
VIDLANKD: Improving Language Understanding via Video-Distilled Knowledge Transfer

Zineng Tang Jaemin Cho Hao Tan Mohit Bansal
UNC Chapel Hill
{tterrnan, jmincho, haotan, mbansal}@cs.unc.edu

Abstract

Since visual perception can give rich information beyond text descriptions for world understanding, there has been increasing interest in leveraging visual grounding for language learning. Recently, vokenization [68] has attracted attention by using the predictions of a text-to-image retrieval model as labels for language model supervision. Despite its success, the method suffers from approximation error of using finite image labels and the lack of vocabulary diversity of a small image-text dataset. To overcome these limitations, we present VIDLANKD, a video-language knowledge distillation method for improving language understanding. We train a multi-modal teacher model on a video-text dataset, and then transfer its knowledge to a student language model with a text dataset. To avoid approximation error, we propose to use different knowledge distillation objectives. In addition, the use of a large-scale video-text dataset helps learn diverse and richer vocabularies. In our experiments, VIDLANKD achieves consistent improvements over text-only language models and vokenization models, on several downstream language understanding tasks including GLUE, SQuAD, and SWAG. We also demonstrate the improved world knowledge, physical reasoning, and temporal reasoning capabilities of our model by evaluating on the GLUE-diagnostics, PIQA, and TRACIE datasets. Lastly, we present comprehensive ablation studies as well as visualizations of the learned text-to-video grounding results of our teacher and student language models.¹

1 Introduction

Language learning can be aided by grounded visual cues, as they provide powerful signals for modeling a vastness of experiences in the world that cannot be documented by text alone [5; 29; 4]. While the recent trend of large-scale language model pretraining indirectly provides some world knowledge from text, most large text corpora (e.g., Wikipedia) do not provide enough multi-modal grounding information. Previous works have explored multiple ways of grounding language to visual information such as constructing a common vector space [38; 7] and supervising the model with token-wise generated vision labels [68]. However, the widely-used image-text datasets (e.g., MS COCO [48]) are much smaller than text-only corpora in terms of word counts and vocabulary diversity for language learning.

The recent method of ‘vokenization’ [68] is a promising initial step towards addressing this problem by supervising language models with weakly-aligned vision-language groundings. Firstly, an image-text matching model retrieves a corresponding image to each text token in a sentence. Then a language model learns to predict the selected image (called ‘voken’) for each text token. This can be seen as a knowledge distillation (KD) process [33] from a vision-language grounding model to a language

¹Code and models: <https://github.com/zinengtang/VidLanKD>




Figure 1: Overview of the proposed VIDLANKD method. We first pretrain a teacher language model on a multi-modal dataset (Sec. 3.2). Then we distill the knowledge of the teacher model (weights frozen) to a student language model on a text dataset (Sec. 3.3).

model. Although the voken classification task helps the language model to improve on natural language understanding (NLU) tasks, there exist several limitations: (1) images cannot faithfully convey word meanings that require more activity-based and physical commonsense knowledge. (2) the voken supervision suffers from approximation/quantization error of the text-to-image retrieval.

To address these problems, we propose a novel Video-and-Language Knowledge Distillation method, named VIDLANKD. Our teacher model consists of a video encoder and a language encoder. They are jointly trained with a video-language contrastive learning objective and a masked language modeling (MLM) objective on a multi-modal dataset (see Fig. 1). Then, we transfer the knowledge of the frozen teacher language encoder to a student language model by minimizing the distance between contextualized text representations of two models on a text dataset. For this, we propose to use different KD objectives including neuron selectivity transfer (NST) [34] and contrastive representation distillation (CRD) [71] that avoid the approximation error from voken assignments [68]. For cross-modal pretraining of our teacher model, we use HowTo100M [54], a large-scale video dataset which has more diverse vocabulary and richer world commonsense (e.g., physical and temporal) knowledge compared to MS COCO image dataset.

In our experiments, student language models learned with the proposed video-language KD objectives outperform the baseline text-pretrained language models and the models distilled with vokenization, on several diverse natural language understanding benchmarks including GLUE [73], SQuAD [61], and SWAG [79]. We also show comprehensive ablation studies on video encoders, student KD objectives, teacher pretraining objectives, and video vs. image-based pretraining. Furthermore, we empirically illustrate that our model successfully learns linguistic world knowledge and physical/temporal commonsense abilities from video, by showing improved performances on the GLUE-diagnostics [73], PIQA [6], and TRACIE [82] datasets.

Overall, our contributions are: (1) a novel cross-modal knowledge distillation method for improving natural language understanding, (2) using rich video-text data which can overcome the limitations of image vokenization, (3) empirical improvements on several language understanding benchmarks and studying different knowledge distillation methods, and (4) analysis on linguistic/physical/temporal knowledge learned from videos and ablation studies on the effectiveness of proposed components.

2 Related Work

2.1 Knowledge Distillation

Knowledge distillation (KD) [33] is the process of transferring knowledge from a teacher model to a student model. It has been successfully used in a wide range of applications, such as machine translation [40], visual recognition [31], speech recognition [10], and recommendation systems [69]. Recent works advanced the field of knowledge distillation by proposing new architectures [77; 80; 1; 55] and objectives [34; 14].

While many KD works study the problem of knowledge transfer within the same modality, cross-modal knowledge distillation [27; 20; 71] tackles the knowledge transfer across different modalities. Gupta et al. [27] transfers the knowledge of a model trained with RGB images to another model for depth maps and optical flow. Do et al. [20] proposes a KD method for visual question answering

[2], where the trilinear (image-question-answer) relational representation of a teacher model is transferred to a bilinear (image-question) student model. Tian et al. [71] combines contrastive learning and knowledge distillation to improve the knowledge transfer between different modalities. Our VIDLANKD transfers the knowledge of a multi-modal teacher model learned from a video dataset to a student language model that tackles natural language understanding tasks.

2.2 Language Pretraining

Large-scale pretraining of contextualized language models has seen huge success in natural language processing in recent years. ELMo [57] proposes to pretrain and fine-tune a large recurrent language model, which improves performance on a diverse set of downstream natural language processing tasks. BERT [19] improves the scalability of the pretrain/fine-tune framework by using a transformer [72] language model with a masked language modeling objective. Since then, pretraining of transformer language models has been extensively explored [49; 78; 44; 22; 64; 59; 16] for various natural language understanding [61; 74; 79; 73] and generation tasks [24; 62; 61; 63].

2.3 Multi-modal Pretraining

Following the success of language pretraining with transformer models, pretraining of image-text [67; 51; 13; 47; 83; 45] and video-text [66; 54; 85; 53; 46; 70] multi-modal transformers have achieved improvements on numerous multi-modal downstream tasks [2; 76; 84]. The multi-modal transformers take both visual and textual inputs and are pretrained on image-text or video-text pairs with multi-modal masked language modeling objectives. Despite the success on multi-modal downstream tasks, Tan and Bansal [68] finds that the multi-modal pretraining does not improve (and sometimes even harms) the language understanding performance. This is because the scale and diversity of text vocabulary of image-text and video-text datasets are usually smaller than those of text datasets. To utilize the rich vocabulary of text dataset, our VIDLANKD transfers the knowledge of pretrained multi-modal model to a student language model with a large text dataset.

2.4 Visually-Grounded Language Learning

A series of works explore using visual information to aid language understanding and generation tasks including co-reference resolution [42; 15], machine translation [81], bilingual lexicon learning [39], and multi-modal contrastive learning [47]. Vokenization [68] proposes the visually-supervised language model, which is closest to our work. Vokenization proposes to supervise a language model to predict a visualized token, called ‘voken’ for each input text token. Vokens are obtained by a contextualized token-to-image matching model, pretrained on a MS COCO image captioning dataset [12]. In this work, we experiment with alternative objectives which avoid the approximation error from finite voken assignments. In addition, we use HowTo100M [54] video dataset, which provides a more diverse vocabulary as well as richer world commonsense and physical reasoning knowledge.

3 Video-Language Knowledge Distillation

3.1 Method Overview

We aim to learn a better language representation with the knowledge distilled from visual information. For this, we leverage two kinds of datasets: the aligned multi-modal dataset, $D_{VL}: \{(\mathbf{x}, \mathbf{v})\}$ (e.g., HowTo100M [54]); and the text dataset, $D_L: \{\mathbf{x}\}$ (e.g., Wikipedia), where \mathbf{x} is a sentence and \mathbf{v} is a video paired with \mathbf{x} . Our knowledge transfer is done in two stages: (1) cross-modal pretraining of a teacher model, M^T , on multi-modal data D_{VL} (Eq. 1) (2) distilling the knowledge of teacher model to a student model, M^S , on text data D_L (Eq. 2). We illustrate our two-stage knowledge transfer method in Fig. 1.

$$\min_{\theta^T} \mathbb{E}_{\mathbf{x}, \mathbf{v} \sim D_{VL}} \mathcal{L}^T(M^T, \mathbf{x}, \mathbf{v}) \tag{1}$$

$$\min_{\theta^S} \mathbb{E}_{\mathbf{x} \sim D_L} \mathcal{L}^{KD}(M^T, M^S, \mathbf{x}) \tag{2}$$

Our teacher model M^T consists a language model LM^T and a visual encoder V^T . Both LM^T and V^T have transformer [72] architectures, where LM^T takes text tokens \mathbf{x} and V^T takes video frame




Figure 2: Cross-modal pretraining of our teacher model on a multi-modal dataset (Sec. 3.2). We train our teacher model with (a) video-language contrastive learning and (b) masked language modeling. For video-language contrastive learning, we only illustrate the negative text samples for brevity.

features \mathbf{v} as inputs. Our student model M^S is a transformer language model LM^S sharing the same architecture with teacher language model LM^T . As illustrated in Fig. 1(a), we first train teacher models LM^T and V^T with contrastive learning and masked language modeling. Then, we distill the knowledge of teacher models to student model LM^S as in Fig. 1(b). In the following subsections, we discuss the detailed training procedure of teacher (Sec. 3.2, Fig. 2) and student models (Sec. 3.3, Fig. 3).

3.2 Teacher Model

We train our teacher model on a multi-modal dataset with two objectives, i.e., video-language contrastive learning (Fig. 2(a)) and masked language modeling (Fig. 2(b)): $\mathcal{L}^T = \mathcal{L}_{CT} + \mathcal{L}_{MLM}$ ²

Architecture. As shown in Figure 2, our teacher model M^T consists of a language encoder LM^T and a visual encoder V^T . Both LM^T and V^T have similar transformer architecture.³ For each sentence \mathbf{x} , we tokenize it and append a special token [CLS] that represents the entire sentence following Devlin et al. [19]. LM^T takes \mathbf{x} and outputs contextualized representation $\mathbf{h}^x = \{\mathbf{h}_{[CLS]}^x, \mathbf{h}_1^x \cdots \mathbf{h}_{|\mathbf{x}|}^x\}$. For each video \mathbf{v} , we extract frame-level features \mathbf{e}^v with an off-the-shelf image encoder (see more details in Sec. 4.2). Note that the parameters of the image encoder are not updated to save computation. We feed the frame features $\mathbf{e}^v = \{\mathbf{e}_1^v \cdots \mathbf{e}_{|\mathbf{v}|}^v\}$ to our visual encoder V^T to get contextualized video frame features $\mathbf{h}^v = \{\mathbf{h}_1^v \cdots \mathbf{h}_{|\mathbf{v}|}^v\}$. We get the final video representation $\bar{\mathbf{h}}^v$ by temporally averaging frame-level features: $\bar{\mathbf{h}}^v = \frac{1}{|\mathbf{v}|} \sum_{i=1}^{|\mathbf{v}|} \mathbf{h}_i^v$. Different from Tan and Bansal [68], both LM^T and V^T parameters are trained from scratch.

Video-Language Contrastive Learning. To learn multi-modal grounding, we use a contrastive learning objective that encourages a closer distance between representations of aligned video-text pairs than unaligned pairs, as shown in Fig. 2 (a). For each \mathbf{x} , we randomly sample another text \mathbf{x}' from its batch with $\mathbf{x}' \neq \mathbf{x}$. Similarly, for each \mathbf{v} , we randomly sample another video \mathbf{v}' from its

²In our experiments, different weights over the objectives did not significantly change the results.

³In our experiments, we use the BERT architecture with two different configurations: 12 layers/768 hidden dimensions ($BERT_{12L/768H} = BERT_{BASE}$) and 6 layers/512 hidden dimensions ($BERT_{6L/512H}$).

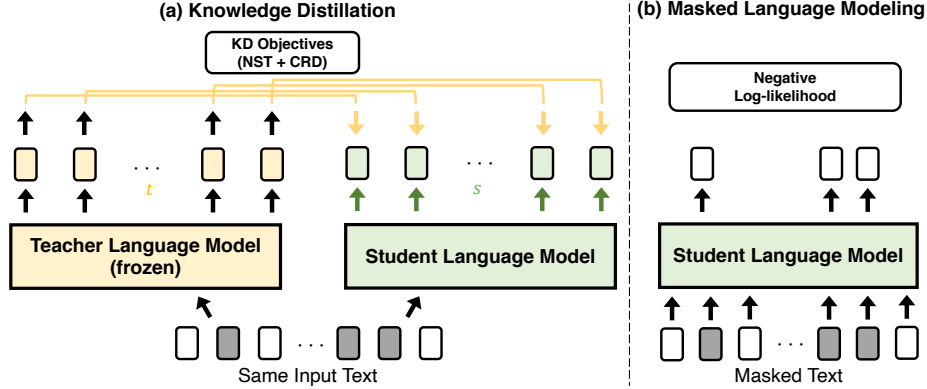


Figure 3: Illustration of our knowledge distillation from teacher language model LM^T to student language model LM^S on a text dataset (Sec. 3.3). We train our student model with (a) knowledge distillation objectives and (b) masked language modeling.

batch with $\mathbf{v}' \neq \mathbf{v}$. Then, we use hinge loss $\max\{0, \alpha - pos + neg\}$ on cosine similarities:

$$\mathcal{L}_{CT}(\mathbf{x}, \mathbf{x}', \mathbf{v}, \mathbf{v}') = \sum_i^{|\mathbf{x}|} [\max\{0, \alpha - \cos(\mathbf{h}_i^x, \overline{\mathbf{h}}^v) + \cos(\mathbf{h}_i^{x'}, \overline{\mathbf{h}}^v)\} + \max\{0, \alpha - \cos(\mathbf{h}_i^x, \overline{\mathbf{h}}^v) + \cos(\mathbf{h}_i^x, \overline{\mathbf{h}}^{v'})\}] \quad (3)$$

where α is the margin between the similarities of a positive pair and a negative pair. Different from previous methods [7; 32] that exploit sentence-level contrastive loss, we follow [68] to construct a token-level contrastive loss (triplet margin loss) that grounds the visual information to each contextualized token output. This fine-grained contrastive loss will help the token-level knowledge distillation in Sec. 3.4.

Masked Language Modeling. For better language understanding in our teacher model, we follow BERT [19] to use masked language modeling (MLM) objective (Fig. 2(b)). By replacing 15% of tokens in \mathbf{x} with a special token [MASK], we obtain a masked text $\mathbf{x}^{\text{masked}}$ with the same length. The model takes $\mathbf{x}^{\text{masked}}$ as input and learns to predict the tokens by minimizing the negative log-likelihoods: $\mathcal{L}_{\text{MLM}}(\mathbf{x}, \mathbf{x}^{\text{masked}}) = -\sum_{i \in \text{Mask}} \log p(\mathbf{x}_i | \mathbf{x}^{\text{masked}})$, where Mask refers to the indices of masked tokens.

3.3 Student Model

After we train a teacher model on a multi-modal dataset, we transfer its knowledge to a student model on a text dataset. Following Kim and Rush [40], we train our student model with a sum of masked language modeling and two knowledge distillation objectives, NST and CRD (see Sec. 3.4):

$$\mathcal{L}^S = \mathcal{L}_{\text{MLM}} + \mathcal{L}_{\text{NST}}^{\text{KD}} + \mathcal{L}_{\text{CRD}}^{\text{KD}} \quad (4)$$

Architecture. As shown in Fig. 3 (a), our student model M^S is a language model LM^S with the same transformer architecture as the teacher language model LM^T . We train LM^S from scratch. Following previous works [11; 26], we introduce a multi-layer perceptron (MLP) distillation head on top of the last hidden states of LM^S . In our ablation study in appendix, we find that adding a distillation head slightly improves the distillation performance.

3.4 Knowledge Distillation Objectives

We next describe the knowledge distillation (KD) objectives used in transferring knowledge from the teacher model LM^T (Sec. 3.2) to this student model LM^S . Note that the weights of the teacher model LM^T are frozen during the knowledge distillation process since the teacher model should not be affected by the student model’s performance. Following Kim and Rush [40], we use the knowledge

distillation objective combined with the MLM objective (Fig. 3 (b)). Concretely, we use the same input text mask for MLM and KD objectives. While we calculate the MLM loss only on masked positions, we calculate KD losses using all hidden states following Clark et al. [16] (Fig. 3).

We study the following KD objectives: Soft-label [33], L2 Regression [3], Neuron Selectivity Transfer (NST) [34], Contrastive Representation Distillation (CRD) [71], and Vokenization [68]. In our experiments comparing different KD objectives (Table 5), NST and CRD perform best, while the combination of them improved the performance even further. Therefore, we propose NST+CRD for our cross-modal knowledge distillation objective.

Soft Label: Hinton et al. [33] proposed a knowledge transfer method by taking a teacher model prediction with temperature scaling as a ‘soft label’. We minimize cross-entropy between $P^T(y|x)$ and $P^S(y|x)$, i.e., the word output probabilities of LM^T and LM^S given the input text \mathbf{x} respectively:

$$\mathcal{L}_{\text{soft-label}}^{\text{KD}}(\mathbf{x}) = - \sum_{i=1}^{|\mathbf{x}|} \sum_{z \in Z} P^T(y_i = z|\mathbf{x}) \log P^S(y_i = z|\mathbf{x}) \quad (5)$$

where Z is the word vocabulary. Following Hinton et al. [33], we divide the softmax logits of LM^T and LM^S by a temperature parameter $\tau = 2.0$. Note that for soft-label KD, we reuse the LM head, instead of learning an additional distillation head.

L2 Regression: Following Ba and Caruana [3] which uses feature regression for KD, we minimize the squared L2 distance between $\mathbf{s}(\mathbf{x})$ and $\mathbf{t}(\mathbf{x})$, the last hidden states of LM^T and LM^S given input text \mathbf{x} :

$$\mathcal{L}_{\text{Regression}}^{\text{KD}}(\mathbf{x}) = \sum_{i=1}^{|\mathbf{x}|} \|\mathbf{s}(\mathbf{x})_i - \mathbf{t}(\mathbf{x})_i\|_2^2 \quad (6)$$

Neuron Selectivity Transfer (NST): NST [34] is a KD method that transfers heatmap like spatial activation patterns of teacher neurons to student neurons. We transfer the sequential activation patterns of $\mathbf{t}(\mathbf{x}) \in \mathbb{R}^{|\mathbf{x}| \times d}$ to $\mathbf{s}(\mathbf{x}) \in \mathbb{R}^{|\mathbf{x}| \times d}$, where $\mathbf{t}(\mathbf{x})$ and $\mathbf{s}(\mathbf{x})$ are the last hidden states of LM^T and LM^S given input text \mathbf{x} , and d is the hidden state dimension (# neurons). Following Huang and Wang [34], we use the squared maximum mean discrepancy (MMD) [25] with kernel trick to measure the distance between the activation patterns of student neurons $\{\mathbf{s}(\mathbf{x})_{*,i}\}_{i=1}^d$ and teacher neurons $\{\mathbf{t}(\mathbf{x})_{*,j}\}_{j=1}^d$:

$$\begin{aligned} \text{MMD}^2(\mathbf{x}) &= \frac{1}{d^2} \sum_{i=1}^d \sum_{i'=1}^d k[\mathbf{s}(\mathbf{x})_{*,i}; \mathbf{s}(\mathbf{x})_{*,i'}] + \frac{1}{d^2} \sum_{j=1}^d \sum_{j'=1}^d k[\mathbf{t}(\mathbf{x})_{*,j}; \mathbf{t}(\mathbf{x})_{*,j'}] \\ &\quad - \frac{2}{d^2} \sum_{i=1}^d \sum_{j=1}^d k[\mathbf{s}(\mathbf{x})_{*,i}; \mathbf{t}(\mathbf{x})_{*,j}] \end{aligned} \quad (7)$$

where we use Gaussian kernel $k[\mathbf{s}; \mathbf{t}] = \exp\left(-\frac{\|\mathbf{s}-\mathbf{t}\|_2^2}{2\sigma^2}\right)$ with $\sigma = 1$. We transfer the teacher activation patterns to the student by minimizing squared MMD: $\mathcal{L}_{\text{NST}}^{\text{KD}}(\mathbf{x}) = \text{MMD}^2(\mathbf{x})$

Contrastive Representation Distillation (CRD): CRD [71] is a KD objective which maximizes the mutual information between the teacher and student representations with contrastive learning. Let’s denote $\mathbf{s} \in S$ and $\mathbf{t} \in T$ as student and teacher representations given \mathbf{x} . We are given 1 positive pair (drawn from the joint distribution) for every N (batch size) negative pairs (drawn from the product of marginals; independent randomly drawn inputs from T and S). Following [71], we maximize the lower bound of mutual information between \mathbf{s} and \mathbf{t} by minimizing the following term:

$$\begin{aligned} \mathcal{L}_{\text{CRD}}^{\text{KD}}(\mathbf{x}) &= -\mathbb{E}_{q(\mathbf{s}, \mathbf{t}|\text{positive})}[\log h(\mathbf{s}, \mathbf{t})] - N \cdot \mathbb{E}_{q(\mathbf{s}, \mathbf{t}|\text{negative})}[\log(1 - h(\mathbf{s}, \mathbf{t}))] \\ h(\mathbf{s}, \mathbf{t}) &= \frac{\exp(f_1(\mathbf{s})^\top f_2(\mathbf{t}))}{\exp(f_1(\mathbf{s})^\top f_2(\mathbf{t})) + \frac{N}{M}} \end{aligned} \quad (8)$$

where M is the cardinality of the dataset, f_1, f_2 are learned linear layers followed by $L2$ normalization, which map the student and teacher representations into a same feature space. Since a large N leads

to a tight mutual information lower bound, following [71], we implement a memory buffer that stores the latent features of each data sample computed from previous batches. Therefore, during training we can efficiently retrieve a large number of negative samples from the memory buffer. Note that since CRD is based on contrastive learning, it is the only KD objective where student and teacher language models can take different inputs.

Vokenization: Vokenization [68] could be viewed as a knowledge distillation method, where token-level text-to-image retrieval results (called ‘vokens’) of a multi-modal matching model are used as labels for a student language model. For the i -th input token \mathbf{x}_i , we calculate cosine similarity between the i -th teacher language model hidden state $\mathbf{t}(\mathbf{x})_i$ and a video feature \mathbf{v} . Out of 30K pre-selected videos, we select a video that maximizes cosine similarity and use it as the voken for \mathbf{x}_i . By denoting the voken of \mathbf{x}_i as \mathbf{voken}_i , we formulate our vokenization-based KD objective as:

$$\mathcal{L}_{\text{Voken}}^{\text{KD}}(\mathbf{x}) = - \sum_{i=1}^{|\mathbf{x}|} \log P_{\text{voken}}^S(y_i = \mathbf{voken}_i | \mathbf{x}) \quad (9)$$

where $P_{\text{voken}}^S(y|\mathbf{x})$ is the voken classification probabilities of LM^S given input text \mathbf{x} . We experiment with vokenization-based KD by retrieving vokens from images and videos (see Table 6 of Sec. 5.2). Note that vokenization suffers from approximation error; it’s hard to cover diverse textual concepts with 30K vokens. This motivates us to experiment with different ‘soft’ KD objectives described in this section (see Table 5 of Sec. 5.2).

4 Experimental Setup

4.1 Datasets

Video-Text Dataset. We use HowTo100M [54] for cross-modal pretraining of our teacher model (Sec. 3.2). HowTo100M has 1.22M videos totaling 136M video clips with total duration of 134,472 hours describing over 23K different visual tasks. There are 138M captions, 568M tokens with 633K distinct tokens.

Text Pretraining Dataset. To transfer the knowledge from our teacher language models to student language models (Sec. 3.3), we follow Tan and Bansal [68] to use English Wikipedia. For ablation studies (Sec. 5.2), we use Wiki103 [52], a widely used subset of English Wikipedia. There are 2.9B tokens and 120M sentences in English Wikipedia, and 111M tokens and 4.2M sentences in Wiki103.

Text Downstream Dataset. Following Tan and Bansal [68], we finetune our models on GLUE [73], SQuAD [61] 1.0 and SQuAD2.0 [60], and SWAG [79] to assess the pretraining performance. Since some smaller tasks in GLUE are reported as unstable in recent papers [21], we evaluate on the four largest datasets of GLUE: SST-2 [9], QNLI [61], QQP [35], and MNLI [74]. In addition, we also evaluate our models on the GLUE diagnostics [73], PIQA [6], and TRACIE [82] to measure its linguistic knowledge, physical reasoning, and temporal reasoning abilities.

4.2 Video Feature Representations

Following Miech et al. [54], we encode video features by concatenating features from a 2D frame-level image encoder and a 3D video encoder in channel dimension. Note that the parameters for 2D image encoder and 3D video encoder are not updated.

For the 2D image encoder, we sample video frames by 1fps (frame/second). The 2D image encoder outputs features for each frame individually. We experiment with ResNet-152 [30] pretrained on ImageNet-1K [18] and CLIP [58] image encoder (ViT-B/32 [23]). In contrast to conventional image encoders trained with image label classification, the CLIP image encoder is trained to match a corresponding natural language description by large-scale contrastive learning. We discuss if this natural language supervision can help our cross-modal KD in Sec. 5.1.

For the 3D video encoder, we use 3D-ResNeXt-152⁴ [75; 28; 36] trained from a combination of publicly available datasets: ActivityNet [8], Kinetics [37], UCF-101 [65], and HMDB-51 [43]. The 3D video encoder processes 24fps videos with 3D convolution and yields features at 1.5fps. Then we sub-sample the features to 1fps to match the frame rate of 2D image encoder.

⁴<https://github.com/kenshohara/3D-ResNets-PyTorch>

Table 1: Cross-modal knowledge distillation results of BERT_{12L/768H} student language model on 7 downstream NLU tasks. In the first block, we include the image-based vokenization (Img-Voken) and its text-only pretrained baseline performance from Tan and Bansal [68]. In the second block, we compare our cross-modal KD method (NST+CRD) to video-based vokenization (Vid-Voken) and a text-only pretrained baseline. †EM refers to ‘Exact Match’.

	SST-2	QNLI	QQP	MNLI	SQuAD v1.1	SQuAD v2.0	SWAG	Avg.
	Acc	Acc	Acc	Acc	EM†	EM	Acc	
BERT _{12L/768H} [68]	89.3	87.9	83.2	79.4	77.0	67.7	65.7	78.6
+ KD (Img-Voken) [68]	92.2	88.6	88.6	82.6	78.8	68.1	70.6	81.4
BERT _{12L/768H}	89.0	88.0	86.2	79.2	77.2	68.0	65.0	78.9
+ KD (Vid-Voken) w/ ResNet	93.4	89.2	88.7	83.0	78.9	68.7	70.0	81.7
+ KD (Vid-Voken) w/ CLIP	94.1	89.8	89.0	83.9	79.2	68.6	71.6	82.3
+ KD (NST+CRD) w/ ResNet	94.2	89.3	89.7	84.0	79.0	68.9	71.8	82.4
+ KD (NST+CRD) w/ CLIP	94.5	89.6	89.8	84.2	79.6	68.7	72.0	82.6

4.3 Implementation Details

For the student distillation head, we use a two-layer MLP with ReLU activation. For both student and teacher language models, following previous works [49; 17; 68], we truncate input text that is longer than 128 tokens. We truncate videos features that are longer than 512 frames. We use an AdamW [41] optimizer with learning rate $2e-4$ and weight decay [50] of 0.01. We reserve 10K samples of the HowTo100M dataset as validation data. We train the teacher model until it converges on validation data. For downstream tasks, we report the results on the validation sets. We train 3 epochs with a learning rate of $1e-4$ and a batch-size of 32 for all downstream tasks. We use hinge loss margin $\alpha = 1.0$ for \mathcal{L}_{CT} (Eq. 3). We implement our models with PyTorch 1.5 [56] and train them with Nvidia GeForce RTX 2080ti GPUs. For teacher pretraining, we use 4 GPUs for BERT_{12L/768H} and BERT_{6L/512H} models for 7 days and 2.5 days respectively. For knowledge distillation, we use 4 GPUs for BERT_{12L/768H} and BERT_{6L/512H} models for 10 days and 3 days respectively.

5 Results and Analysis

5.1 Primary Downstream Task Results

In the first block of Table 1, we include the image-based vokenization (Img-Voken) and their text-only pretrained baseline from Tan and Bansal [68].⁵ Given our reproduced text-only baseline shows a similar average performance (78.9 vs 78.6), our student models distilled from NST+CRD are much better (82.6 vs 81.4). We discuss the comparison between video-based and image-based KD in detail in the following ablation study in comparison to vokenization (Table 6).

In the second block of Table 1, we compare our proposed cross-modal KD method (NST+CRD) to video-based vokenization (Vid-Voken) and a non-KD baseline (BERT_{12L/768H}) which is only pretrained on text. We can see both cross-modal KD methods (i.e., KD and Vid-Voken) significantly outperform the text-only baseline across all 7 downstream tasks. We also experiment with different 2D frame encoders (Sec. 4.2): ResNet and CLIP. For both Vid-Voken and NST+CRD, we observe CLIP further improves the performance results over ResNet, indicating using a strong visual encoder helps the teacher training and thus benefits the knowledge distillation.

5.2 Ablation Studies

In this section, we conduct a comprehensive ablation study to show the effectiveness of our proposed methods. For all ablation experiments, we use BERT_{6L/512H} architecture for student and teacher language models. We use ResNet-152 for 2D frame encoder and 3D-ResNeXt-152 for 3D frame encoder (Sec. 4.2). Wiki103 [52] is used for student model training. We also perform ablation experiments on the effect of additional distillation head in appendix.

⁵Vokenization uses pretrained BERT checkpoint for its ‘teacher’ (vokenizer) model but we train our teacher language model fully from scratch.

Text-only Pretraining. Our cross-modal KD improves the performance on downstream NLU tasks significantly (Sec. 5.1). Where does the improvement come from, video or text? To answer this question, we conduct text-only pretraining of BERT_{6L/512H} on Wiki103 text (111M tokens), HowTo100M captions (568M tokens) and compare them to a no-pretrain baseline. In Table 2, while both pretrained models improve the performance over the no-pretrain baseline, Wiki103-trained model outperforms HowTo100M-trained model (which has more tokens) significantly. This indicates that our KD methods improve NLU performance because of multimodal grounding, instead of just the larger corpus.

Table 2: Text-only pretraining results of BERT_{6L/512H} pretrained on Wiki103, HowTo100M captions, and no-pretrain baseline.

Pretrained on	SST-2	QNLI	QQP	MNLI
No-Pretrain	79.6	61.5	72.7	61.6
Wiki103 (Formal language)	88.8	84.9	85.3	77.4
HowTo100M (ASR captions)	83.3	78.5	83.7	71.5

Effect of Teacher Training Objectives. We here analyze the teacher training objectives by comparing the corresponding distilled student model results. In Table 3, the teacher model trained solely with MLM (+KD from T^{MLM}) does not significantly change the student model performance. At the same time, the teacher model trained with only visual supervision, i.e., contrastive objective (+KD from T^{CT}), improves the result. This illustrates the motivation to perform knowledge transfer from a visually supervised MLM model. Lastly, combining the MLM and the contrastive objective (+KD from T^{MLM+CT}) in teacher model training shows the best student results.

Table 3: Ablation results showing the effect of the teacher model’s training objectives. NST is used for knowledge distillation.

	SST-2	QNLI	QQP	MNLI
BERT _{6L/512H}	88.8	84.9	85.3	77.4
+KD from T^{MLM}	88.1	83.1	85.6	77.4
+KD from T^{CT}	88.9	85.2	86.2	77.5
+KD from T^{MLM+CT}	91.1	85.0	87.4	78.4

Two-stage PT vs. Cross-modal KD. In Table 4, we compare two-stage pretraining with a single model to our proposed cross-modal KD approach. For single model baselines, we use text-only (MLM on Wiki103), video-only (MLM+CT on HowTo100M), and two-stage (video-then-text) pretraining. While the two-stage pretraining shows better results than the text/video-only pretraining, our VIDLANKD outperforms all baselines on GLUE tasks, especially on SST-2 and MNLI.

Table 4: Comparison of pretraining on text, video, both (Two-stage PT), and our VidLanKD.

Model	SST-2	QNLI	QQP	MNLI
Text PT	88.8	84.9	85.3	77.4
Video PT	84.0	78.9	84.2	73.1
Two-Stage PT	90.3	85.0	87.2	76.9
VIDLANKD	91.1	85.0	87.4	78.4

KD Objectives Comparison. In Table 5, we compare different knowledge distillation objectives introduced in Sec. 3.4. The student models trained with NST [34] and CRD [71] show the best finetuning performance on downstream tasks. When combining NST and CRD, performance further improves with marginal additional computation cost, hence we propose to use NST+CRD for our cross-modal knowledge distillation.

Table 5: Ablation of knowledge distillation objectives.

	SST-2	QNLI	QQP	MNLI
BERT _{6L/512H}	88.8	84.9	85.3	77.4
+KD-Soft label	87.2	84.4	86.4	76.6
+KD-Regression	88.8	84.8	87.1	78.1
+KD-Vid Voken	89.7	85.5	86.5	77.8
+KD-NST	91.1	85.0	87.4	78.4
+KD-CRD	90.0	85.5	87.3	78.3
+KD-NST+CRD	91.5	85.8	87.4	78.7

Comparison to Vokenization. In Table 6, we compare NST [34] and Vokenization [68] in both image and video-level teacher model supervision. For video-level supervision, we provide our visual encoder with the whole video features (Sec. 4.2). For image-level supervision, we provide our visual encoder only with 2D features of the middle frame for each video clip. With image-level super-

Table 6: Comparison between vokenization (Voken) and NST with image and video-level supervision.

	SST-2	QNLI	QQP	MNLI
BERT _{6L/512H}	88.8	84.9	85.3	77.4
+KD-Voken (Image)	89.3	84.4	86.0	77.5
+KD-NST (Image)	88.9	85.0	86.3	77.2
+KD-Voken (Video)	89.7	85.5	86.5	77.8
+KD-NST (Video)	91.1	85.0	87.4	78.4

vision (first block), Vokenization and NST show comparable performance. However, with video-level supervision (second block), NST outperforms Vokenization on 3 out of 4 tasks. The gap in the video domain might come from voken approximation error, where each image or video input is approximated with one of 30K predefined vokens. Since videos usually contain more diverse contents than images, the voken approximation error would be amplified in video-level supervision, whereas our NST distillation avoids this issue.

5.3 Analyzing the Knowledge Learned from Video

In this subsection, we analyze the knowledge that our language models learn from video via cross-modal knowledge distillation. To measure linguistic knowledge and physical/temporal reasoning ability, we show results of our models on the GLUE diagnostics [73], the Physical Interaction Question Answering (PIQA) [6], and TRACIE [82]. In addition, we visualize the learned multi-modal grounding ability of our model with text-to-video retrieval.

Table 7: Finetuning performance on GLUE diagnostics [73], PIQA [6] and TRACIE [82] datasets, which measure the linguistic knowledge, physical and temporal reasoning capabilities of language models, respectively.

	GLUE diagnostics				PIQA	TRACIE
	Lexicon	Predicate	Logic	Knowledge		
BERT _{6L/512H}	53.0	64.2	44.5	44.0	56.9	63.4
+ KD-NST	53.3 (+0.3)	63.7 (-0.5)	44.8 (+0.3)	48.6 (+4.6)	60.0 (+3.1)	66.7 (+3.3)

Linguistic Knowledge. GLUE diagnostics dataset [73] evaluates sentence understanding through natural language inference (NLI) problems. The dataset consists of sentence pairs labeled with their entailment relations (entailment, contradiction, or neutral) in both directions and tagged with a set of entailment labels. Each example in the dataset is labeled with 4 categories of linguistic phenomena: (1) lexical semantics, (2) predicate-argument structure, (3) logic, and (4) knowledge (including common sense). In Table 7, we compare the baseline language model (BERT_{6L/512H} pretrained on Wiki103) to our NST-distilled model. We finetune the models on MNLI [74] that has the same format and test on GLUE diagnostics. We observe a large gain on the knowledge category (which involves common sense and external world knowledge) while there are no significant differences on other categories. This suggests that our student model learns the external, grounded world knowledge in the teacher model and the video-text dataset.

Physical and Temporal Reasoning. PIQA [6] is a question answering dataset evaluating physical interactions and commonsense reasoning. TRACIE [82] is a temporal reasoning benchmark on implicit events, which are not mentioned explicitly in natural language text but can be inferred from it. In Table 7, our BERT_{6L/512H} distilled with NST significantly outperform the text-only pretrained baseline on both benchmarks. The finding suggests (consistent with the GLUE diagnostics findings above) that video knowledge distillation also helps improve the physical and temporal reasoning capabilities of the language model. See appendix for the more detailed discussion on the PIQA and TRACIE experiment.

6 Conclusion

We introduce VIDLANKD, a novel cross-modal knowledge distillation method to help general language understanding. Our teacher model is first trained on a video-text dataset, and then we transfer its knowledge to a student language model with a text dataset. Via the distillation objectives and video-text datasets, our method overcomes the limitations of the recent vokenization method. We empirically demonstrate that our VIDLANKD improves on several NLU tasks over models trained by pure-text or vokenization. We conduct comprehensive ablation analysis to show the effectiveness of each proposed component. We also illustrate the linguistic knowledge and physical/temporal commonsense reasoning learned from videos, and visualize our model’s multi-modal grounding ability.

Acknowledgments

We thank the reviewers for their helpful comments. We thank Yixin Nie and Gabriel Ilharco for useful dataset suggestions. This work was supported by ARO-YIP Award W911NF-18-1-0336, DARPA MCS Grant N66001-19-2-4031, DARPA KAIROS Grant FA8750-19-2-1004, Google Focused Research Award, and Bloomberg Data Science Ph.D. Fellowship. The views, opinions, and/or findings contained in this article are those of the authors and not of the funding agency.

References

- [1] Rohan Anil, Gabriel Pereyra, Alexandre Passos, Robert Ormandi, George E Dahl, and Geoffrey E Hinton. 2018. Large scale distributed neural network training through online distillation. *arXiv preprint arXiv:1804.03235*.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *ICCV*.
- [3] Lei Jimmy Ba and Rich Caruana. 2014. Do deep nets really need to be deep? In *NeurIPS*.
- [4] Emily M. Bender and Alexander Koller. 2020. **Climbing towards NLU: On meaning, form, and understanding in the age of data.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- [5] Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, et al. 2020. Experience grounds language. In *EMNLP*.
- [6] Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *AAAI*, pages 7432–7439.
- [7] Patrick Bordes, Eloi Zablocki, Laure Soulier, Benjamin Piwowarski, and Patrick Gallinari. 2019. Incorporating visual semantics into sentence representations within a grounded space. In *EMNLP-IJCNLP*.
- [8] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970.
- [9] Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. In *SemEval*.
- [10] Yevgen Chebotar and Austin Waters. 2016. Distilling knowledge from ensembles of neural networks for speech recognition. In *Interspeech*, pages 3439–3443.
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607. PMLR.
- [12] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- [13] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Learning universal image-text representations. In *ECCV*.
- [14] Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. 2018. Darkrank: Accelerating deep metric learning via cross sample similarities transfer. In *AAAI*.
- [15] Gordon Christie, Ankit Laddha, Aishwarya Agrawal, Stanislaw Antol, Yash Goyal, Kevin Kochersberger, and Dhruv Batra. 2016. Resolving language and vision ambiguities together: Joint segmentation & prepositional attachment resolution in captioned scenes. In *EMNLP*.
- [16] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In *ICLR*.
- [17] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *ACL*.
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.

- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- [20] Tuong Do, Thanh-Toan Do, Huy Tran, Erman Tjiputra, and Quang D Tran. 2019. Compact trilinear interaction for visual question answering. In *ICCV*, pages 392–401.
- [21] Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith. 2020. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *arXiv preprint arXiv:2002.06305*.
- [22] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *NeurIPS*.
- [23] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*.
- [24] David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.
- [25] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. **A kernel two-sample test**. *JMLR*, 13(25):723–773.
- [26] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. In *NeurIPS*.
- [27] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. 2016. Cross modal distillation for supervision transfer. In *CVPR*, pages 2827–2836.
- [28] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *CVPR*.
- [29] Stevan Harnad. 1990. **The symbol grounding problem**. *Physica D: Nonlinear Phenomena*, 42(1):335–346.
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778.
- [31] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. 2019. Bag of tricks for image classification with convolutional neural networks. In *CVPR*, pages 558–567.
- [32] Jack Hessel, Lillian Lee, and David Mimno. 2019. Unsupervised discovery of multimodal links in multi-image, multi-sentence documents. In *EMNLP-IJCNLP*, pages 2034–2045.
- [33] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2014. Distilling the knowledge in a neural network. In *NeurIPS Deep Learning Workshop*.
- [34] Zehao Huang and Naiyan Wang. 2017. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*.
- [35] Shankar Iyer, Nikhil Dandekar, and Kornel Csernai. 2017. **First quora dataset release: Question pairs**.
- [36] Hirokatsu Kataoka, Tenga Wakamiya, Kensho Hara, and Yutaka Satoh. 2020. Would mega-scale datasets further enhance spatiotemporal 3d cnns? *arXiv preprint arXiv:2004.04968*.
- [37] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- [38] Douwe Kiela, Alexis Conneau, Allan Jabri, and Maximilian Nickel. 2018. Learning visually grounded sentence representations. In *NAACL*.
- [39] Douwe Kiela, Ivan Vulic, and Stephen Clark. 2015. Visual bilingual lexicon induction with transferred convnet features. In *EMNLP*. ACL; East Stroudsburg, PA.
- [40] Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. In *ACL*.
- [41] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *ICLR*.

- [42] Chen Kong, Dahua Lin, Mohit Bansal, Raquel Urtasun, and Sanja Fidler. 2014. What are you talking about? text-to-image coreference. In *CVPR*, pages 3558–3565.
- [43] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. 2011. Hmdb: a large video database for human motion recognition. In *ICCV*, pages 2556–2563. IEEE.
- [44] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. In *ICLR*.
- [45] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, Daxin Jiang, and Ming Zhou. 2020. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*, pages 11336–11344.
- [46] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. 2020. Hero: Hierarchical encoder for video+ language omni-representation pre-training. In *EMNLP*.
- [47] Wei Li, Can Gao, Guocheng Niu, Xinyan Xiao, Hao Liu, Jiachen Liu, Hua Wu, and Haifeng Wang. 2020. Unimo: Towards unified-modal understanding and generation via cross-modal contrastive learning. *arXiv preprint arXiv:2012.15409*.
- [48] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer.
- [49] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- [50] Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *ICLR*.
- [51] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*.
- [52] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. Pointer sentinel mixture models. In *ICLR*.
- [53] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. 2020. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*.
- [54] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, pages 2630–2640.
- [55] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. 2020. Improved knowledge distillation via teacher assistant. In *AAAI*.
- [56] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*.
- [57] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL*.
- [58] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*.
- [59] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*.
- [60] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *ACL*.
- [61] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *EMNLP*.
- [62] Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. **A neural attention model for abstractive sentence summarization**. In *EMNLP*.

- [63] Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *ACL*.
- [64] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. Mass: Masked sequence to sequence pre-training for language generation. In *ICML*.
- [65] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. Ucf101: A dataset of 101 human actions classes from videos in the wild. *CoRR*.
- [66] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. Videobert: A joint model for video and language representation learning. In *ICCV*.
- [67] Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP*.
- [68] Hao Tan and Mohit Bansal. 2020. Vokenization: improving language understanding with contextualized, visual-grounded supervision. In *EMNLP*.
- [69] Jiayi Tang and Ke Wang. 2018. Ranking distillation: Learning compact ranking models with high performance for recommender system. In *SIGKDD*, pages 2289–2298.
- [70] Zineng Tang, Jie Lei, and Mohit Bansal. 2021. Decembert: Learning from noisy instructional videos via dense captions and entropy minimization. In *NAACL-HLT*, pages 2415–2426.
- [71] Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2020. Contrastive representation distillation. In *ICLR*.
- [72] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.
- [73] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*.
- [74] Adina Williams, Nikita Nangia, and Samuel R Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL*.
- [75] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. 2018. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*.
- [76] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*.
- [77] Zheng Xu, Yen-Chang Hsu, and Jiawei Huang. 2018. Training shallow and thin networks for acceleration via knowledge distillation with conditional adversarial networks. In *ICLR Workshop*.
- [78] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*.
- [79] Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. In *EMNLP*.
- [80] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. 2018. Deep mutual learning. In *CVPR*, pages 4320–4328.
- [81] Zhuosheng Zhang, Kehai Chen, Rui Wang, Masao Utiyama, Eiichiro Sumita, Zuchao Li, and Hai Zhao. 2019. Neural machine translation with universal visual representation. In *ICLR*.
- [82] Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal, and Dan Roth. 2021. **Temporal reasoning on implicit events from distant supervision**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1361–1371, Online. Association for Computational Linguistics.
- [83] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *AAAI*.
- [84] Luowei Zhou, Chenliang Xu, and Jason J Corso. 2017. Towards automatic learning of procedures from web instructional videos. In *AAAI*.
- [85] Linchao Zhu and Yi Yang. 2020. Actbert: Learning global-local video-text representations. In *CVPR*.