

Composable Augmentation Encoding for Video Representation Learning

Chen Sun^{1,2} Arsha Nagrani¹ Yonglong Tian³ Cordelia Schmid¹
¹ Google Research ² Brown University ³ MIT

{chensun, anagrani, cordelias}@google.com yonglong@mit.edu

Abstract

We focus on contrastive methods for self-supervised video representation learning. A common paradigm in contrastive learning is to construct positive pairs by sampling different data views for the same instance, with different data instances as negatives. These methods implicitly assume a set of representational invariances to the view selection mechanism (e.g., sampling frames with temporal shifts), which may lead to poor performance on downstream tasks which violate these invariances (fine-grained video action recognition that would benefit from temporal information). To overcome this limitation, we propose an ‘augmentation aware’ contrastive learning framework, where we explicitly provide a sequence of augmentation parameterisations (such as the values of the time shifts used to create data views) as composable augmentation encodings (CATE) to our model when projecting the video representations for contrastive learning. We show that representations learned by our method encode valuable information about specified spatial or temporal augmentation, and in doing so also achieve state-of-the-art performance on a number of video benchmarks.

1. Introduction

We focus on contrastive learning [16] of self-supervised video representations. The contrastive objective is simple: it pulls the latent representations of positive pairs close to each other, while pushing negative pairs apart. It is a natural fit for self-supervised learning, where positive and negative pairs can be constructed from the data itself, without the need for additional annotations.

Amongst different positive pair generation techniques, a particularly successful one has been augmentation-invariant contrastive learning [8, 20, 61], which has shown impressive results for image representation learning. In this instance discrimination framework, positive pairs are constructed by applying artificial aggressive *photo-geometric* data augmentations to create different versions of the same instances – learned representations are thus encouraged to

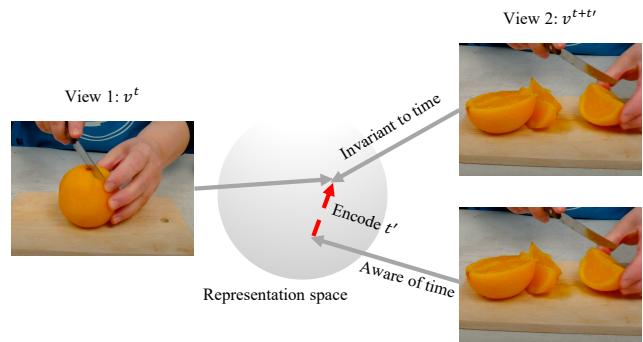


Figure 1. Standard self-supervised contrastive learning methods are *augmentation-invariant*, they ingest two augmented views of the same instance and encourage their latent representations to be similar. For videos, if the two views were sampled with a temporal shift, this approach would learn representations that are invariant to time changes, losing valuable temporal dynamics information (e.g. the action of slicing oranges) for downstream tasks. We instead project latent representations by *additionally encoding* the *relative transformation* of the two views (e.g. the time shift by t') with a composable augmentation encoding (CATE) to make the representation *augmentation-aware*.

be invariant to these data augmentations.

A number of self-supervised works also extend this idea to videos, where instead of artificial augmentations, they treat temporal shift as a *natural* data augmentation [46, 57], where two frames from the same video with a temporal shift form a positive pair. While this is useful for capturing high-level semantic information, e.g. knowledge of object categories, it can remove fine-grained information that may be useful depending on the downstream task – as previously studied in images [8, 63]. Consider the case of the oranges in Figure 1 – for object classification, invariance to small time shifts in a video that can result in view point changes or shape deformations may be useful and even desirable. However, for a downstream task involving reasoning about temporal relationships, such as recognising the state transition between a whole orange and orange slices caused by the act of cutting fruit, a representation invariant to temporal shifts may have lost valuable information.

We argue that augmentation-aware information can be

retained if the relative augmentations of the two views are known to the contrastive learning framework. For the oranges in Figure 1, an encoder could keep the shape information of the sliced orange view if it is aware that the other view is t' behind in time (and likely to be a whole orange).

We therefore propose a generalised framework for self-supervised video representation learning as follows. For notational convenience, we use the term *data augmentation* to consider all parameterisable data transformations, including shifts in space or in time. We then apply these generalised data augmentations to create different views of the same data, as is done in previous augmentation-invariant contrastive learning. However, instead of directly applying a contrastive loss on these views, we apply an additional projection head that optionally also *encodes the augmentations* that were used to create the views in the first place. For example, given two views of an image obtained via cropping, this can be an encoding of the bounding box coordinates of the cropping spatial transformation. For videos, this encoding can also include information about the temporal relationship between the views (*e.g.* a 5-second shift), in addition to spatial transformations. In the case of missing or occluded sequential data, we can also specify the particular locations to be predicted. When no such encodings are provided, our framework becomes standard augmentation-invariant contrastive learning.

We formulate this framework as a prediction task, given sequential data as input. The input sequence in this case contains encoded visual representations to be learned, and optionally a set of encoded data transformations for prediction. By using transformers, we can easily compose multiple encoded data transformations in the input sequence. We refer to this projection as a Composable AugmenTation Encoding (CATE) model. When data transformations are explicitly encoded in this way, our training objective can motivate a model to utilise such information if it helps with learning (*e.g.* learning the temporal dynamics that oranges can be cut into slices). We choose to always sample negative pairs across different instances, so the model has the freedom to ignore the transformation encodings if they do not help in reducing the contrastive loss. However, empirical results show that they are almost always utilised. We conduct thorough evaluations to test the efficacy of our framework, and in doing so make the following contributions: (1) we propose Composable AugmenTation Encoding (CATE) to learn augmentation-aware representations, and validate that CATE learns representations that preserve useful information (*e.g.* location, arrow of time) more effectively than a view-invariant baseline without augmentation encoding; (2) we perform a number of ablations on augmentation type and parameterisation, and observe that different downstream tasks favour the awareness of different augmentations, *e.g.* temporal awareness is particu-

larly helpful for fine-grained action recognition. We also find that encoding both the arrow of time and the absolute value of the temporal shift outperforms using just the arrow of time while encoding temporal information; (3) we set a new state-of-the-art for self-supervised learning on Something-Something [15], a dataset designed for fine-grained action recognition, and finally (4) we also achieve state-of-the-art performance on standard benchmarks such as HMDB51 [31] and UCF101 [52].

2. Related Works

Contrastive Learning. Recently the most competitive self-supervised representation learning methods have used contrastive learning [8, 9, 20, 21, 23, 42, 55, 61]. This idea dates back to [16], where contrastive learning was formulated as binary classification with margins. Modern contrastive approaches rely on a large number of negatives [20, 55, 61] and therefore a common technique is to employ the k-pair InfoNCE loss [42, 51]. By choosing different positive pairs or ‘view’ of the data, contrastive learning can encourage different representational invariances, *e.g.*, luminance and chrominance [55], rotation [39], image augmentations [8, 20], temporal shifts [17, 42, 47, 49, 69], text and its contexts [30, 33, 38], and multi-modal invariance [10, 41, 44, 48].

A number of works have highlighted the issues with these inbuilt representational invariances. InfoMin [56] demonstrates that different invariances favour different downstream tasks, and proposes that optimal views for a given task should only be invariant to irrelevant factors of that task. In [46], occlusion-invariance is shown to benefit the downstream task of object detection. However, designing task-dependent invariances requires knowing the downstream task beforehand and may make the learned representations less general. To overcome this, [63] learns multiple embedding spaces, each invariant to all but one augmentation. This means that the model (projection head) complexity grows linearly with the cardinality of the set of invariances. Instead, our transformer based projection head takes in a sequence of composable encodings, and can modulate the invariances that we want with a fixed model complexity.

Self-supervised Learning for Video. Videos offer additional opportunities for learning representations beyond those of images, by exploiting spatio-temporal and multi-modal [3, 43] information. Some works extend spatial tasks to the space-time dimensions of videos *e.g.* by using rotation [26] or jigsaw solving [29], while others use temporal information by ordering frames or clips [14, 25, 32, 34, 40, 60, 65], future prediction [17, 18, 58], speed prediction [5, 59], or motion [1, 12]. TaCo [4] consolidates these works by combining different temporal augmentations. Similar to us, they try to build augmentation awareness, but do so by adding a different *pretext task* head besides the projection

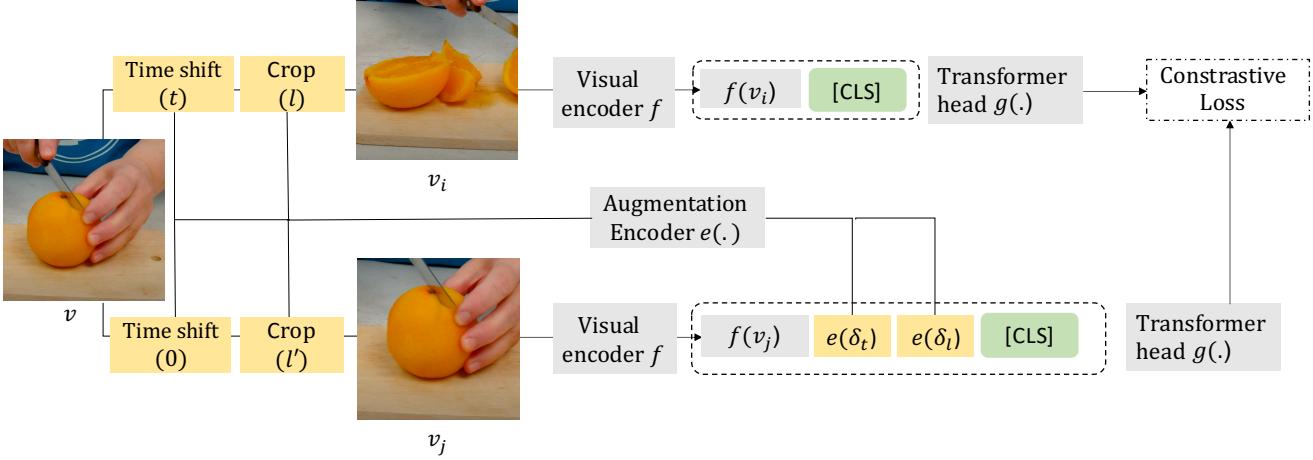


Figure 2. Overview of our contrastive learning framework CATE: Positive pairs are constructed from the same instance. For each view, a random set of data augmentations (e.g. temporal shifting, spatial cropping) is sampled and applied. The views are then encoded by a shared visual encoder (e.g. 3D ConvNets for videos). Encoded visual features, along with parameterised and embedded data augmentations, are then passed to a transformer head (this contains multiple layers, only the input layer is shown for simplicity) which summarises the input sequence and generates projected features for contrastive learning. In this example, the bottom transformer head is tasked to predict the features knowing the temporal augmentation (predict features t seconds ahead in time) and spatial augmentation (shift of box coordinates) relative to v_i . The visual encoder f is transferred to the downstream tasks.

head for each temporally transformed video, which once again grows linearly with the cardinal set of augmentations.

3. Method

This section describes our unified CATE framework for contrastive learning. We begin by providing an overview of contrastive learning. We then describe two paradigms in contrastive learning - ‘view-invariant’ and ‘predictive coding’, and show how our framework consolidates both using a transformer projection head. We also discuss how existing contrastive learning approaches can be viewed as special cases of our framework. An illustrative overview of our proposed framework is provided in Figure 2.

3.1. Contrastive Learning

Contrastive learning methods learn representations by maximising agreement between differently augmented views of the same data example (a positive pair) while pushing apart different data samples (negative pairs). The construction of views for a positive pair can be very general, *e.g.*, via a stochastic data augmentation module or by sampling co-occurring modalities from multi-sensory data. Formally, given two random variables \mathbf{x} and \mathbf{y} – contrastive learning seeks to learn a function that discriminates samples from the joint distribution $p(\mathbf{x}, \mathbf{y})$ and samples from the product of marginals $p(\mathbf{x})p(\mathbf{y})$. Therefore, there is a natural connection with mutual information maximisation, and as shown in [42] contrastive learning can be viewed as maximising a lower bound on the mutual information between representations of \mathbf{x} and \mathbf{y} . Specifically, given an

anchor point \mathbf{x} , its paired \mathbf{y}^+ , and a set of negatives \mathbf{y}_i^- ($i = 1, 2, \dots, K$), the InfoNCE loss is defined as follows:

$$\mathcal{L}_{\text{NCE}} = -\mathbb{E} \left[\log \frac{e^{h(\mathbf{x}, \mathbf{y}^+)}}{e^{h(\mathbf{x}, \mathbf{y}^+)} + \sum_{i=1}^K e^{h(\mathbf{x}, \mathbf{y}_i^-)}} \right] \quad (1)$$

The critical function $h(\cdot, \cdot)$ typically consists of one or more backbone networks [8, 20, 55], projection heads [8, 61], and a cosine similarity function. This function is optimised to assign a high score to positive pairs $(\mathbf{x}, \mathbf{y}^+)$ and a low score to negative pairs $(\mathbf{x}, \mathbf{y}_i^-)$. Minimising this InfoNCE loss is equivalent to maximising a lower bound I_{NCE} on the mutual information between \mathbf{x} and \mathbf{y} , denoted as $I(\mathbf{x}; \mathbf{y})$:

$$I(\mathbf{x}; \mathbf{y}) \geq \log(K) - \mathcal{L}_{\text{NCE}} = I_{\text{NCE}}(\mathbf{x}; \mathbf{y}) \quad (2)$$

3.2. Invariant and Predictive Coding

For visual representation learning, positive pairs can be constructed from unlabeled data in a self-supervised fashion. Following the notation of [8], a popular approach is instance discrimination: where two views $\mathbf{v}_i, \mathbf{v}_j$ of the same instance \mathbf{v} are generated by applying independently sampled random data augmentations:

$$\begin{aligned} \mathbf{v}_i &= a(\mathbf{v}; \tau_i) & \mathbf{v}_j &= a(\mathbf{v}; \tau_j) \\ \mathbf{x} &= g(f(\mathbf{v}_i)) & \mathbf{y}^+ &= g(f(\mathbf{v}_j)) \end{aligned}$$

where $a(\cdot)$ is the data augmentation operation parameterised by τ , $f(\cdot)$ is the visual encoder usually implemented using ConvNets whose outputs are used for transfer learning, and $g(\cdot)$ is a projection head implemented by a Multi-layer Perceptron (MLP). Note here that both $f(\cdot)$ and $g(\cdot)$

are shared across views. We refer to this approach henceforth as invariant coding.

Predictive coding, on the other hand provides an alternative approach, where certain regions within the same instance are masked out, such as different words in a sentence, different objects in an image or different frames in a video. This creates two different views - observed and ‘missing’ regions of the instance $\mathbf{v}_i, \mathbf{v}_j$:

$$\begin{aligned}\mathbf{v}_i &= m(\mathbf{v}; \tau_i) \\ \mathbf{v}_j &= m(\mathbf{v}; \tau_j)\end{aligned}$$

where $m(\cdot)$ is the masking function, parameterised by τ . Without loss of generality, we assume τ to be additive. This is a reasonable assumption for visual data which can be considered as sequential, and the indices (*e.g.* timestamp for video, pixel location for image) to be masked are additive. The representations to be contrasted are then computed as:

$$\begin{aligned}\mathbf{x} &= g(f(\mathbf{v}_i); \mathbf{0}) \\ \mathbf{y}^+ &= g(f(\mathbf{v}_j); \tau_j - \tau_i)\end{aligned}$$

Unlike with view-invariant coding, the projection head $g(\cdot)$ now also takes τ as input, *i.e.* it is tasked to predict the representation conditioned on τ (*e.g.* which could represent the time information used to generate the views – ‘five seconds ahead’), and $\mathbf{0}$ indicates the identity operation. In practice, $g(\cdot)$ can be easily implemented with a recurrent neural network as in [42] or a transformer with positional embeddings as in [11].

We note that the predictive projection head $g(\cdot)$ does not necessarily need to take τ into account, assuming $f(\cdot)$ already encodes information invariant of τ . However, arguably a much simpler task is to have $f(\cdot)$ encode the information useful for prediction and $g(\cdot)$ serve as the predictor (*e.g.* generating new words for text, or learning the temporal dynamics for videos.).

3.3. A Unified CATE Framework

Both $a(\cdot)$ and $m(\cdot)$ can be seen as atomic data operations that ‘transform’ the data based on τ , and multiple data operations can be chained together. We can therefore unify invariant and predictive coding using the following notation:

$$\begin{aligned}\mathbf{v}_i &= t(\mathbf{v}; \tau_i) & \mathbf{v}_j &= t(\mathbf{v}; \tau_j) \\ \mathbf{x} &= g(f(\mathbf{v}_i); \mathbf{0}) & \mathbf{y}^+ &= g(f(\mathbf{v}_j); e(\tau_j - \tau_i))\end{aligned}$$

where $\tau_i = [\tau_i^1, \tau_i^2, \dots]$ is a sequence of atomic data transformations, $t(\cdot)$ applies the sequence of data augmentations on the input data, and $e(\cdot)$ is an encoder that takes τ as input.

Given a sequence of atomic data operations, we can control which are view-invariant and which are predictive with

$e(\cdot)$: when we instruct $e(\cdot)$ to ignore certain type of operations, our method is ‘view-invariant’ to such operations, otherwise we claim that such operations are predictive.

We implement $g(\cdot)$ using a transformer. Our transformer model takes a set of inputs, among them encoded visual features $f(\mathbf{v}_i)$. We project the features to have the same size as the hidden size of the transformer with a linear layer. Additionally, each selected data operation is encoded into an embedding by an encoder $e_i(\tau_i)$ dedicated to the operation type. For example, for cropping, the inputs would be the differences between the coordinates of cropped boxes (for implementation details see Sec. 4.2). Finally, we have a special [CLS] token which ‘summarises’ the information from visual features and augmentation embeddings, and outputs a single embedding for each instance. The output embedding is then projected with a linear layer to the desired output size. A big advantage of our framework is that the transformer projection head can elegantly deal with variable length inputs, and hence multiple augmentation encodings can be composed with a fixed model capacity. We demonstrate this in the experiments (Sec. 4.3) that composing both crop and time encoding improves performance.

Note that other common contrastive based learning techniques can be expressed as special cases of our generalised framework: if $e(\cdot)$ is set to always return $\mathbf{0}$ (the identity), our model is akin to SimCLR [8] using a transformer-based projection head. When the inputs τ to $e(\cdot)$ correspond to the masked indices from a sequence, our formulation coincides with that of BERT [11] or GPT [6].

4. Experiments

We first describe the datasets and their experimental setup (Sec. 4.1), and then delve into the implementation details (Sec. 4.2) of our framework. We then investigate a number of model ablations to better understand the design choices of CATE, described in Sec. 4.3. To further analyse our framework, we also design and evaluate on two proxy task that require knowledge of time in video - predicting time shifts and early action classification (Sec. 4.4). Finally, we compare performance to the state of the art on regular action classification benchmarks.

4.1. Datasets

Something-something [15] v1: This is a video dataset focused on human object interactions. The dataset has 108,499 videos with 174 categories of fine-grained human-object interactions. The action categories are designed to be focused on temporal information - time is needed to distinguish between *picking up something* and *putting down something*, and hence it has been observed [64] that for this dataset, capturing subtle temporal changes is important for good performance. We use Something-Something v1 (SSv1) for our main ablation experiments.

Something-something [15] v2: SSv2 is built on top of SSv1 by expanding the dataset size to 220,847 videos. The dataset is then augmented by [35] with object bounding box annotations. We use SSv2 to compare with previously published results. For both SSv1 and SSv2, we adopt the linear evaluation protocol used by self-supervised learning on ImageNet, where both pretraining and evaluation are done on the same dataset. At pre-training, videos in the training split are used to learn the representations. At the evaluation stage, we first train a supervised linear classifier on top of the frozen representations in the training split, then report classification accuracy on the validation set. To compare with others, we report results on Something-Else, which defines a split of SSv2 for few shot classification.

Kinetics-400: This dataset consists of 240K 10-second clips from YouTube videos with action labels covering 400 classes. We follow the standard practice of training on the trimmed clips in the training split but ignoring their action labels. To evaluate representation learned on Kinetics-400, we follow standard practice and report results on the following two datasets:

HMDB51: HMDB51 [31] contains 6,766 video clips from 51 action classes. Evaluation is performed using average classification accuracy over three train/test splits from [24], each with 3,570 train and 1,530 test videos.

UCF101: UCF101 [52] contains 13K videos downloaded from YouTube spanning over 101 human action classes. Similar to HMDB51 and as is standardly done, evaluation is performed using average classification accuracy over three train/test splits. We pretrain on Kinetics and then evaluate on HMDB51 and UC101 in the following two ways:

- (i) **Standard action classification:** We report performance with both (a) linear evaluation on frozen features and (b) finetuning. This is to compare with the state-of-the-art.
- (ii) **Early action classification:** This is to further understand the value of CATE, and results are provided in Sec. 4.4. Here we predict high level actions that may be performed in the future, given noisy visual evidence at the current time stamp, similar to what has been explored previously for action detection [54]. For this task, we train on just the first frame of the video in UCF and HMDB, and use only this single frame at test time as well.

4.2. Implementation Details

Base Model: Our implementation is based on the SimCLR [8] code. Unless otherwise mentioned, we use a standard 3D ResNet-50 following the architecture of the ‘slow’ branch in the SlowFast Networks [13]. Global batch normalization is used during contrastive pre-training, and local batch normalization is used during transfer learning. We use the standard data augmentations used by SimCLR [8]: random cropping, color jittering and Gaussian blur. For spatial cropping, we found it beneficial to limit the range of

cropped area to (0.16, 0.81) of the original image area. For videos, we use time shifting as an additional atomic data operation. All the above spatial augmentations are applied consistently over time to avoid corrupting temporal continuity. We use a lightweight transformer encoder as the projection head, with a hidden size of 768 units, intermediate size of 3072 and number of attention heads 12. We use 4 transformer layers in total. We add a linear projection layer after the transformer with output size of 256. As investigated in Section 4.3, when used standalone with no encoded data augmentations, the transformer head performs comparably with the nonlinear projection head used by SimCLR. This allows us to focus on the impact of augmentation encoding.

Augmentation Encoding: In this work we encode two augmentations - spatial cropping and temporal shift, as we found these to be most effective empirically for the downstream tasks of action classification. As noted by SimCLR, invariance to augmentations like color jittering and gaussian blur is beneficial for classification, and hence we do not encode these. We note here that a method that could automatically select which augmentations to encode would be interesting, and we leave this for future work. Augmentations are encoded by $e(\cdot)$ as follows: for cropping we record the 4 scalar values representing the bounding box (x_1, y_1, h, w) of the crop, we then compute the relative distance of the cropped boxes between two views $\delta_{x,y}$ and project it to 768-dim with a linear layer. For temporal shifting we encode the binary indicator for arrow of time, and then a single scalar representing the number of frames shifted by. Each is projected to 768-dim with an embedding lookup table respectively, then summed together.

Pretraining: We feed 16 input frames to the ResNet-50-3D backbone at pre-training, with a frame sampling stride of 4 for Kinetics, and 2 for SSv1 and SSv2. All frames are cropped and resized to 224×224 . The transformer projection head is jointly trained with the ResNet-3D backbone during pre-training. We use the LARS optimiser with an initial learning rate of $4.8 (= 0.15 * \sqrt{(\text{BatchSize})})$, and weight decay of 10^{-4} . Unless otherwise mentioned, we pre-train for 500 epochs, with a batch size of 1024.

Evaluation: For linear evaluation, we freeze the pre-trained visual encoders $f(\cdot)$, extract the 2048-dim output features and train a linear classifier on top. The transformer projection head is not used at this stage. We sample 16 frames with a stride of 1 during training, and up to 8 sliding windows of 32 frames for evaluation, which covers the entire video span. This ‘multi-crop in time’ evaluation protocol is standard practice used by prior work [13]. For linear evaluation, we use a momentum optimiser of learning rate 0.16 and batch size of 256. For finetuning, we lower the learning rate to 0.02 and reduce the batch size to 128. All models are trained for 50 epochs.

Aug.	Time	Enc.	Projection	Top-1 Acc.	Top-5 Acc.
✗	✗		MLP	17.1	40.9
✓	✗		Linear	20.9	45.9
✓	✗		MLP	26.4	55.2
✓	✗		Transformer	26.5	55.9
✓	✗		Tr. + MLP	24.6	53.2
✓	✓		Transformer	31.2	61.4

Encoded	τ	Dropout	Top-1 Acc.	Top-5 Acc.
No	-	-	26.5	55.9
Crop	$\delta_{x,y}$	✗	27.2	56.7
Crop	$\delta_{x,y}$	✓	28.1	58.0
Time	$\text{sgn}(\delta_t)$	✗	28.1	57.9
Time	δ_t	✗	31.3	62.4
Time	δ_t	✓	31.2	61.4

Table 1: **Left:** Ablation showing the value of encoding time with different transformer projection heads and **Right:** Effect of regularisation on encodings. All results are on SSv1 with linear eval on frozen features. Top row: Vanilla simCLR Row four: SimCLR++, simCLR with time augmentation and transformer projection head. Last Row: When we encode time, we get a large boost in performance. On the right we show the impact of adding dropout to two aug. encodings - crop and time shifts. For crop encoding, dropout makes the pre-training task harder (higher contrastive loss at convergence), and also improves downstream accuracy. For time encoding, it is important to encode not only the arrow of time, but also the relative distance. However, dropout regularisation does not help in this setting.

4.3. Model Ablations

In this section we perform 5 ablations, all on the SSv1 dataset. We pretrain on SSv1 train set without labels and then train a single linear layer. We evaluate on the SSv1 validation set. We first ablate the two key design choices of CATE: (1) the use of a transformer projection head compared to an MLP or linear layer (Table 1, left), and (2) the way we parameterise and regularise augmentation encodings (Table 1, right). In particular, we assess the impact of encoding both time shifts and their direction (arrow of time) (Table 1, right). We then show (3) that it is possible to compose multiple augmentations in our framework, and finally we ablate some low level details such as (4) the number of layers in the transformer head and (5) the number of epochs used for training.

1. Different Projection Head Types. In this section, we start with the vanilla SimCLR model, and then vary the following (i) adding in temporal augmentations while training our model, (ii) encoding time or not under our CATE framework and (ii) whether we use a linear, MLP or transformer projection head. Results are shown in Table 1 (left).

SimCLR: Vanilla SimCLR [8] is applied to video frames, with no temporal data augmentation. Spatial augmentations are applied on the same frames to create views, and an MLP projection head is used. Results on SSv1 can be seen in the first row of Table 1 (left).

SimCLR++: In addition to spatial augmentation, we sample frames at different times in the same video to create views (row 3). We then also replace the MLP projection head with a transformer projection head which takes only the encoded visual representation $f(\cdot)$. This baseline is shown in the fourth row of Table 1 (left), highlighted in blue, and henceforth is referred to as SimCLR++. It is a strong baseline for us to compare to. From Table 1 (left), it is clear that temporal data augmentation is essential for video representation learning, leading to a 9% gain on top-1 accuracy with

an MLP projection head (rows 1 and 3). When no temporal augmentation is encoded, we observe that a nonlinear MLP projection head gives similar performance to the transformer projection head (rows 2 and 3), which validates that any performance improvements in CATE are not due solely to direct replacing the MLP head with a transformer. We also observe that combining the transformer head with the MLP head leads to slightly worse performance (row 5).

2. Encoding Augmentations. We first observe in Table 1 (left) that adding temporal encoding improves the top-1 accuracy by nearly 5% from 26.5% to 31.2% over SimCLR++ (last row vs fourth row in blue).

To further explore the efficacy of our augmentation encoding, in Table 1 (right), we ablate on the augmentation type using two augmentations - cropping and temporal shifts (time). We also explore the method of parameterisation for temporal augmentation (using just the arrow of time $\text{sgn}(\delta_t)$ or the distance in time δ_t , which includes both the absolute value of the temporal shift and its direction), and also the regularisation on the encodings (dropout or no dropout). We observe that both crop encoding and time encoding on their own improve the classification accuracy over the baseline, with time encoding providing a slightly bigger boost. By comparing the fourth row and the fifth row, we can see that the parameterisation of the temporal augmentation also matters, and it is beneficial to pass the distance in time along with the arrow of time to the augmentation encoder. Finally, we observe that dropout regularisation helps crop encoding but not time encoding. We hypothesise that crop encoding might make the contrastive task too easy and stronger regularisation is needed; time encoding does not suffer from this issue as there is more variation to be learnt from different frames in a video, and it actually benefits from a more informative encoding (from $\text{sgn}(\delta_t)$ to δ_t).

3. Composing multiple augmentations. In Table 2 we

show results for composing both crop and time encodings before feeding them to the transformer head. We can see for that for SSv1 (similar results hold for SSv2 and can be found in Table A4, appendix), using crop and time encodings individually leads to improved performance over the no-encoding baseline (*i.e.* SimCLR++, first row), and composing them together leads to further improvement.

Enc. Crop	Enc. Time	Top-1 Acc.	Top-5 Acc.
✗	✗	26.5	55.9
✓	✗	28.1	58.0
✗	✓	31.2	61.4
✓	✓	32.2	62.4

Table 2: **Composing spatial (crop) and temporal encodings** for Something-Something v1. Each individual encoding outperforms the no encoding baseline (SimCLR++). Composing them together yields the best performance.

4. Number of transformer layers. We experiment with a number of Transformer layers (1,2,4,8) for our projection head (with time encoding), and observe that the performance begins to saturate at four layers (Table A2, appendix). We use four layers in all other experiments.

5. Number of training epochs. We study the impact of the number of epochs used for pre-training on SSv1 and Kinetics-400. For evaluation we use SSv1, HMDB and UCF. Similar to SimCLR [8], we observe improved performance by increasing the number of epochs initially, which then saturates at around 500 epochs (Table A3, appendix).

4.4. Further Analysis

In this section we further analyse the effect of encoding augmentations on learned representations.

Per-class breakdown on SSv1: We conjecture that augmentation encoding is helpful for the downstream tasks that need to be aware of the corresponding spatial and temporal augmentations. To verify this conjecture, in Table 3 and Table A8 (appendix) we list the SSv1 classes that benefit the most and the least from time and crop encoding, respectively. We sort the classes by computing their per-class Average Precision. We can clearly see that the top classes for time encoding are typically sensitive of temporal ordering by definition, such as *lift up then drop down*, and *move closer*, where changing the arrow of time would lead to the opposite action (*e.g. move farther away*). Similarly, for crop encoding, the classes that benefit the most are those that require some level of spatial reasoning (*e.g. lift up, drop down, pull from right to left*, and *move down*), while the bottom classes typically do not require spatial reasoning.

Predicting Time Shifts: In previous experiments, we confirm empirically that encoding augmentations during pre-training leads to better downstream performance. As a san-

Label	ΔAP
Lifting something up completely, then letting it drop down	21.0
Pulling two ends of something so that it gets stretched	19.8
Moving something and something closer to each other	18.5
Taking one of many similar things on the table	17.2
Pushing something so that it almost falls off but doesn't	16.7
Poking something so lightly that it doesn't move	-4.6
Pretending to pour something out of something	-5.4
Poking a stack of something without the stack collapsing	-5.5
Pretending to spread air onto something	-7.8

Table 3: Classes that benefit the most and the least with **time encoding** on SSv1. We sort the classes by their differences on Average Precision.

ity check however, we also further design a proxy task to verify that the representations are indeed storing the encoded information and not discarding it. To analyse the time encodings, we design a time shift classification experiment based on the SSv1 dataset. For each video, we sample two 16-frame clips and use their relative distance in time as the classification label. The label space is *quantised* every 6 frames (0.5 seconds). During training and evaluation, we take the frozen representations of the two clips, concatenate them channel-wise and pass them to a linear classifier on top. Table 4 shows the results. We can see that by providing the encoded time augmentation δ_t during pre-training, CATE learns representations that maintain temporal shift information, solving the task with near perfect accuracy. Providing only the arrow of time $\text{sgn}(\delta_t)$ retains some information, while the no encoding baseline performs poorly on this probing task.

Encode Time	τ	Time Offset Acc.
✗	-	5.7
✓	$\text{sgn}(\delta_t)$	65.7
✓	δ_t	99.9

Table 4: **Time Shift Classification on SSv1.** Encoding time significantly helps on this proxy task, validating the intuition that our model retains useful time information.

Early Action Classification: We use this benchmark to investigate the impact of the number of observed frames on action recognition. Table A5 in Section 6.2 (Appendix) reports results on UCF-101 and HMDB-51. We observe that the gain of encoding time is bigger for early action classification than for full video classification, which indicates that time information is more important for the task.

4.5. Comparison with State-of-the-Art

Finally, we present comparisons with previous state-of-the-art methods on SS, UCF101 and HMDB51.

For SS, we compare our self-supervised representations with other weakly- and fully-supervised representations.

For evaluation, all representations are frozen, and a linear classifier is trained on the labeled training examples from the target dataset. In Table 5, we compare CATE with competitive weakly-supervised methods. CATE is pretrained on the train split of SSv1, and all weakly-supervised representations are pretrained by the authors of [66] on 19M public videos with hashtag supervision. The target dataset is SSv1. Despite training with only 0.1M videos without using their labels, our method is able to outperform these weakly-supervised approaches by large margins.

Method	Supervision	Top-1 Acc.
Distillation [22]	Weak	15.6
Prototype [50]	Weak	20.3
ClusterFit [66]	Weak	20.6
SimCLR++ [8]*	Self	26.4
CATE	Self	32.2

Table 5: **Comparison to SoTA on the Something-something v1 val set.** We use linear evaluation on frozen features. We compare to weakly-supervised baselines by [66]. *: re-implemented with temporal augmentation.

In Table 6, we compare with fully-supervised Spatial-Temporal Interaction Networks (STIN) [35] for few-shot action classification. Both CATE and STIN are pretrained on the ‘Base’ split of Something-Else, which contains half of the videos. STIN uses its labels as supervision while CATE does not. The target few-shot dataset contains 5, or 10 examples per class, across 86 classes. This is a more challenging setup than the 5-way classification setup used by [7]. We found CATE achieves on par or better performance than the supervised STIN.

Method	Pretrain	5-shot Acc.	10-shot Acc.
STIN+OIE+NL [35]	Supervised	17.7	20.7
SimCLR++ [8]*	Self-sup.	14.4	19.8
CATE	Self-sup.	18.0	22.9

Table 6: **Comparison to SoTA on Something-Else, a split of Something Something-v2 for few shot classification.** *: our re-implementation with temporal augmentation.

We also compare to the state-of-the-art on both HMDB51 and UCF101 in Table 7. Using frozen features, our model outperforms all other works that pretrain using RGB frames only – on UCF we even outperform a large number of works that use end-to-end finetuning. Additionally, our model outperforms AVSlowFast [62] which uses additional supervision from audio, and both MemDPC [18] and the recently proposed CoCLR [19], which use additional information from pre-extracted optical flow. Our model also compares favourably with MIL-NCE, XDC and ELO that are trained on orders of magnitude more training data – IG65M consists of 21 years of data (XDC), HTM,

Method	Modalities	Dataset	Frozen	UCF	HMDB
Shuff&Lrn* [40]	V	UCF	✓	26.5	12.6
3DRotNet [27]	V	K400	✓	47.7	24.8
CBT [53]	V	K600	✓	54.0	29.5
MemDPC [18]	V	K400	✓	54.1	30.5
TaCo [4]	V	K400	✓	59.6	26.7
CATE	V	K400	✓	84.3	53.6
MemDPC [18]	V+F	K400	✓	58.5	33.6
CoCLR [19]	V+F	K400	✓	74.5	46.1
AVSlowFast [62]	V+A	K400	✓	77.4	44.1
MIL-NCE [37]	V+T	HTM	✓	83.4	54.8
XDC [2]	V+A	IG65M	✓	85.3	56.0
ELO [45]	V+A	YT8M	✓	–	64.5
Shuff&Lrn* [40]	V	UCF	✗	50.2	18.1
CMC [55]	V	UCF	✗	59.1	26.7
OPN [32]	V	UCF	✗	59.6	23.8
ClipOrder [65]	V	UCF	✗	72.4	30.9
3DRotNet [27]	V	K400	✗	66.0	37.1
DPC [17]	V	K400	✗	75.7	35.7
CBT [53]	V	K600	✗	77.0	47.2
MemDPC [18]	V	K400	✗	78.1	41.2
SpeedNet [5]	V	K400	✗	81.1	48.8
VTHCL [67]	V	K400	✗	82.1	49.2
TaCo [4]	V	K400	✗	85.1	51.6
CATE	V	K400	✗	88.4	61.9
MemDPC [18]	V+F	K400	✗	86.1	54.5
CoCLR [19]	V+F	K400	✗	87.9	54.6
MIL-NCE [37]	V+T	HTM	✗	91.3	61.0
ELO [45]	V+A	YT8M	✗	93.8	67.4
XDC [2]	V+A	IG65M	✗	94.2	67.4

Table 7: **Comparison to the state of the art on UCF101 [52] and HMDB51 [31].** *reimplemented by [53]. **Frozen ✓** means the pretrained representation is fixed and classified with a linear layer, while **✗** means all layers are finetuned end-to-end. Rows highlighted in light blue use modalities beyond RGB frames as sources of supervision. Modalities are **V**: RGB frames only, **T**: text from ASR, **F**: pre-extracted optical flow, **A**: audio.

15 years (MIL-NCE), and YouTube 8M, 13 years (ELO). In contrast, Kinetics400 contains only 28 days of video data.

On finetuning we note that the gaps are smaller, however we still outperform all previously published works that use RGB frames only. We note that methods that use additional information from other modalities and train on orders more training data (MIL-NCE, XDC and ELO) are able to almost saturate performance on the UCF dataset.

5. Conclusion

We propose a general framework for contrastive learning, that allows us to build augmentation awareness in video representations. Our method consists of an elegant transformer head to encode augmentation information in a composable manner, and achieves state-of-the-art results for video representation learning. Future work will include evaluating on structured video understanding tasks and measuring the extent of equivariance learned by the representations.

References

- [1] Pulkit Agrawal, Joao Carreira, and Jitendra Malik. Learning to see by moving. In *CVPR*, 2015. 2
- [2] Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. *NeurIPS*, 2020. 8
- [3] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *ICCV*, 2017. 2
- [4] Yutong Bai, Haoqi Fan, Ishan Misra, Ganesh Venkatesh, Yongyi Lu, Yuyin Zhou, Qihang Yu, Vikas Chandra, and Alan Yuille. Can temporal information help with contrastive self-supervised learning? *arXiv preprint arXiv:2011.13046*, 2020. 2, 8
- [5] Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos. In *CVPR*, 2020. 2, 8, 13
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. In *NeurIPS*, 2020. 4
- [7] Kaidi Cao, Jingwei Ji, Zhangjie Cao, Chien-Yi Chang, and Juan Carlos Niebles. Few-shot video classification via temporal alignment. In *CVPR*, 2020. 8
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 1, 2, 3, 4, 5, 6, 7, 8, 12
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 2
- [10] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang. Perfect match: Improved cross-modal embeddings for audio-visual synchronisation. In *ICASSP*, 2019. 2
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 4
- [12] Ali Diba, Vivek Sharma, Luc Van Gool, and Rainer Stiefelhagen. Dynamonet: Dynamic action and motion network. In *ICCV*, 2019. 2
- [13] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *CVPR*, 2019. 5
- [14] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *CVPR*, 2017. 2
- [15] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *ICCV*, 2017. 2, 4, 5
- [16] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 1, 2
- [17] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding. *arXiv preprint arXiv:1909.04656*, 2019. 2, 8
- [18] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. *arXiv preprint arXiv:2008.01065*, 2020. 2, 8, 12, 13
- [19] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *NeurIPS*, 2020. 8
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 3
- [21] Olivier J Hénaff, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019. 2
- [22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 8
- [23] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR*, 2019. 2
- [24] Haroon Idrees, Amir R Zamir, Yu-Gang Jiang, Alex Gorban, Ivan Laptev, Rahul Sukthankar, and Mubarak Shah. The THUMOS challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding*, 2017. 5
- [25] Simon Jenni, Givi Meishvili, and Paolo Favaro. Video representation learning by recognizing temporal transformations. In *ECCV*, 2020. 2, 13
- [26] Longlong Jing and Yingli Tian. Self-supervised spatiotemporal feature learning by video geometric transformations. *arXiv preprint arXiv:1811.11387*, 2018. 2
- [27] Longlong Jing, Xiaodong Yang, Jingen Liu, and Yingli Tian. Self-supervised spatiotemporal feature learning via video rotation prediction. *arXiv preprint arXiv:1811.11387*, 2018. 8
- [28] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017. 13, 14
- [29] Dahyun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *AAAI*, 2019. 2
- [30] Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. A mutual information maximization perspective of language representation learning. In *ICLR*, 2020. 2

- [31] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: A large video database for human motion recognition. In *ICCV*, 2011. 2, 5, 8, 13
- [32] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Kumar Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *ICCV*, 2017. 2, 8, 12, 13
- [33] Lajanugen Logeswaran and Honglak Lee. An efficient framework for learning sentence representations. In *ICLR*, 2018. 2
- [34] Dezhao Luo, Chang Liu, Yu Zhou, Dongbao Yang, Can Ma, Qixiang Ye, and Weiping Wang. Video cloze procedure for self-supervised spatio-temporal learning. In *AAAI*, 2020. 2, 13
- [35] Joanna Materzynska, Tete Xiao, Roei Herzig, Huijuan Xu, Xiaolong Wang, and Trevor Darrell. Something-else: Compositional action recognition with spatial-temporal interaction networks. In *CVPR*, 2020. 5, 8
- [36] Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017. 13, 14
- [37] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, 2020. 8
- [38] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NeurIPS*, 2013. 2
- [39] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *CVPR*, 2020. 2
- [40] Ishan Misra, C. Lawrence Zitnick, and Martial Hebert. Shuffle and learn: Unsupervised learning using temporal order verification. In *ECCV*, 2016. 2, 8
- [41] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. *arXiv preprint arXiv:2004.12943*, 2020. 2
- [42] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 3, 4, 12
- [43] Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Ambient sound provides supervision for visual learning. In *ECCV*, 2016. 2
- [44] Mandela Patrick, Yuki M Asano, Ruth Fong, João F Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations. *arXiv preprint arXiv:2003.04298*, 2020. 2
- [45] AJ Piergiovanni, Anelia Angelova, and Michael S Ryoo. Evolving losses for unsupervised video representation learning. In *CVPR*, 2020. 8
- [46] Senthil Purushwalkam and Abhinav Gupta. Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. *arXiv preprint arXiv:2007.13916*, 2020. 1, 2
- [47] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. *arXiv preprint arXiv:2008.03800*, 2020. 2
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. 2
- [49] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, and Sergey Levine. Time-contrastive networks: Self-supervised learning from video. In *ICRA*, 2018. 2
- [50] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *NeurIPS*, 2017. 8
- [51] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NeurIPS*, 2016. 2
- [52] K. Soomro, A. Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. Technical Report CRCV-TR-12-01, 2012. 2, 5, 8, 13
- [53] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Contrastive bidirectional transformer for temporal representation learning. *arXiv preprint arXiv:1906.05743*, 2019. 8
- [54] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Rahul Sukthankar, Kevin Murphy, and Cordelia Schmid. Relational action forecasting. In *CVPR*, 2019. 5
- [55] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 2, 3, 8
- [56] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning. In *NeurIPS*, 2020. 2
- [57] Michael Tschannen, Josip Djolonga, Marvin Ritter, Aravindh Mahendran, Neil Houlsby, Sylvain Gelly, and Mario Lucic. Self-supervised learning of video-induced visual invariances. In *CVPR*, 2020. 1
- [58] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *CVPR*, 2016. 2
- [59] Jiangliu Wang, Jianbo Jiao, and Yun-Hui Liu. Self-supervised video representation learning by pace prediction. *arXiv preprint arXiv:2008.05861*, 2020. 2
- [60] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *CVPR*, 2018. 2
- [61] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 1, 2, 3
- [62] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual slowfast networks for video recognition. *arXiv preprint arXiv:2001.08740*, 2020. 8
- [63] Tete Xiao, Xiaolong Wang, Alexei A Efros, and Trevor Darrell. What should not be contrastive in contrastive learning. *arXiv preprint arXiv:2008.05659*, 2020. 1, 2
- [64] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning for video understanding. In *ECCV*, 2018. 4, 12

- [65] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yuetong Zhuang. Self-supervised spatiotemporal learning via video clip order prediction. In *CVPR*, 2019. [2](#), [8](#)
- [66] Xuetong Yan, Ishan Misra, Abhinav Gupta, Deepti Ghadiyaram, and Dhruv Mahajan. Clusterfit: Improving generalization of visual representations. In *CVPR*, 2020. [8](#)
- [67] Ceyuan Yang, Yinghao Xu, Bo Dai, and Bolei Zhou. Video representation learning with visual tempo consistency. *arXiv preprint arXiv:2006.15489*, 2020. [8](#)
- [68] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruyssen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, Lucas Beyer, Olivier Bachem, Michael Tschannen, Marcin Michalski, Olivier Bousquet, Sylvain Gelly, and Neil Houlsby. The visual task adaptation benchmark. *arXiv preprint arXiv:1910.04867*, 2019. [13](#)
- [69] Chengxu Zhuang, Alex Andonian, and Daniel Yamins. Unsupervised learning from video with deep neural embeddings. *arXiv preprint arXiv:1905.11954*, 2019. [2](#)