

Understanding Video Transformers via Universal Concept Discovery

Matthew Kowal^{1,3,4*} Achal Dave³ Rares Ambrus³

Adrien Gaidon³ Konstantinos G. Derpanis^{1,2,4} Pavel Tokmakov³

¹York University, ²Samsung AI Centre Toronto, ³Toyota Research Institute, ⁴Vector Institute

Project page: yorkucvil.github.io/VTCD

Abstract

This paper studies the problem of concept-based interpretability of transformer representations for videos. Concretely, we seek to explain the decision-making process of video transformers based on high-level, spatiotemporal concepts that are automatically discovered. Prior research on concept-based interpretability has concentrated solely on image-level tasks. Comparatively, video models deal with the added temporal dimension, increasing complexity and posing challenges in identifying dynamic concepts over time. In this work, we systematically address these challenges by introducing the first Video Transformer Concept Discovery (VTCD) algorithm. To this end, we propose an efficient approach for unsupervised identification of units of video transformer representations - concepts, and ranking their importance to the output of a model. The resulting concepts are highly interpretable, revealing spatiotemporal reasoning mechanisms and object-centric representations in unstructured video models. Performing this analysis jointly over a diverse set of supervised and self-supervised representations, we discover that some of these mechanisms are universal in video transformers. Finally, we show that VTCD can be used for fine-grained action recognition and video object segmentation.

1. Introduction

Understanding the hidden representations within neural networks is essential for addressing regulatory concerns [14, 34], preventing harms in deployment [8, 33], and can aid innovative model designs [16]. This problem has been studied extensively for images, both for convolutional neural networks (CNNs) [5, 25, 29, 37] and, more recently, vision transformers (ViTs) [55, 68], resulting in multiple key insights. For example, image classification models extract low-level positional and texture cues at early layers and gradually combine them into higher-level, semantic concepts at later layers [5, 27, 49].

*Work completed during an internship at Toyota Research Institute




Figure 1. Heatmap predictions of the TCOW model [66] for tracking through occlusions (top), together with concepts discovered by our VTCD (bottom). We can see that the model encodes positional information in early layers, identifies containers and collision events in mid-layers and tracks through occlusions in late layers. Only one video is shown, but the discovered concepts are shared between many dataset samples.

However, while video transformers do share their overall architecture with image-level ViTs, the insights obtained in existing works do very little to explain their inner mechanisms. Consider, for example, the recent approach for tracking occluded objects [66] shown in Figure 1 (top). To accurately reason about the trajectory of the invisible object inside the pot, texture or semantic cues alone would not suffice. What, then, are the *spatiotemporal* mechanisms used by this approach? Are any of these mechanisms *universal* across video models trained for different tasks?

To answer these questions, in this work we present the Video Transformer Concept Discovery algorithm (VTCD) - the first concept discovery methodology for interpreting the representations of deep video transformers. We focus on concept-based interpretability [24, 25, 29, 74] due to its

capacity to explain the decision-making process of a complex model’s distributed representations in high-level, intuitive terms. Our goal is to decompose a representation at any given layer into human-interpretable ‘concepts’ without any labelled data (*i.e.* concept discovery) and then rank them in terms of their importance to the model output.

Concretely, we first group *model features* at a given layer into spatiotemporal tubelets via SLIC clustering [1], which serve as a basis for our analysis (Section 3.1.1). Next, we cluster these tubelets across videos to discover high-level concepts [17, 24, 25, 42, 74] (Section 3.1.2). The resulting concepts for an occluded object tracking method [66] are shown in Figure 1 (bottom) and span a broad range of cues, including spatiotemporal ones that detect events, like collisions, or track the containers.

To better understand the decision-making mechanisms of video transformers, we then quantify the importance of concepts for the model’s predictions. Inspired by previous work on saliency maps [53], we propose a novel, noise-robust approach to estimate concept importance (Section 3.2). Unlike existing techniques that rely on gradients [37], or concept occlusion [24], our approach effectively handles redundancy in self-attention heads in transformer architectures.

Next, we use VTCD to study whether there are any universal mechanisms in video transformer models, that emerge irrespective of their training objective. To this end, we extend recent work [19] to automatically identify *important* concepts that are shared between several models in Section 4.1. We then analyze a diverse set of representations (*e.g.* supervised, self-supervised, or video-language) and make a number of discoveries: (i) many concepts are indeed shared between models trained for different tasks; (ii) early layers tend to form a spatiotemporal basis that underlines the rest of the information processing; (iii) later layers form object-centric video representations, even in models trained in a self-supervised way.

We also show how VTCD can be applied for downstream tasks. Firstly, pruning the heads of an action classification model according to their estimated importance yields a 4.3% *increase* in accuracy while reducing computation by 33%. Secondly, object-centric concepts discovered by VTCD can be used for video-object segmentation (VOS) and achieve strong performance on the DAVIS’16 benchmark [52] even for self-supervised representations.

2. Related work

Our work introduces a novel *concept-based interpretability* algorithm that focuses on *transformer-based representations* for *video understanding*. Below, we review the most relevant works in each of these fields.

Concept-based interpretability is a family of neural network interpretability methods used to understand, post-hoc, the representations that a model utilizes for a given task.

Closed-world interpretability operates under the premise of having a labeled dataset of concepts [5, 37]. However, for videos, it is unclear what concepts may exist and also difficult to densely label videos even if they were known a priori.

In contrast, unsupervised concept discovery makes no assumptions on the existence of semantic concepts and uses clustering to partition data into interpretable components within the model’s feature space. ACE [29] and CRAFT [25] segment input images into superpixels and random crops, respectively, before applying clustering at a given layer. In videos, however, the potential tubelets far outnumber image crops, prompting us to introduce a more efficient method for segmenting videos into proposals in Section 3.1.1.

A necessary component of concept-based interpretability is measuring the importance (*i.e.* fidelity) of the discovered concepts to the model. However, the aforementioned methods [24, 25, 29, 37, 74] were developed for CNNs, and are not readily applicable to transformers. The main challenge of ranking concepts in attention heads is due to the transformers’ robustness to minor perturbations in self-attention layers. To address this limitation, we introduce a new algorithm to rank the significance of any architectural unit, covering both heads and intra-head concepts in Section 3.2.

Recent work [19] identifies neurons that produce similar activation maps across various image models (including transformers). However, neurons are unable to explain the full extent of a models’ distributed [21] representation. In contrast, our method works on features of arbitrary dimension and is applied to video models.

Interpretability of transformers has received significant attention recently, due to its success in a variety of computer vision tasks. Early work [55] contrasted vision transformer representations with CNNs (representational differences per layer, receptive fields, localization of information, etc). Other work aims to generate saliency heatmaps based on attention maps of a model [11, 12]. Later works focused on understanding the impact of different training protocols [51, 68] (*e.g.* self-supervised learning (SSL) vs. supervised) and robustness [54, 75]. The features of a specific SSL vision transformer, DINO [2], were explored in detail and shown to have surprising utility for part-based segmentation tasks. However, none of these works address concept-based interpretability or study video representations.

Independently, studies in natural language processing (NLP) analyzed self-attention layers [20, 67] and found that heads are often specialized to capture different linguistic or grammatical phenomenon. This is qualitatively seen in vision works that show dissimilar attention maps for different self-attention heads [15, 36]. Moreover, other NLP works [48, 67] explore the impact of removing heads and find that only a small number need to be kept to produce similar performance. Our findings agree with evidence from

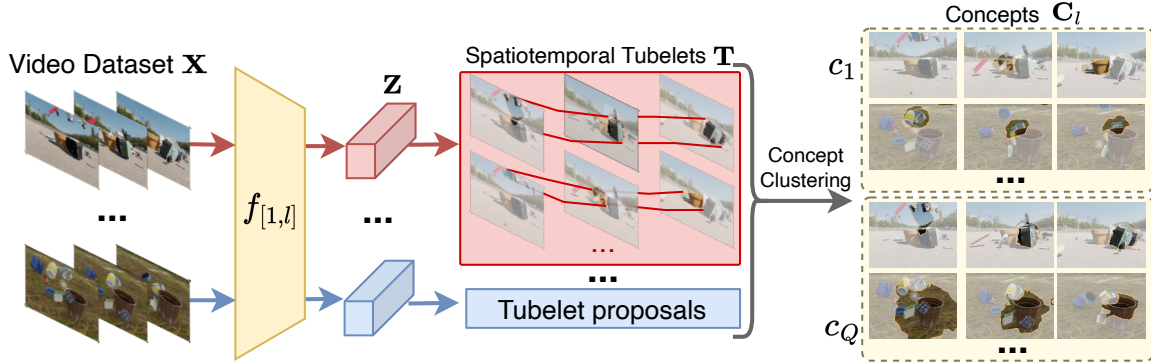


Figure 2. Video Transformer Concept Discovery (VTCD) takes a dataset of videos, \mathbf{X} , as input and passes them to a model, $f_{[1,l]}$ (shown in yellow). The set of video features, \mathbf{Z} , are then parsed into spatiotemporal tubelet proposals, \mathbf{T} (shown in red), via SLIC clustering in the feature space. Finally, tubelets are clustered across the videos to discover high-level units of network representation - concepts, \mathbf{C} (right).

these works and in Section 5 we further demonstrate that targeted pruning of unimportant heads from video transformers can actually *improve* a model’s performance.

Video model interpretability is an under-explored area of research considering the recent successes of deep learning models in action recognition [39], video object segmentation [43, 52, 63, 66], or self-supervised approaches [23, 56, 62, 64, 69]. Efforts have used proxy tasks to measure the degree to which models use dynamic information [28, 32, 35] or scene bias [13, 44, 45]. One method quantifies the static and dynamic information contained in a video model’s intermediate representation [40, 41]. However, these methods can only measure one or two predefined concepts (*i.e.* static, dynamic, or scene information) while our approach is not restricted to a subset of concepts. Another work visualizes videos that activate single neurons (or filters) in 3D CNN’s via activation maximization with temporal regularization [22]. While this method has no restrictions on what a neuron can encode, it only applies to 3D CNNs and does not truly capture ‘concepts’ across distributed representations (*i.e.* feature space directions that generalize across videos).

3. Video transformer concept discovery

We study the problem of decomposing a video representation into a set of high-level open-world concepts and ranking their importance for the model’s predictions. We are given a set of (RGB) videos, $\mathbf{X} \in \mathbb{R}^{N \times 3 \times T \times H \times W}$, where N , T , H , and W denote the dataset size, time, height, and width, respectively, and an L layer pretrained model, f . Let $f_{[r,l]}$ denote the model from layer r to l , with $f_{[1,l]}(\mathbf{X}) = \mathbf{Z}_l \in \mathbb{R}^{N \times C \times T' \times H' \times W'}$ being the intermediate representation at layer l . To decompose \mathbf{Z}_l into a set of human-interpretable concepts, $\mathbf{C}_l = \{c_1, \dots, c_Q\}$, existing, image-level approaches [25, 29] first parse the N feature maps into a set of M proposals, $\mathbf{T} \in \mathbb{R}^{M \times C}$ ($M > N$), where each T_m corresponds to a region of the input image. These proposals are then clustered into $Q \ll M$ concepts in the feature space of the model to form an assignment ma-

trix $W \in \mathbb{R}^{M \times Q}$. The importance of each concept c_i to the model’s prediction is quantified by a score $s_i \in [0, 1]$. Performing this analysis over all layers in f produces the entire set of concepts for a model, $\mathbf{C} = \{\mathbf{C}_1, \dots, \mathbf{C}_L\}$, together with their corresponding importance scores.

Existing approaches are not immediately applicable to video transformers because they do not scale well and are focused on 2D CNN architectures. In this work, we extend concept-based interpretability to video representations. To this end, we first describe a computationally tractable proposal generation method (Section 3.1.1) that operates over space-time feature volumes and outputs spatiotemporal tubelets. Next (Section 3.1.2), we adapt existing concept clustering techniques to video transformer representations. Finally, in Section 3.2 we propose a novel concept importance estimation approach applicable to any architecture units, including transformer heads.

3.1. Concept discovery

3.1.1 Tubelet proposals

Previous proposal methods [24, 25, 29] use superpixels or crops in RGB space to generate segments; however, the number of possible segments is exponentially greater for videos. Moreover, proposals in color space are unrestricted and may not align with the model’s encoded information, leading to many irrelevant or noisy segments. To address these drawbacks, we instantiate proposals in *feature space*, which naturally partitions a video based on the information contained within each layer (shown in Figure 2, left).

We produce tubelets per-video via Simple Linear Iterative Clustering [1] (SLIC) on the spatiotemporal features as

$$\mathbf{T} = \text{GAP}(\mathbf{B} \odot \mathbf{Z}) = \text{GAP}(\text{SLIC}(\mathbf{Z}) \odot \mathbf{Z}), \quad (1)$$

where $\mathbf{T} \in \mathbb{R}^{M \times C}$ is the set of tubelets for the dataset, $\mathbf{B} \in \{0, 1\}^{C \times M \times N \times T' \times H' \times W'}$ are spatiotemporal binary support masks obtained from SLIC, M is the number of tubelets for all N videos ($M \gg N$), and GAP represents global average pooling over the space and time dimensions.




Figure 3. A visual representation of concept masking for a single concept. Given a video \mathbf{x}_i and a concept, c_l , we mask the tokens of the intermediate representation $\mathbf{z}_i = f_{[1,l]}(\mathbf{x}_i)$ with the concepts’ binary support masks, \mathbf{B}_{c_l} , to obtain the perturbed prediction, \hat{y}_i .

SLIC is an extension of the K-Means algorithm [46] that controls a trade-off between cluster support regularity and adaptability, and also constrains cluster support masks to be connected. Together these properties produce non-disjoint tubelets that are easier to interpret for humans because they reduce the need to attend to multiple regions in a video at a time. Further, the pruning step in SLIC makes it more robust to the hyperparameter that controls the desired number of clusters, as it automatically prunes spurious, disconnected tubelets. Next, we describe our approach for grouping individual tubelets into higher-level concept clusters.

3.1.2 Concept clustering

Recent work [24, 25, 74] has used Non-Negative Matrix Factorization (NMF) [17] to cluster proposals into concepts. Given a non-negative data matrix, $\mathbf{T}^+ \in \mathbb{R}^{M \times C}$, NMF aims to find two non-negative matrices, $\mathbf{W}^+ \in \mathbb{R}^{M \times Q}$ and $\mathbf{C}^+ \in \mathbb{R}^{Q \times C}$, such that $\mathbf{T}^+ = \mathbf{W}^+ \mathbf{C}^+$, where \mathbf{W}^+ is the cluster assignment matrix. Unfortunately, NMF cannot be applied to transformers as they use GeLU non-linearities, rather than ReLU, resulting in negative activations.

We solve this problem by leveraging Convex Non-negative Matrix Factorization [17] (CNMF). Despite the name, CNMF extends NMF and allows for negative input values. This is achieved by constraining the factorization such that the columns of \mathbf{W} are convex combinations of the columns of \mathbf{T} , *i.e.* each column of \mathbf{W} is a weighted average of the columns of \mathbf{T} . This constraint can be written as

$$\mathbf{W} = \mathbf{T}\mathbf{G}, \quad (2)$$

where $\mathbf{G} \in [0, 1]^{C \times Q}$ and $\sum_j \mathbf{G}_{i,j} = 1$. To cluster a set of tubelets, \mathbf{T} , into corresponding concepts, we optimize

$$(\mathbf{G}^*, \mathbf{C}^*) = \arg \min_{\mathbf{C} > 0, \mathbf{G} > 0} \|\mathbf{T} - \mathbf{T}\mathbf{G}\mathbf{C}\|^2, \quad (3)$$

where the final set of concepts are the rows of the matrix \mathbf{C} , *i.e.* concept centroid c_i is the i^{th} row of \mathbf{C} (Figure 2, right).

3.2. Concept importance

Given a set of discovered concepts, we now aim to quantify their impact on model performance. One approach,

shown in Figure 3, is to mask out each concept independently and rank the importance based on the drop in performance [24]. Formally, let c_l be a single target concept, and $\mathbf{B}_{c_l} \in \{0, 1\}^{C \times M \times N \times T' \times H' \times W'}$ the corresponding binary support masks over \mathbf{X} . It can then be masked in layer l via

$$\hat{y} = f_{[l,L]}(\mathbf{Z}_l \odot (1 - \mathbf{B}_{c_l})). \quad (4)$$

While this approach works well for CNNs [24], transformers are robust to small perturbations within self-attention layers [48, 67]. Therefore, single concept masking has little effect on performance (shown by results in Figure 4). Instead, we mask a *high percentage* of sampled concepts in parallel (across all layers and heads) and then empirically validate in Section 5.1 that averaging the results over thousands of samples produces valid concept rankings.

Formally, we propose **Concept Randomized Importance Sampling (CRIS)**, a robust method to compute importance for any unit of interest. To this end, we first randomly sample K different concept sets, such that each $\mathbf{C}^k \subset \mathbf{C}$. We then define \mathbf{C}_l^k as the set of concepts in \mathbf{C}^k discovered at layer l , with $\mathbf{B}_{\mathbf{C}_l^k}$ denoting the corresponding binary support masks. We mask out every concept at every layer via

$$\hat{y}_k = g(\tilde{\mathbf{B}}_{\mathbf{C}_L^k} \odot f_{[L-1,L]}(\cdots (\tilde{\mathbf{B}}_{\mathbf{C}_1^k} \odot f_{[0,1]}(\mathbf{X}))), \quad (5)$$

where $g(\cdot)$ is the prediction head (*e.g.* an MLP) and $\tilde{\mathbf{B}}$ denotes the inverse mask (*i.e.* $1 - \mathbf{B}$). Finally, we calculate the importance of each concept, c_i , via

$$s_i = \frac{1}{K} \sum_k (\mathbb{D}(\tilde{y}, y) - \mathbb{D}(\hat{y}_k, y)) \mathbb{1}_{c_i \in \mathbf{C}^k}, \quad (6)$$

where \tilde{y} is the original prediction without any masking and \mathbb{D} is a metric quantifying performance (*e.g.* accuracy).

4. Understanding transformers with VTCD

Our algorithm facilitates the identification of concepts within any unit of a model and quantifying their significance in the final predictions; however, this is insufficient to fully represent the computations performed by a video transformer. It is also crucial to understand how these concepts are employed in the model’s information flow.

As several recent works have shown [20, 50], the residual stream of a transformer serves as the backbone of the information flow. Each self-attention block then reads information from the residual stream with a linear projection, performs self-attention operations to process it, and finally writes the results back into the residual stream. Crucially, self-attention processing is performed individually for each head with several studies showing, both in vision [2, 18, 36] and NLP [67], that different self-attention heads capture distinct information. In other word, heads form the basis of the transformer representation.

A closer analysis of the concepts found in the heads of the TCOW model [66] with VTCD allows us to identify several information processing patterns in that model. In particular, Figure 1 shows that the heads in early layers group input tokens based on their spatiotemporal positions. This information is then used to track objects and identify events in mid-layers, and later layers utilize mid-layer representations to reason about occlusions. Next, we study which of these mechanisms are *universal* across video transformers trained on different datasets and objectives (cf. [19]).

4.1. Rosetta concepts

Inspired by previous work [19], we mine for *Rosetta concepts* that are shared between models and represent the same information. The key to identifying Rosetta units is a robust metric, R , where a higher R -score corresponds to the two units having a larger amount of shared information. This previous work [19] focused on finding such neurons in image models based on correlating their activation maps. We instead measure the similarity between concepts (*i.e.* distributed representations) via the mean Intersection over Union (mIoU) of the concepts’ support.

Formally, we mine Rosetta concepts by first applying VTCD to a set of D models $\{f^1, \dots, f^D\}$, resulting in discovered concepts, $\mathbf{C}^j = \{c_1^j, \dots, c_i^j\}$, and importance scores, $\mathbf{S}^j = \{s_1^j, \dots, s_i^j\}$, for each model f^j . We then aim to measure the similarity between all concept D -tuples from the models. Given a set of D concepts, $\{c_i^1, \dots, c_i^D\}$, and corresponding binary support masks, $\{\mathbf{B}_i^1, \dots, \mathbf{B}_i^D\}$, we define the similarity score of these concepts as

$$R_i^D = \frac{|\mathbf{B}_i^1 \cap \dots \cap \mathbf{B}_i^D|}{|\mathbf{B}_i^1 \cup \dots \cup \mathbf{B}_i^D|}. \quad (7)$$

Naively computing the similarity between all D -tuples results in an exponential number of computations and is intractable for even small D . To mitigate these issues, we exclude two types of concepts: (i) unimportant ones and (ii) those with a low R -score among d -tuples, where $d < D$. More specifically, we only consider the most important $\epsilon\%$ of concepts from each model. We then iterate over $d \in \{2, \dots, D\}$ and filter out concepts that have R -scores less than δ for all d -tuples in which it participates. Formally, the filtered Rosetta d -concept scores are defined as

$$\mathbf{R}_{\epsilon, \delta}^d = \{R_i^d \mid R_i^d > \delta \forall R_i^d \in \mathbf{R}_\epsilon^d\}, \quad (8)$$

where \mathbf{R}_ϵ^d is the set of all R -scores among d concepts after the ϵ importance filtering. This results in a significantly smaller pool of candidates for the next stage, $d+1$, reducing the overall computational complexity of the algorithm. Finally, as some concepts may reside in a subset of the models but are still interesting to study, we examine the union of all important and confident Rosetta d -concepts corresponding to R -scores $\mathbf{R}_{\epsilon, \delta}^2 \cup \dots \cup \mathbf{R}_{\epsilon, \delta}^D$.

Model	TCOW		VideoMAE	
	Positive ↓	Negative ↑	Positive ↓	Negative ↑
Baseline + Occ	0.174	0.274	0.240	0.300
Baseline + CRIS	0.166	0.284	0.157	0.607
VTCD (Ours)	0.102	0.288	0.094	0.625

Table 1. Comparison of our tubelet proposal approach to the widely used random crop tubelets [24, 25, 72] with occlusion-based importance [24] (Baseline + Occ) for both TCOW [66] and VideoMAE [64]. Our tubelets result in concepts that are more faithful to the model’s representations even when the baseline is equipped with our concept scoring algorithm (Baseline + CRIS).

5. Experiments

We evaluate our concept discovery algorithm quantitatively and qualitatively across a variety of models and tasks.

Datasets. We primarily use two datasets in our experiments: TCOW Kubric [66] and Something-Something-v2 (SSv2) [30]. The former is a synthetic, photo-realistic dataset of 4,000 videos with randomized object location and motion, based on the Kubric synthetic video generator [31]. This dataset is intended for semi-supervised VOS (semi-VOS) through occlusions. SSv2 contains 220,847 real videos intended for fine-grained action recognition. Each sample is a crowdsourced video of a person-object interaction (*i.e.* doing something to something). Unlike other video classification benchmarks [9, 60], temporal reasoning is fundamental for distinguishing SSv2 actions, making it an ideal choice for analyzing spatiotemporal mechanisms.

Models. We evaluate four models with public pretrained checkpoints: (i) TCOW [66] trained on Kubric for semi-VOS, (ii) VideoMAE [64] trained on SSv2 for action classification (Supervised VideoMAE), (iii) VideoMAE self-supervised on SSv2 (SSL VideoMAE), and (iv) Intern-Video [70], a video-text model trained contrastively on 12M clips from eight video datasets and 100M image-text pairs from LAION-400M [58]. As TCOW requires a segmentation as input, when applying it to SSv2, we manually label the most salient object in the initial frame. We focus our analysis on the first two models, and use the last two to validate the universality of our Rosetta concepts.

Implementation details. For all experiments, we run VTCD on 30 randomly sampled videos and discover concepts from all heads and layers. We focus on keys in the self-attention heads, as they produce the most meaningful clusters per prior work [2]. We present results with queries and values on the project page, with further implementation details.

VTCD target metrics. To discover and rank concepts, VTCD requires a target evaluation metric. For TCOW Kubric, we use the intersection-over-union (IoU) between the predicted and the groundtruth masks. For SSv2, we use classification accuracy for a target class.

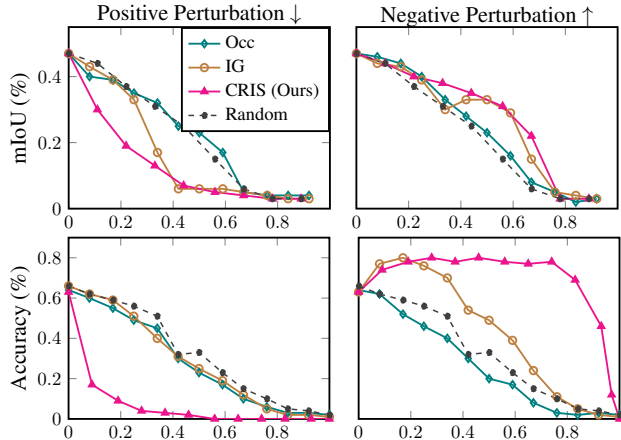


Figure 4. Attribution curves for every layer of TCOW trained on Kubric (top) and VideoMAE trained on SSv2 (bottom). We remove concepts from most-to-least (left) or least-to-most important (right). CRIS produces better concept importance than methods based on single concept occlusions (Occ) or gradients (IG).

5.1. Quantitative evaluation

This section presents a quantitative validation of VTCD’s key components compared to various baselines. We first assess our tubelet proposal methodology, followed by a comparison of CRIS with other concept importance methods. We follow the standard evaluation protocol, and measure the *fidelity* of the discovered concepts [25, 29, 74]. To this end, we calculate attribution curves [24, 25, 29], where concepts are removed from a model in either most-to-least order (*positive* perturbation), or least-to-most (*negative* perturbation). The intuition is that concepts, and corresponding importance scores, with higher fidelity will have a steeper performance decrease when removing the most important concepts, and vice-versa for the reverse order.

Tubelet validation. Unlike previous methods that partition the inputs into proposals in the pixel space, VTCD generates the proposals via SLIC [1] clustering in the model’s feature space. We ablate this design choice by comparing VTCD with a strong *random cropping* baseline used by recent image-based concept discovery methods (ConceptSHAP [72], CRAFT [25], Lens [24]), shown as ‘Baseline + Occ’ in Table 1. In all cases, VTCD yields concepts that are more faithful to the model’s representation. To further isolate the effect of proposals on the performance, we equip the baseline with our concept importance estimation approach (shown as ‘Baseline + CRIS’). We observe that VTCD still outperforms this strong baseline in all settings, validating that generating tubelet proposals in the feature space of the transformer indeed yields concepts that are more faithful to the model’s representation.

Concept important evaluation. In Figure 4, we plot concept attribution curves of our method for TCOW for Kubric (top) and supervised VideoMAE for SSv2, targeting 10 ran-

Model	Accuracy \uparrow	GFLOPs \downarrow
Baseline	37.1	180.5
VTCD 33% Pruned	41.4	121.5
VTCD 50% Pruned	37.8	91.1

Table 2. Pruning unimportant heads in VideoMAE results in improved efficiency and accuracy when targeting a subset of classes. Here, we target the six SSv2 classes containing types of spills.

domly sampled classes and averaging results (bottom). In addition, we report several baselines: (i) concept removal in a random order, (ii) standard, occlusion-based concept importance estimation [24], and (iii) a gradient based approach [24, 37]. In all cases, CRIS produces a more viable importance ranking, dramatically outperforming both random ordering and the occlusion baseline. The integrated gradients method performs similarly to ours for TCOW, but is much worse for the action classification VideoMAE.

Notably, we observe that the performance actually *increases* for VideoMAE when up to 70% of the least important concepts are removed. Recall that SSv2 is a fine-grained action classification dataset. VTCD removes concepts that are irrelevant for the given class, hence increasing the robustness of the model’s predictions. We further quantify this effect in Table 2, where we focus on the six classes where a ‘spill’ happens (listed in the appendix). We then use CRIS to rank the *heads* in VideoMAE according to their effect on performance using the training set and report the results for pruning the least important ones on the validation set. Pruning 33% of the heads actually *improves* the accuracy by 4.3% while reducing FLOPs from 180.5 to 121.5. Further removing 50% of the heads retains the original performance (+0.7%) and reduces FLOPs to 91.1. These results show that one can tune the trade-off between performance and computation using VTCD for pruning.

5.2. Qualitative analysis

We have seen that the importance assigned to concepts discovered by VTCD aligns with the accuracy of the model. We now assess the concepts qualitatively. To this end, Figure 5 shows two representative videos for the top three most important concepts for the TCOW and VideoMAE models for the class *dropping something into something*.

For TCOW, the most important concept occurs in layer five and tracks the target object. Interestingly, the same concept highlights objects with similar appearance and 2D position to the target. This suggests that the model solves the disambiguation problem by first identifying possible distractors in mid-layers (*i.e.* five) and then using this information to more accurately track the target in final layers. In fact, the second most important concept, occurring in layer nine, tracks the target object throughout the video.

For VideoMAE, the most important concept highlights

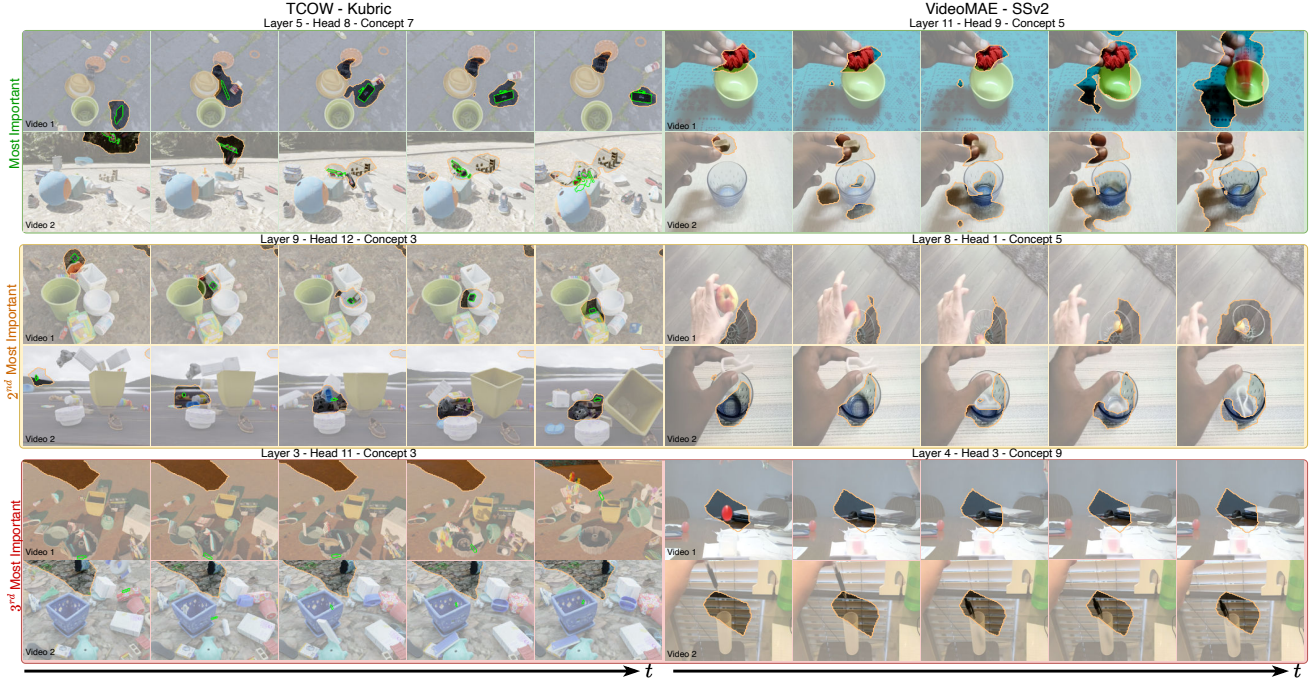


Figure 5. The top-3 most important concepts for the TCOW model trained on Kubric (left) and VideoMAE trained on SSv2 for the target class *dropping something into something* (right). Two videos are shown for each concept and the query object is denoted with a green border in Kubric. For TCOW, the 1st and 2nd (top-left, middle-left) most important concepts track multiple objects including the target and the distractors. For VideoMAE, the top concept (top-right) captures the object and dropping event (*i.e.* hand, object and container) while the 2nd most important concept (middle-right) captures solely the container. Interestingly, for both models and tasks, the third most important concept (bottom) is a temporally invariant tubelet. Section 5.3 provides a further discussion.




Figure 6. A sample Rosetta concept found in four models trained for different tasks. Interestingly, we find object-centric representations in all the models.

the object until it is dropped, after which the object and its container are highlighted. The second concept clearly captures the container, but notably not the object itself, making a ring-like shape. These concepts identify an important mechanism for differentiating similar classes (*e.g.* *dropping*

something into/behind/in-front of something).

The third concept for each model captures similar information, occurring in early layers: a temporally invariant, spatial support. This corroborates research [2, 27] suggesting that positional information processing occurs early in the model, acting as a reference frame between semantic information and the tokens themselves. We now turn to Rosetta concepts: concepts *shared* across multiple models.

5.3. Rosetta concepts

We begin by applying VTCD to the four models, targeting two classes chosen due to their dynamic nature: *rolling something across a flat surface* and *dropping something behind something*. We then mine Rosetta concepts following Section 4.1, with $\delta = 0.15$ and $\epsilon = 15\%$ in all experiments. The resulting Rosetta 4-concepts contains 40 tuples with an average R -score (Equation 7) of 17.1. Note the average R -score between all possible 4-concepts is 0.6, indicating the significance of the selected matches. Figure 6 visualizes one mined Rosetta 4-concept, which captures an object tracking representation in the late layers. This demonstrates that universal representations indeed exist between all four models. Our project website shows more shared concepts as videos.

Next, we qualitatively analyze all Rosetta d-concept with $d \in \{2, 3, 4\}$ at various layers. Figure 7 shows a represen-

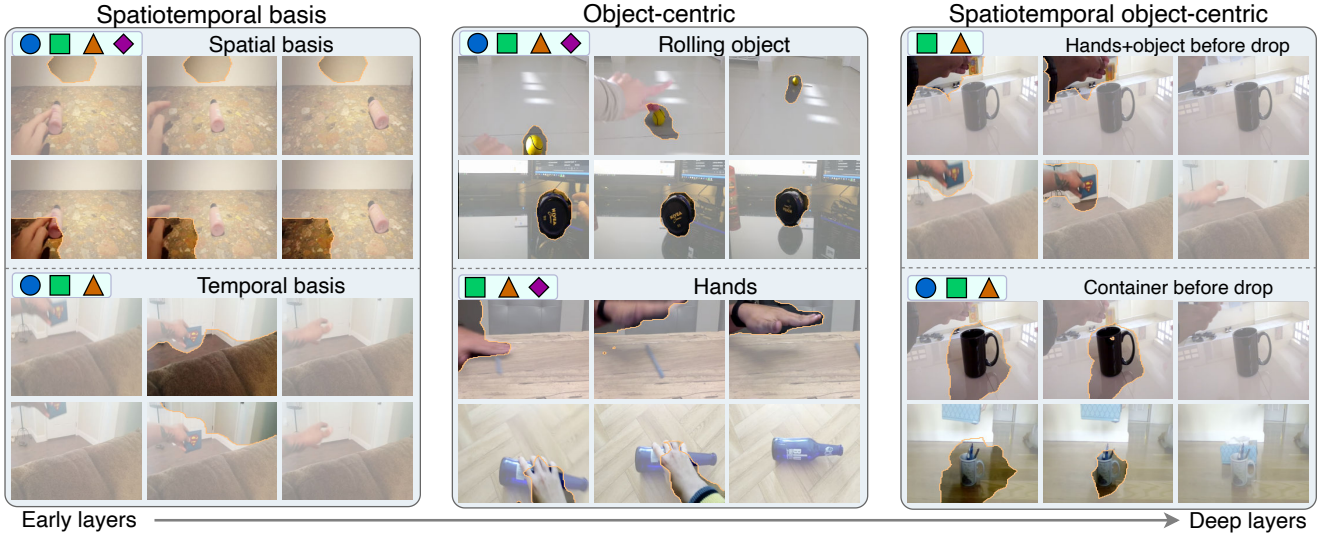


Figure 7. Universal concepts emerge in video transformers despite being trained for different tasks. Early layers encode spatiotemporal positional information. Middle layers track various objects. Deep layers capture fine-grained spatiotemporal concepts, *e.g.* related to occlusion reasoning. Legend: ● TCOW, ■ Supervised VideoMAE, ▲ SSL VideoMAE, ◆ InternVideo.

Features	VTCD	VTCD + SAM [38]
VideoMAE-SSL	45.0	68.1
VideoMAE	43.1	66.6
InternVideo	45.8	68.0

Table 3. We apply VTCD to discover object tracking concepts and evaluate them with mIoU on the DAVIS’16 validation set.

tative sample. In early layers, we find that the models learn spatiotemporal basis representations (Figure 7, left): they decompose the video space-time volume into connected regions, facilitating higher-level reasoning in later layers. This is consistent with prior works that show spatial position is encoded in early layers of image transformers [2, 27].

In the mid-layers (Figure 7, middle), we find that, among other things, all the models learn to localize and track individual objects. This result introduces a new angle to the recently developed field of object-centric representation learning [3, 26, 47, 59]: it invites us to explore how specialized approaches contribute, given that object concepts naturally emerge in video transformers. In addition, all models, except for synthetically trained TCOW, develop hand tracking concepts, confirming the importance of hands for action recognition from a bottom-up perspective [61, 73].

We validate the accuracy of the discovered object-centric concepts on the DAVIS’16 video-object segmentation (VOS) benchmark [52]. To this end, we run VTCD on the training set and select concepts that have the highest overlap with the groundtruth tracks. We then use them to track new objects on the validation set and report results in Table 3 (see appendix for implementation details). In this analysis, we solely evaluate representations that are not directly trained for VOS. While solid performance is achieved by all representations, segmentation accuracy is limited by the low resolution of the concept masks. To mitigate this,

we introduce a simple refinement step that samples random points inside the masks to prompt SAM [38]. The resulting approach, shown as ‘VTCD + SAM’, significantly improves performance. We anticipate that future developments in video representation learning will automatically lead to better VOS methods with the help of VTCD.

Finally, deeper layers contain concepts that build on object-centric representations to capture spatiotemporal events. For example, three models learn to identify containers before an object is dropped into them and two models track the object in the hand until it is dropped. One notable exception is InternVideo [70], which is primarily trained on images with limited spatiotemporal modeling. Perhaps surprisingly, these concepts are also found in the self-supervised VideoMAE [64], which was never trained to reason about object-container relationships. This raises the question: Can intuitive physics models [10, 71] be learned via large-scale video representation training?

6. Conclusion

In this work, we introduced VTCD, the first algorithm for concept discovery in video transformers. We empirically demonstrated that it is capable of extracting human-interpretable concepts from video understanding models and quantifying their importance for the final predictions. Using VTCD, we discovered shared concepts among several models with varying objectives, revealing common processing patterns, such as a spatiotemporal basis in early layers. In later layers, useful, higher-level representations universally emerge, such as those responsible for object tracking. Large-scale video representation learning is an active area of research at the moment [4, 6, 7] and our approach can serve as a key to unlocking its full potential.

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *TPAMI*, 34(11):2274–2282, 2012. 2, 3, 6
- [2] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep VIT features as dense visual descriptors. In *ECCV Workshops*, 2022. 2, 4, 5, 7, 8, 1
- [3] Zhipeng Bao, Pavel Tokmakov, Allan Jabri, Yu-Xiong Wang, Adrien Gaidon, and Martial Hebert. Discovering objects that can move. In *CVPR*, 2022. 8
- [4] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido Assran, and Nicolas Ballas. V-JEPA: Latent video prediction for visual representation learning, 2024. 8
- [5] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *CVPR*, 2017. 1, 2
- [6] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, and Adam Letts. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 8
- [7] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. 2024. 8
- [8] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *ACM FAccT*, 2018. 1
- [9] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *CVPR*, 2017. 5
- [10] Michael Chang, Tomer Ullman, Antonio Torralba, and Joshua Tenenbaum. A compositional object-based approach to learning physical dynamics. In *ICLR*, 2016. 8
- [11] Hila Chefer, Shir Gur, and Lior Wolf. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *ICCV*, 2021. 2
- [12] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *CVPR*, 2021. 2
- [13] Jinwoo Choi, Chen Gao, C. E. Joseph Messou, and Jia-Bin Huang. Why can't I dance in the mall? Learning to mitigate scene bias in action recognition. In *NeurIPS*, 2019. 3
- [14] European Commission. Laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts. *European Commission*, 2021. 1
- [15] Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. On the relationship between self-attention and convolutional layers. In *ICLR*, 2020. 2
- [16] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023. 1
- [17] Chris HQ Ding, Tao Li, and Michael I Jordan. Convex and semi-nonnegative matrix factorizations. *TPAMI*, 32(1):45–55, 2008. 2, 4
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, and Sylvain Gelly. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 4
- [19] Amil Dravid, Yossi Gandelsman, Alexei A Efros, and Assaf Shocher. Rosetta neurons: Mining the common units in a model zoo. In *ICCV*, 2023. 2, 5
- [20] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova Das-Sarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. <https://transformer-circuits.pub/2021/framework/index.html>. 2, 4, 1
- [21] Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of superposition. *Transformer Circuits Thread*, 2022. https://transformer-circuits.pub/2022/toy_model/index.html. 2
- [22] Christoph Feichtenhofer, Axel Pinz, Richard P Wildes, and Andrew Zisserman. Deep insights into convolutional networks for video recognition. *IJCV*, 128:420–437, 2020. 3
- [23] Christoph Feichtenhofer, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *NeurIPS*, 2022. 3
- [24] Thomas Fel, Victor Boutin, Mazda Moayeri, Rémi Cadène, Louis Bethune, Mathieu Chalvidal, and Thomas Serre. A holistic approach to unifying automatic concept extraction and concept importance estimation. *NeurIPS*, 2023. 1, 2, 3, 4, 5, 6
- [25] Thomas Fel, Agustin Picard, Louis Bethune, Thibaut Boissin, David Vigouroux, Julien Colin, Rémi Cadène, and Thomas Serre. CRAFT: Concept recursive activation factorization for explainability. In *CVPR*, 2023. 1, 2, 3, 4, 5, 6
- [26] Ruohan Gao, Dinesh Jayaraman, and Kristen Grauman. Object-centric representation learning from unlabeled videos. In *ACCV*, 2017. 8
- [27] Amin Ghiasi, Hamid Kazemi, Eitan Borgnia, Steven Reich, Manli Shu, Micah Goldblum, Andrew Gordon Wilson, and Tom Goldstein. What do vision transformers learn? a visual exploration. *arXiv preprint arXiv:2212.06727*, 2022. 1, 7, 8
- [28] Amir Ghodrati, Efstratios Gavves, and Cees G. M. Snoek. Video time: Properties, encoders and evaluation. In *BMVC*, 2018. 3
- [29] Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. Towards automatic concept-based explanations. In *NeurIPS*, 2019. 1, 2, 3, 6

- [30] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thureau, Ingo Bax, and Roland Memisevic. The “something something” video database for learning and evaluating visual common sense. In *ICCV*, 2017. 5
- [31] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanaprasam, Florian Golemo, and Charles Herrmann. Kubric: A scalable dataset generator. In *CVPR*, 2022. 5
- [32] Isma Hadji and Richard P Wildes. A new large scale dynamic texture dataset with application to convnet understanding. In *ECCV*, 2018. 3
- [33] Sven Ove Hansson, Matts-Åke Belin, and Björn Lundgren. Self-driving vehicles-An ethical overview. *Philosophy & Technology*, pages 1–26, 2021. 1
- [34] The White House. President Biden issues executive order on safe, secure, and trustworthy artificial intelligence. *The White House*, 2023. 1
- [35] Filip Ilic, Thomas Pock, and Richard P Wildes. Is appearance free action recognition possible? In *ECCV*, 2022. 3
- [36] Rezaul Karim, He Zhao, Richard P. Wildes, and Mennatullah Siam. MED-VT: Multiscale encoder-decoder video transformer with application to object segmentation. In *CVPR*, 2023. 2, 4, 1
- [37] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, and Fernanda Viegas. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In *ICML*, 2018. 1, 2, 6
- [38] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *ICCV*, pages 4015–4026, 2023. 8, 2
- [39] Yu Kong and Yun Fu. Human action recognition and prediction: A survey. *IJCV*, 130(5):1366–1401, 2022. 3
- [40] Matthew Kowal, Mennatullah Siam, Md Amirul Islam, Neil DB Bruce, Richard P Wildes, and Konstantinos G Derpanis. A deeper dive into what deep spatiotemporal networks encode: Quantifying static vs. dynamic information. In *CVPR*, 2022. 3
- [41] Matthew Kowal, Mennatullah Siam, Md Amirul Islam, Neil DB Bruce, Richard P Wildes, and Konstantinos G Derpanis. Quantifying and learning static vs. dynamic information in deep spatiotemporal networks. *arXiv preprint arXiv:2211.01783*, 2022. 3
- [42] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999. 2
- [43] Fuxin Li, Taeyoung Kim, Ahmad Humayun, David Tsai, and James M Rehg. Video segmentation by tracking many figure-ground segments. In *ICCV*, 2013. 3
- [44] Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *CVPR*, 2019. 3
- [45] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *ECCV*, 2018. 3
- [46] S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982. 4
- [47] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *NeurIPS*, 2020. 8
- [48] Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? *NeurIPS*, 2019. 2, 4
- [49] Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom in: An introduction to circuits. *Distill*, 2020. <https://distill.pub/2020/circuits/zoom-in>. 1
- [50] Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022. <https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>. 4, 1
- [51] Namuk Park, Wonjae Kim, Byeongho Heo, Taekyung Kim, and Sangdoon Yun. What do self-supervised vision transformers learn? In *ICLR*, 2023. 2
- [52] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 2, 3, 8
- [53] Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized input sampling for explanation of black-box models. In *BMVC*, 2018. 2
- [54] Yao Qin, Chiyuan Zhang, Ting Chen, Balaji Lakshminarayanan, Alex Beutel, and Xuezhi Wang. Understanding and improving robustness of vision transformers through patch-based negative augmentation. *NeurIPS*, 2022. 2
- [55] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *NeurIPS*, 2021. 1, 2
- [56] Kanchana Ranasinghe, Muzammal Naseer, Salman Khan, Fahad Shahbaz Khan, and Michael S Ryoo. Self-supervised video transformer. In *CVPR*, 2022. 3
- [57] Peter J Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987. 2
- [58] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. LAION-400M: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 5
- [59] Maximilian Seitzer, Max Horn, Andrii Zadaianchuk, Dominik Zietlow, Tianjun Xiao, Carl-Johann Simon-Gabriel, Tong He, Zheng Zhang, Bernhard Schölkopf, and Thomas Brox. Bridging the gap to real-world object-centric learning. In *ICLR*, 2023. 8

- [60] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5
- [61] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *ECCV*, 2018. 8
- [62] Xinyu Sun, Peihao Chen, Liangwei Chen, Changhao Li, Thomas H Li, Minghui Tan, and Chuang Gan. Masked motion encoding for self-supervised video representation learning. In *CVPR*, 2023. 3
- [63] Pavel Tokmakov, Jie Li, and Adrien Gaidon. Breaking the “object” in video object segmentation. In *CVPR*, 2023. 3
- [64] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *NeurIPS*, 2022. 3, 5, 8
- [65] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. scikit-image: image processing in python. *PeerJ*, 2:e453, 2014. 1
- [66] Basile Van Hoorick, Pavel Tokmakov, Simon Stent, Jie Li, and Carl Vondrick. Tracking through containers and occluders in the wild. In *CVPR*, 2023. 1, 2, 3, 5
- [67] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *ACL*, 2019. 2, 4, 1
- [68] Matthew Walmer, Saksham Suri, Kamal Gupta, and Abhinav Shrivastava. Teaching matters: Investigating the role of supervision in vision transformers. In *CVPR*, 2023. 1, 2
- [69] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. VideoMAE v2: Scaling video masked autoencoders with dual masking. In *CVPR*, 2023. 3
- [70] Yi Wang, Kunchang Li, Yizhuo Li, Yinan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, and Zun Wang. InternVideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. 5, 8
- [71] Jiajun Wu, Ilker Yildirim, Joseph J Lim, Bill Freeman, and Josh Tenenbaum. Galileo: Perceiving physical object properties by integrating a physics engine with deep learning. *NeurIPS*, 2015. 8
- [72] Chih-Kuan Yeh, Been Kim, Sercan Arik, Chun-Liang Li, Tomas Pfister, and Pradeep Ravikumar. On completeness-aware concept-based explanations in deep neural networks. In *NeurIPS*, 2020. 5, 6, 1
- [73] Chuhan Zhang, Ankush Gupta, and Andrew Zisserman. Helping hands: An object-aware ego-centric video recognition model. In *ICCV*, 2023. 8
- [74] Ruihan Zhang, Prashan Madumal, Tim Miller, Krista A Ehinger, and Benjamin IP Rubinstein. Invertible concept-based explanations for CNN models with non-negative concept activation vectors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021. 1, 2, 4, 6
- [75] Daquan Zhou, Zhiding Yu, Enze Xie, Chaowei Xiao, Animesh Anandkumar, Jiashi Feng, and Jose M Alvarez. Understanding the robustness in vision transformers. In *ICML*, 2022. 2