model.

4.1 Probabilistic Error Upper Bound

Theorem 1. Let $Z = (X, Y) \in \mathbb{R} \times \mathbb{R}$ be a random pair of scalars satisfying Assumption 1 for a Gaussian mixture of $T > 1$ equi-probable concepts. Denoting by $(A, B) \in \{1, \dots, T\}^2$ the concepts to which the pair Z belongs, $X|A = a \sim \mathcal{N}(\mu_a, \sigma_a^2)$ and $Y|B = b \sim \mathcal{N}(\mu'_b, \sigma'_b{}^2)$. We further assume that each component of the mixture is 6σ -separated, i.e., $|\mu_i - \mu_j| > 6\sigma_{\max}$ and $|\mu'_i - \mu'_j| > 6\sigma'_{\max}$ for every $i \neq j$, where $\sigma_{\max} \triangleq \max_t \{\sigma_t\}$ and $\sigma'_{\max} \triangleq \max_t \{\sigma'_t\}$.

Let $\Pr(A = a \neq B = b) = \frac{\eta}{T(T-1)}$ and $\Pr(A = B = a) = \frac{1-\eta}{T}$ for every $a, b \in \{1, \dots, T\}$. The binary indicator



(a) Multimodal data visualization. Each initial monomodal embedding space contains somewhat dense clusters of ‘concepts’, where a ‘concept’ could be a specific object or action (e.g., ‘cutting’, ‘knife’, ‘check’, ‘tire’, ‘oven’, etc.). It is likely that *correctly associated* (Definition 1) pairs form dense clusters in both modalities that contain pairs that are also associated with each other and of the same ‘concept’ (GREEN, $z_1 - z_6, z_9 - z_{11}$). In contrast, a *wrongly associated* (Definition 1) pair may still belong to dense clusters in both modalities but those clusters are not likely to contain pairs that are associated with each other (RED, z_7 and z_8). Best viewed in color.

(b) Multimodal space defined by (1). Each point above ($\{z_i\}_{i=1}^{11}$) represents a single pair from the left sub-figure. The distance between points that is visualized is computed based on (1). Given Assumption 1, in the multimodal space above, *correctly associated* pairs are correlated with high density and vice-versa. Best viewed in color.

Figure 1: Noise estimation using multimodal density estimation

P denoting whether the pair (X, Y) is correctly associated consequently has $Pr(P = 1) = Pr(A = B) = 1 - \eta$, where η is the noise ratio of the dataset.

Let $\{z_i = (x_i, y_i)\}_{i=1}^M$ be a finite sample of pairs drawn independently from the described model, and let \bar{S}_i be the average similarity between z_i and its K nearest neighbors as defined in (3), with $S(z_i, z_j) \triangleq \frac{s(x_i, x_j) + s(y_i, y_j)}{2}$, and $s(x, x') \triangleq -|x - x'|$. Then, the following bounds hold for every τ and $t > 0$,

$$P(\bar{S}_i \geq \tau \mid p_i = 0) \leq \frac{\mathcal{M}_{\bar{S}_i \mid p_i=0}(t)}{e^{t\tau}}, \quad (5)$$

$$P(\bar{S}_i \leq \tau \mid p_i = 1) \leq \frac{\mathcal{M}_{\bar{S}_i \mid p_i=1}(-t)}{e^{-t\tau}}, \quad (6)$$

where, $\mathcal{M}_{\bar{S}_i \mid p_i}(t)$ is the moment generating function of $\bar{S}_i \mid p_i$ defined in the extended version (Amrani et al. 2020).

The proof is provided in the extended version (Amrani et al. 2020). It is important to remark that for the simplicity of analysis we assumed the pairs to be formed of scalars. While, from the first glance, this assumption might severely limit the possible configurations of the concepts in each of the modalities, in our analysis we made no assumptions whatsoever on the way the concepts are collocated in space, except the 6σ -separation that can hold in any number of dimensions. The validity of the presented analysis in the multidimensional case is corroborated by the toy dataset example in Section 6.1.

4.2 Numerical Simulations and Analysis

In this section, we present numerical simulations of the probabilistic error bounds in (5) and (6). Our goal is to gain

insight into: (a) why the method works; (b) the effect of the design choice (K) and dataset properties (η, M, T) on the performance of the model; and (c) analyze possible failure cases. More specifically, we: (a) set $\tau = \tau^*$ such that $P(\bar{S}_i \geq \tau^* \mid p_i = 0) = P(\bar{S}_i \leq \tau^* \mid p_i = 1)$, i.e., $\tau^* = \sqrt{\mathcal{M}_{\bar{S}_i \mid p_i=0}(t) \cdot \mathcal{M}_{\bar{S}_i \mid p_i=1}(-t)}$; (b) sweep a single parameter at a time, while the rest are fixed; and (c) optimize for t over $[1, 100]$.

Discussion. In Fig. 2a we study the effect of K . As expected, increasing K decreases the error bound initially and from a certain value ($\triangleq K_0$), the error bound increases. Not surprisingly, $K_0 \approx \frac{M}{T} \cdot (1 - \eta)$, which is the average number of correctly associated pairs per concept. Throughout Figures 2a, 2b, 2c, 2d we see the error bound is influenced greatly by this equality, i.e., when $K > K_0$, the model performs well and when $K \leq K_0$, it starts to fail. Specifically, in Fig. 2c we show that the error bound goes to zero as the size of the dataset (M) increases (another point of view is that K_0 is increased). For this reason we mark the point where $KT = M(1 - \eta)$ by a red dashed line in Fig. 2. It is clear that for real-world data, concepts are usually not equi-probable and thus assigning a global value for K is sub-optimal. However, this finding allows us to better understand such failure cases and thus choose a reasonable K value. An additional instance where the method fails, gives us insight into why usually the method succeeds. The method will fail for a small number of concepts (T) regardless of K_0 , because for a small number of concepts there is a higher chance that two or more *wrongly associated* pairs belong to the same pair of concepts in both modalities. Fortunately, in real-world data T is almost always large, and additionally as T increases, this problem is




Figure 2: Numerical simulations of our probabilistic error upper bound. From left to right: sweep over K , η , M and T . We mark with a red dashed vertical line where $K = \frac{M}{T} \cdot (1 - \eta)$ holds.

alleviated by a factor of $\mathcal{O}(T^2)$ (see extended version (Amrani et al. 2020)). A simulation of this case is presented in the extended version (Amrani et al. 2020).

5 Experimental Settings

5.1 Implementation Details

Model. For a fair comparison to the baseline model HTM (Miech et al. 2019), we use the same class of non-linear embedding functions: $f(\mathbf{v}) = (W_1^v \mathbf{v} + b_1^v) \circ \sigma(W_2^v (W_1^v \mathbf{v} + b_1^v) + b_2^v)$, $g(\mathbf{c}) = (W_1^c \mathbf{c} + b_1^c) \circ \sigma(W_2^c (W_1^c \mathbf{c} + b_1^c) + b_2^c)$, where $W_1^v \in \mathbb{R}^{d \times d_v}$, $W_1^c \in \mathbb{R}^{d \times d_c}$, $W_2^v, W_2^c \in \mathbb{R}^{d \times d}$, $b_1^v, b_2^v, b_1^c, b_2^c \in \mathbb{R}^d$ are the learnable parameters, σ is an element-wise sigmoid activation and \circ is the element-wise multiplication. We use $d_v = 4096$, $d_c = 300$, and $d = 6144$.

Training dataset. We train our model using the HowTo100M (Miech et al. 2019) narrated video dataset that consists of more than 1.2M videos accompanied with automatically generated speech transcription. Similarly to (Miech et al. 2019), we use the provided transcription to create pairs of clip-caption defined by each caption time stamp, where each video clip shorter than 5 seconds is extended symmetrically in time so that the duration is at least 5 seconds. Note that we only use 1.16M videos since some of the videos are no longer available for download.

Input caption features. For the word representations, we use the standard GoogleNews pre-trained word2vec embedding model (Mikolov et al. 2013). For the input sentence representations used in Section 3.3 we simply average word representation over each sentence.

Input visual features. We extract 2D features using ImageNet pre-trained Resnet-152 (He et al. 2016) at a rate of 1 frame per second. We extract 3D features using Kinetics (Carreira and Zisserman 2017) pre-trained ResNeXt-101 16-frames (Hara, Kataoka, and Satoh 2018) at a rate of 24 frames per second. After temporal max pooling we concatenate 2D and 3D features to form a single feature vector per video clip.

Loss & Optimization. We train our model using the Soft Max Margin loss function described in Section 3.4. We use the ADAM (Kingma and Ba 2015) optimizer with a fixed learning rate of 10^{-3} .

Time complexity. Using FAISS (Johnson, Douze, and Jégou 2019), computation of the Multimodal Density Estimation described in Section 3.3 is done in less than 15 hours over 10 CPUs. Training the model on the large HowTo100M

dataset is done on a single V100 GPU and takes less than 24 hours.

Additional implementation details are included in the extended version (Amrani et al. 2020).

5.2 Downstream Tasks

Video Visual Question Answering (VQA). The Video VQA task comprises answering questions about videos presented in natural language (Antol et al. 2015). Essentially, an instance of VQA includes an input video and a free-form textual query regarding the content in the video, and an expected textual answer. To accommodate this task we fine-tune our learned multimodal representations and evaluate our model on two datasets: MSRVT-TQA and MSVD-QA (Xu et al. 2017). These datasets are based on existing video description datasets. See extended version (Amrani et al. 2020) for detailed statistics of each dataset.

Most VQA models use a video and question as input, and the answer is presented as the output of an LSTM unit (Hochreiter and Schmidhuber 1997) or a softmax layer over a set of predetermined answers. However, these types of architectures do not fully utilize the information which exists in coordinated representations, i.e., the representation of the correct answer might likely be closely embedded to the visual representation, given the question. To better utilize our learned multimodal representations specifically for the VQA task, we use a similar architecture to (Hu, Chao, and Sha 2018), but for video. We learn two more sets of embeddings on top of the pre-trained embeddings that were learned with the HowTo100M dataset: a question+video embedding, i.e., we embed a concatenation of the question and video to a single feature vector; and an answer embedding. We train the model with a max margin ranking loss function to embed an answer close to its question+video. Inference is performed simply with a nearest neighbor search over the set of predetermined answers in the joint video+question and answer space. This model is very simple compared to most VQA models, yet as we show in Table 2 it is very powerful when built on effective self-supervised pre-trained joint embeddings.

Text-To-Video Retrieval. Text-To-Video Retrieval includes retrieval of video clips based on textual description (Liu et al. 2019; Mithun et al. 2018; Song and Soleymani 2019). With a learned joint representation space, retrieval is performed with a nearest neighbor search over the joint embedding space. To evaluate our model we use three different datasets: MSRVT-T, MSVD and LSMDC (Xu et al. 2016;




Figure 3: Toy dataset Visualizations. (a, c): T-SNE of input embeddings. As we can see, separating between samples with $p_i = 1$ and samples with $p_i = 0$ is non-trivial. (b): Cumulative Distribution Function of \hat{p}_i of toy dataset. The green solid line is the CDF of $\hat{p}_i | p_i = 1$, while the red dashed line is the inverse CDF of $\hat{p}_i | p_i = 0$. Assuming a binary prediction is made based on a hard threshold, it is possible to extract the precision and recall for each threshold from the figure above. For example, for the threshold 0.48, both precision and recall are $\simeq 0.9$. Best viewed in color.

Chen and Dolan 2011; Rohrbach et al. 2015). We use the standard evaluation metrics: recall at K ($R@K$) for $K = 1, 5, 10$ and median recall (MR). See extended version (Amrani et al. 2020) for detailed statistics of each dataset.

6 Experiments and Analysis

A toy dataset illustrative results are presented in Section 6.1. Comparison to ablative baselines and state-of-the-art models is presented in Section 6.2. A design choice analysis and qualitative examples are presented in the extended version (Amrani et al. 2020).

6.1 Multidimensional Toy Dataset

In this section, we demonstrate the effectiveness of our method visually using a toy (synthesized) dataset of mixture of Gaussians of T components in each modality, $\{\mathcal{N}(\mu_t^c, \Sigma_t^c)\}_{t=1}^T$ and $\{\mathcal{N}(\mu_t^v, \Sigma_t^v)\}_{k=1}^T$ for caption and video, respectively. *Correctly associated* pairs are represented by $v_i \sim \mathcal{N}(\mu_t^v, \Sigma_t^v)$, $c_i \sim \mathcal{N}(\mu_t^c, \Sigma_t^c)$, such that $t \in [T], i \in [M]$. *Wrongly associated* pairs are represented by $v_i \sim \mathcal{N}(\mu_m^v, \Sigma_m^v)$, $c_i \sim \mathcal{N}(\mu_n^c, \Sigma_n^c)$, such that $m \neq n$, $\{m, n\} \in [T], i \in [M]$. $\eta \in [0, 1]$ is the noise ratio, such that a *wrongly associated* pair is sampled with probability η , and a *correctly associated* pair is sampled with probability $1 - \eta$.

In our experiment, $v_i \in \mathcal{R}^{d_v}$, $c_i \in \mathcal{R}^{d_c}$, $d_v = 128$, $d_c = 128$; $\mu_t^c \in \mathcal{R}^{d_c}$, $\mu_t^v \in \mathcal{R}^{d_v}$, $\forall t \in [T]$ are sampled from a uniform multivariate distribution $U_{d_c}(0, 1)$, $U_{d_v}(0, 1)$, respectively for caption and video; $\Sigma_t^c \in \mathbb{R}^{d_c \times d_c}$, $\Sigma_t^v \in \mathbb{R}^{d_v \times d_v}$ are diagonal matrices where the diagonals are sampled from a multivariate uniform distribution, $U_{d_c}(0, 0.3)$, $U_{d_v}(0, 0.3)$, respectively for caption and video; $\eta = 0.5$, $M = 1250$, $T = 50$; $k = 4$ (k -NN parameter).

In Figures 3a and 3c we visualize T-SNE graphs for caption and video embedding spaces, respectively. In Fig. 3b we visualize the empirical cumulative distribution function of \hat{p}_i (2), the estimated probability of being *correctly associated*. Additionally, in the extended version (Amrani et al. 2020) we empirically reproduce the theoretical graphs in Fig. 2

using the multidimensional toy dataset. These results corroborate the validity of the analysis presented in Thm. 1 in the multidimensional case.

6.2 Ablative Baselines and SOTA Models

We compare our proposed model against two ablative baselines and multiple task-specific state-of-the-art (SOTA) models:

HTM-PT (Miech et al. 2019). The model (architecture and loss function) used in (Miech et al. 2019). This baseline is pre-trained (PT) on the HowTo100M dataset. It is the exact architecture described in Section 5.1 for our model. The only differentiating element is the loss function. (Miech et al. 2019) use the max margin ranking loss function, while we use our proposed Soft Max Margin ranking loss function. Since the model is also trained identically to our own model it is clear that any gain in performance over this baseline is due to our novel noise estimation based density estimation component.

HTM-no-PT (Miech et al. 2019). The same model from above, but without pre-training (no-PT) on the HowTo100M dataset. The (under) performance of this baseline on downstream tasks demonstrates the potential gain of utilizing self-supervised speech-visual correspondence training.

Task specific state-of-the-art models. After fine-tuning for downstream tasks we compare our proposed model to state-of-the-art models for each task and each dataset. (Gao et al. 2018; Xu et al. 2017; Fan et al. 2019; Jang et al. 2017) for VQA, and (Liu et al. 2019; Mithun et al. 2018; Yu, Kim, and Kim 2018; Miech, Laptev, and Sivic 2018; Miech et al. 2019, 2020) for Text-To-Video Retrieval.

Tables 1 and 2 show the result for Text-To-Video Retrieval and Video Question Answering, respectively. Results for Zero-Shot Text-To-Video Retrieval under ‘unfair’ settings are in the extended version (Amrani et al. 2020). We summarize key insights below:

- Our model consistently outperforms the baselines (HTM-PT, HTM-no-PT (Miech et al. 2019)) in both Visual Question Answering and Text-To-Video Retrieval on five differ-

Method	MSRVTT				LSMDC				MSVD			
	R@1	R@5	R@10	MR	R@1	R@5	R@10	MR	R@1	R@5	R@10	MR
Zero-Shot												
Random	0.1	0.5	1.0	500.0	0.1	0.5	1.0	500.0	0.15	0.75	1.49	335
MIL-NCE [†] (Miech et al. 2020)	9.9	24.0	32.4	29.5	–	–	–	–	–	–	–	–
HTM-PT* (Miech et al. 2019)	7.5	21.2	29.6	38.0	4.0	9.8	14.0	137.0	12.86	33.06	45.83	13.0
Ours	8.0	21.3	29.3	33.0	4.2	11.6	17.1	119.0	13.66	35.7	47.74	12.0
Fine-Tuned												
CE [‡] (Liu et al. 2019)	18.2	46.0	60.7	7.0	11.2	26.9	34.8	25.0	19.8	49.0	63.8	6.0
JEMC (Mithun et al. 2018)	7.0	20.9	29.7	38.0	–	–	–	–	20.3	47.8	61.1	6.0
JSFusion (Yu, Kim, and Kim 2018)	10.2	31.2	43.2	13.0	9.1	21.2	34.1	36	–	–	–	–
MoEE (Miech, Laptev, and Sivic 2018)	14.2	39.2	53.8	9	10.1	25.6	34.6	27	–	–	–	–
HTM-no-PT (Miech et al. 2019)	12.4	36.0	52.0	10.0	5.8	18.8	28.4	45.0	13.0	37.43	52.41	10.0
HTM-PT* (Miech et al. 2019)	14.9	40.2	52.8	9.0	7.1	19.6	27.9	40.0	15.52	40.93	55.7	8.0
Ours	17.4	41.6	53.6	8.0	6.4	19.8	28.4	39.0	20.3	48.97	63.26	6.0

Table 1: Text-To-Video retrieval. Zero-Shot: training was done only with HowTo100M dataset. Fine-Tuned: model was fine-tuned with the relevant benchmark dataset. For MR the lower the better. *: results for MSRVTT and LSMDC are from (Miech et al. 2019), while results for MSVD have been reproduced. †: MIL-NCE (Miech et al. 2020) use additional clip-caption pairs. ‡: CE (Liu et al. 2019) use extra labeled data in the form of pre-trained semantic embeddings which include ‘general’ features such as motion, appearance, scene features and OCR.

	MSRVTT-QA [%]	MSVD-QA [%]
ST-VQA (2017)	30.09	31.3
Co-Mem (2018)	32.0	31.7
AMU (2017)	32.5	32.0
HMEMA (2019)	33.0	33.7
HTM-no-PT (2019)	27.05	33.8
HTM-PT (2019)	34.38	34.83
Ours	35.06	35.13

Table 2: Video Question Answering. Results of ST-VQA and Co-Mem taken from (Fan et al. 2019)

ent datasets.

- We set a new state-of-the-art performance for two Visual Question Answering datasets: MSRVTT-QA and MSVD-QA.
- We set a new state-of-the-art performance for Zero-Shot Text-To-Video Retrieval on two datasets: LSMDC and MSVD. For MSRVTT, we hypothesize that MIL-NCE (Miech et al. 2020) performs better due to two reasons: i) they train end-to-end, unlike our baseline model, which operates on pre-processed embeddings; and ii) they utilize additional clip-caption pairs by associating a clip with its adjacent (in time) captions. In fact, when those additional pairs are not used our model outperforms theirs.
- We set a new state-of-the-art performance for (fine-tuned) Text-To-Video Retrieval on two datasets: MSRVTT and MSVD.
- We demonstrate that our model outperforms or is at least on par with the performance of HTM-PT (Miech et al. 2019) even given a setting which is a clear disadvantage

such as training it without 3D features (i.e., only 2D). See extended version (Amrani et al. 2020). This shows:

- (i) The power of our noise estimation method and its potential.
- (ii) Integrating our multimodal density estimation component allows saving time and/or computation power by training and running inference with only 2D features, without (or with minor) performance degradation.

7 Summary

In this work, we showed that the problem of noise estimation in multimodal data can be effectively reduced to a multimodal density estimation task. Based on this efficient noise estimation we proposed a novel building block for noise robust multimodal representation learning that can be integrated into many multimodal learning models and improve their performance instantly. We demonstrated how to integrate our building block into the max margin ranking loss function (Soft Max Margin) and it can similarly be integrated into various architectures and losses. We trained Soft Max Margin on the self-supervised proxy task of speech-visual correspondence that is known to be highly noisy. We further evaluated Soft Max Margin on two different downstream tasks: Visual Question Answering and Text-to-Video Retrieval; and achieved comparable state-of-the-art performance on five different datasets. For supporting the empirical results and analyzing failure cases, we provided a theoretical probabilistic error bound. These results emphasize the importance of self-supervised multimodal representation learning for advancing the state of the art in challenging multimodal artificial intelligence tasks.

References

- Amrani, E.; Ben-Ari, R.; Hakim, T.; and Bronstein, A. 2019. Learning to Detect and Retrieve Objects From Unlabeled Videos. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, 3713–3717. IEEE.
- Amrani, E.; Ben-Ari, R.; Rotman, D.; and Bronstein, A. 2020. Noise estimation using density estimation for self-supervised multimodal learning. *arXiv preprint arXiv:2003.03186*.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Lawrence Zitnick, C.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.
- Arandjelovic, R.; and Zisserman, A. 2017. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, 609–617.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6299–6308.
- Chen, D. L.; and Dolan, W. B. 2011. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, 190–200. Association for Computational Linguistics.
- Fan, C.; Zhang, X.; Zhang, S.; Wang, W.; Zhang, C.; and Huang, H. 2019. Heterogeneous memory enhanced multimodal attention model for video question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1999–2007.
- Frome, A.; Corrado, G. S.; Shlens, J.; Bengio, S.; Dean, J.; Ranzato, M.; and Mikolov, T. 2013. Devise: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, 2121–2129.
- Gao, J.; Ge, R.; Chen, K.; and Nevatia, R. 2018. Motion-appearance co-memory networks for video question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6576–6585.
- Ghosh, A.; Kumar, H.; and Sastry, P. 2017. Robust loss functions under label noise for deep neural networks. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Goldberger, J.; and Ben-Reuven, E. 2017. Training deep neural networks using a noise adaptation layer. In *ICLR*.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 6546–6555.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8): 1735–1780.
- Hu, H.; Chao, W.-L.; and Sha, F. 2018. Learning answer embeddings for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5428–5436.
- Jang, Y.; Song, Y.; Yu, Y.; Kim, Y.; and Kim, G. 2017. Tgif-qa: Toward spatio-temporal reasoning in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2758–2766.
- Jiang, L.; Zhou, Z.; Leung, T.; Li, L.; and Fei-Fei, L. 2018. Mentor-Net: Learning Data-Driven Curriculum for Very Deep Neural Networks on Corrupted Labels. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, 2309–2318. PMLR.
- Johnson, J.; Douze, M.; and Jégou, H. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR 2015*.
- Koller, D.; and Friedman, N. 2009. *Probabilistic graphical models: principles and techniques*. MIT press.
- Korbar, B.; Tran, D.; and Torresani, L. 2018. Cooperative learning of audio and video models from self-supervised synchronization. In *Advances in Neural Information Processing Systems*, 7763–7774.
- Li, J.; Wong, Y.; Zhao, Q.; and Kankanhalli, M. S. 2019. Learning to learn from noisy labeled data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5051–5059.
- Li, Y.; Yang, J.; Song, Y.; Cao, L.; Luo, J.; and Li, L.-J. 2017. Learning from noisy labels with distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, 1910–1918.
- Liu, Y.; Albanie, S.; Nagrani, A.; and Zisserman, A. 2019. Use What You Have: Video Retrieval Using Representations From Collaborative Experts. In *British Machine Vision Conference*.
- Mao, J.; Xu, W.; Yang, Y.; Wang, J.; Huang, Z.; and Yuille, A. 2014. Deep captioning with multimodal recurrent neural networks (m-rnn). *arXiv preprint arXiv:1412.6632*.
- McLachlan, G. J.; and Krishnan, T. 2007. *The EM algorithm and extensions*, volume 382. John Wiley & Sons.
- Miech, A.; Alayrac, J.-B.; Smaira, L.; Laptev, I.; Sivic, J.; and Zisserman, A. 2020. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9879–9889.
- Miech, A.; Laptev, I.; and Sivic, J. 2018. Learning a text-video embedding from incomplete and heterogeneous data. *arXiv preprint arXiv:1804.02516*.
- Miech, A.; Zhukov, D.; Alayrac, J.-B.; Tapaswi, M.; Laptev, I.; and Sivic, J. 2019. Howto100M: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the IEEE International Conference on Computer Vision*, 2630–2640.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Misra, I.; Zitnick, C. L.; and Hebert, M. 2016. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, 527–544. Springer.
- Mithun, N. C.; Li, J.; Metze, F.; and Roy-Chowdhury, A. K. 2018. Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, 19–27.
- Moriya, Y.; Sanabria, R.; Metze, F.; and Jones, G. J. 2019. Grounding Object Detections With Transcriptions. In *ICML Workshop 2019*.
- Norozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, 69–84. Springer.

- Papamakarios, G.; Pavlakou, T.; and Murray, I. 2017. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, 2338–2347.
- Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2536–2544.
- Reed, S.; Akata, Z.; Yan, X.; Logeswaran, L.; Schiele, B.; and Lee, H. 2016. Generative adversarial text to image synthesis. *arXiv preprint arXiv:1605.05396*.
- Reed, S.; Lee, H.; Anguelov, D.; Szegedy, C.; Erhan, D.; and Rabinovich, A. 2015. Training deep neural networks on noisy labels with bootstrapping. In *ICLR Workshop 2015*.
- Rohrbach, A.; Rohrbach, M.; and Schiele, B. 2015. The long-short story of movie description. In *German conference on pattern recognition*, 209–221. Springer.
- Rohrbach, A.; Rohrbach, M.; Tandon, N.; and Schiele, B. 2015. A dataset for movie description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3202–3212.
- Salakhutdinov, R.; and Hinton, G. 2009. Deep boltzmann machines. In *Artificial intelligence and statistics*, 448–455.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 815–823.
- Silverman, B. W. 2018. *Density estimation for statistics and data analysis*. Routledge.
- Socher, R.; Karpathy, A.; Le, Q. V.; Manning, C. D.; and Ng, A. Y. 2014. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics* 2: 207–218.
- Song, Y.; and Soleymani, M. 2019. Polysemous visual-semantic embedding for cross-modal retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1979–1988.
- Srivastava, N.; Mansimov, E.; and Salakhutdinov, R. 2015. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, 843–852.
- Srivastava, N.; and Salakhutdinov, R. 2012. Learning representations for multimodal data with deep belief nets. In *International conference on machine learning workshop*, volume 79.
- Stone, C. J. 1994. The use of polynomial splines and their tensor products in multivariate function estimation. *The Annals of Statistics* 118–171.
- Sukhbaatar, S.; Estrach, J. B.; Paluri, M.; Bourdev, L.; and Fergus, R. 2015. Training convolutional networks with noisy labels. In *ICLR Workshop 2015*.
- Sun, C.; Baradel, F.; Murphy, K.; and Schmid, C. 2019a. Contrastive bidirectional transformer for temporal representation learning. *arXiv preprint arXiv:1906.05743*.
- Sun, C.; Myers, A.; Vondrick, C.; Murphy, K.; and Schmid, C. 2019b. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 7464–7473.
- Tanaka, D.; Ikami, D.; Yamasaki, T.; and Aizawa, K. 2018. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5552–5560.
- Terrell, G. R.; and Scott, D. W. 1992. Variable kernel density estimation. *The Annals of Statistics* 1236–1265.
- Tian, Y.; Krishnan, D.; and Isola, P. 2019. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*.
- Van Rooyen, B.; Menon, A.; and Williamson, R. C. 2015. Learning with symmetric label noise: The importance of being unhinged. In *Advances in Neural Information Processing Systems*, 10–18.
- Vendrov, I.; Kiros, R.; Fidler, S.; and Urtasun, R. 2016. Order-Embeddings of Images and Language. In Bengio, Y.; and LeCun, Y., eds., *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Venugopalan, S.; Xu, H.; Donahue, J.; Rohrbach, M.; Mooney, R.; and Saenko, K. 2015. Translating Videos to Natural Language Using Deep Recurrent Neural Networks. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1494–1504.
- Vondrick, C.; Pirsaviash, H.; and Torralba, A. 2016. Generating videos with scene dynamics. In *Advances in neural information processing systems*, 613–621.
- Wang, J.; Song, Y.; Leung, T.; Rosenberg, C.; Wang, J.; Philbin, J.; Chen, B.; and Wu, Y. 2014. Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1386–1393.
- Wang, L.; Li, Y.; and Lazebnik, S. 2016. Learning deep structure-preserving image-text embeddings. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5005–5013.
- Wang, X.; and Wang, Y. 2015. Nonparametric multivariate density estimation using mixtures. *Statistics and Computing* 25(2): 349–364.
- Wang, Z.; and Scott, D. W. 2019. Nonparametric density estimation for high-dimensional data—Algorithms and applications. *Wiley Interdisciplinary Reviews: Computational Statistics* 11(4): e1461.
- Wei, D.; Lim, J. J.; Zisserman, A.; and Freeman, W. T. 2018. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8052–8060.
- Weston, J.; Bengio, S.; and Usunier, N. 2010. Large scale image annotation: learning to rank with joint word-image embeddings. *Machine learning* 81(1): 21–35.
- Xiao, T.; Xia, T.; Yang, Y.; Huang, C.; and Wang, X. 2015. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2691–2699.
- Xu, D.; Zhao, Z.; Xiao, J.; Wu, F.; Zhang, H.; He, X.; and Zhuang, Y. 2017. Video question answering via gradually refined attention over appearance and motion. In *Proceedings of the 25th ACM international conference on Multimedia*, 1645–1653.
- Xu, J.; Mei, T.; Yao, T.; and Rui, Y. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5288–5296.
- Yu, Y.; Kim, J.; and Kim, G. 2018. A joint sequence fusion model for video question answering and retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 471–487.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *European conference on computer vision*, 649–666. Springer.