

Made to Order: Discovering monotonic temporal changes via self-supervised video ordering

Charig Yang¹ , Weidi Xie^{1,2} , and Andrew Zisserman¹ 

¹ Visual Geometry Group, University of Oxford

² CMIC, Shanghai Jiao Tong University

{charig, weidi, az}@robots.ox.ac.uk

<https://charigyang.github.io/order/>

Abstract. Our objective is to discover and localize monotonic temporal changes in a sequence of images. To achieve this, we exploit a simple proxy task of ordering a shuffled image sequence, with ‘time’ serving as a supervisory signal, since only changes that are monotonic with time can give rise to the correct ordering. We also introduce a transformer-based model for ordering of image sequences of arbitrary length with built-in attribution maps. After training, the model successfully discovers and localizes monotonic changes while ignoring cyclic and stochastic ones. We demonstrate applications of the model in multiple domains covering different scene and object types, discovering both object-level and environmental changes in unseen sequences. We also demonstrate that the attention-based attribution maps function as effective prompts for segmenting the changing regions, and that the learned representations can be used for downstream applications. Finally, we show that the model achieves the state-of-the-art on standard benchmarks for image ordering.

Keywords: Ordering · Change detection · Self-supervised learning

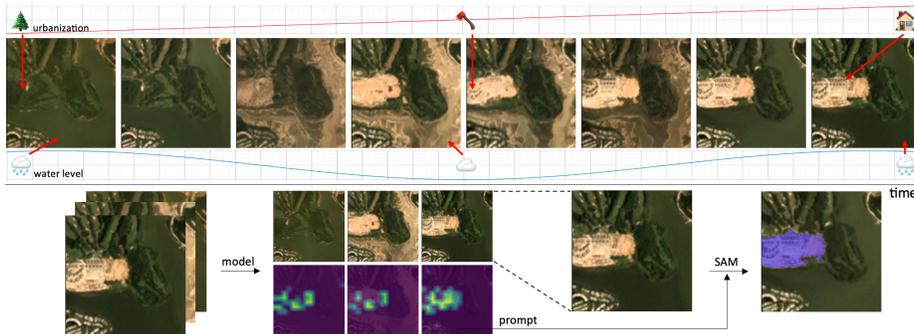


Fig. 1: Localizing monotonic temporal changes. Top: satellite images (ordered left to right) taken months apart, containing several changes – some are monotonic (*e.g.* urbanization), while others are seasonal/cyclic (*e.g.* water level). Bottom: Our model’s attribution map prediction on the sequence is able to localize the regions with monotonic temporal changes (in green), while being invariant to the seasonal and sporadic changes. The model is trained with no manual supervision, generalises to unseen sequences (as here), and the attribution map can also be used as a prompt to obtain segmentation.

1 Introduction

In the image sequence in Figure 1, there exists numerous changes over time, though many of these are seasonal, and hence distracting for long-term monitoring applications. As humans, we not only observe what is changing, but also reason about which changes are correlated monotonically with time and which ones are not. In this paper, we introduce a new task of automatically identifying temporally correlated changes in an image sequence, while being invariant to other changes. More specifically we wish to *discover* and *localize* the image regions where the change is correlated with time. Our motivation is to go beyond just detecting changes, but also discovering what changes are relevant over a period of time, and to explore the potential applications that this task could enable.

To achieve this, we propose a self-supervised proxy task, with time serving as a supervisory signal: the task is simply to order a shuffled sequence of video frames. The insight is that **frames are only orderable if changes are monotonic**, therefore if a model can order the video frames, it would have learned to identify relevant (monotonic temporal-correlated) cues while disregarding other changes. The trained model can then be employed for video analysis applications where the goal is to identify **changes** over time, such as developments or deforestation in satellite imagery (whilst ignoring seasonal variations) (see Figure 1) or aging signs in medical images. It can also be employed for detecting and tracking monotonic object **motion** over time, such as shadows caused by the movement of the sun or animals moving smoothly across the scene.

It is worth noting that temporal ordering has previously been used as a proxy task for self-supervised representation learning, with the learnt representations then finetuned for downstream tasks, such as video action recognition [19,40,47,65] (though due to the disparity between the proxy and downstream tasks, the effectiveness of learnt visual representation from temporal ordering has been unsatisfactory compared to other proxy tasks [15,25,28,29]). In contrast, the objective in this paper is to use ordering as a proxy task to directly train a model for discovering and localizing monotonic changes in video sequences, without any subsequent supervised finetuning. In this sense, we are similar in spirit to previous works that use self-supervision to directly solve tasks, such as [9,31,35,67] that targets tracking and segmentation by training on proxy tasks.

In order to harness the ordering proxy task, we introduce a transformer-based model that is able (i) to perform *ordering* on natural images sequences, and, more importantly, (ii) to provide *attribution* by localizing the evidence that gives rise to its prediction. Specifically, we use a DETR-like transformer encoder-decoder architecture where the queries in the decoder are cast as an *ordering index*. The architecture is designed to allow an attribution map to be obtained directly as part of the proxy training. Once the model has been trained for a particular setting, such as change detection in satellite image sequences, then it can generalize to unseen videos in the same setting, requiring only a forward pass to predict the localization of the monotonic temporal changes from the attribution map.

To summarize, in this paper, we make the following four contributions: (i) we introduce a new task of discovering and localizing monotonic temporal changes

in image sequences, and establish the link that temporal ordering can be used as a self-supervised proxy task for training; (ii) we introduce a flexible transformer-based architecture for ordering that can also automatically localize monotonic changes; (iii) once trained on a setting (such as satellite images), the model is able to discover the changes correlating monotonically with time in unseen image sequences in the same setting, and we demonstrate several situations where this can be applied. Finally, (iv) we demonstrate that the trained model is able to order novel image sequences surpassing the performance of previous approaches for ordering on standard benchmarks.

2 Related Work

Video self-supervised learning has become increasingly popular in computer vision. Most research in this area focuses on representation learning, with an emphasis on downstream performance after supervised fine-tuning or linear probing. In contrast, there is a less explored direction that goes beyond representation learning to directly learning a useful task under the self-supervised learning setting, such as depth [23, 71], optical flow [42, 46], correspondence [64] and sound localization [2, 3, 14, 41]. Our work in this paper builds upon such paradigm, that enables to deploy the model to downstream tasks without the need for labels.

Self-supervised learning from time. This work involves using temporal ordering as a supervisory signal. Alternative sources of supervision is to leverage other cues, such as speed [7], uniformity [68], and the arrow of time [65]. The closest kin to our work involves using temporal sequencing as supervision [19, 20, 40, 47], though their primary focus is on representation learning. In this work, we focus on advocating temporal ordering as a useful task on its own, showing that localization can emerge by using time as supervisory signal. We highlight the differences between our work and others in the Supplementary Material.

Ordering has been a longstanding task in computer science, dating back to sorting algorithms. In machine learning, it is a relevant task in both language [16, 17] and vision [5, 47, 59, 72]. For images, ordering has also been treated as a pretext for self-supervised pretraining, such as jigsaw puzzles [50], or as a task of interest, for example, image sequencing [5, 6, 32, 58, 69]. There is also some interest in the literature that focuses on differentiable sorting algorithms [13, 18, 26, 52, 53], though they mostly focus on algorithmic developments, such as differentiable loss function and black-box combinatorial optimisation. While in this paper, we make contributions towards the architecture by introducing a transformer-based ordering model, which allows ordering of arbitrary-length image sequences, with built-in visualisation via attribution maps.

Attribution localization, specifically in the case where explicit supervision is not given, has been of interest in the vision community. In ConvNets, attribution methods attempt to look into the network to find out where it is seeing [22, 70]. In transformers, several methods have been proposed to look into the attention on the [CLS] tokens [1]. Instead of these implicit localization, several methods have also carefully designed the architecture so that localization emerges explicitly

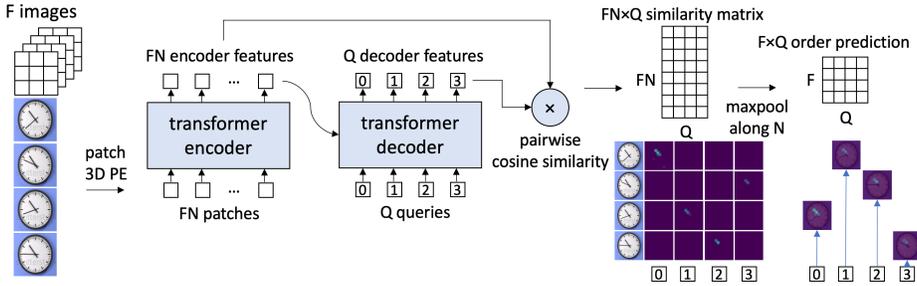


Fig. 2: Network architecture. For an unordered sequence of F frames each with N patches, the transformer encoder takes in all FN patches as input, and outputs FN features. The transformer decoder takes in Q learnable queries, each corresponding to an ordinal position, and the encoder output for cross-attention, resulting in Q features for output. A $FN \times Q$ cosine similarity matrix is constructed between all pairs of features from the encoder and decoder outputs, and the spatial max-pooling over this matrix reveals the $F \times Q$ order predictions. The ordering can simply then be obtained by taking an argmax along each query axis. In the example sequence, the hour hand is correlated monotonically with time, and appears in the attribution map.

despite not being trained on, such as in sound localization [2, 3, 14, 41]. This work follows the latter paradigm, while using the self-supervision from ordering.

Change detection has also been studied in computer vision. Many works look at changes in the image domain [54, 56], and across different applications from construction monitoring [60], satellite imaging [45], to medical imaging [51]. Other works associate short-term changes with motion, and use motion as a cue to discover moving objects [10, 11, 36, 37, 66, 67]. We differ from these lines of work in that we are mostly concerned with temporally coherent changes at different timescales, which may go beyond object level and not associated with motion.

3 Method

3.1 Problem Formulation

Our goal is to train a vision model to localize the changes in an image sequence that correlate monotonically with time. As a subsidiary goal, the model should also be able to order the image sequence.

Formally, given set of images, the model should output an *attribution map* $\mathbf{S}_{\text{att}} \in \mathbb{R}^{F \times H \times W}$ and an ordering $y_{\text{order}} \in \mathbb{Z}^F : y_{\text{order},i} \in \{0, 1, 2, \dots, F-1\}$ as:

$$y_{\text{order}}, \mathbf{S}_{\text{att}} = \Phi(\mathcal{I}_0, \mathcal{I}_1, \dots, \mathcal{I}_{F-1})$$

where $\mathcal{I} \in \mathbb{R}^{C \times H \times W}$ represents the input images. We show that we can train the model Φ via self-supervised learning on a proxy task, namely, ordering an arbitrary sequence of F images shuffled from a temporal sequence.

3.2 Ordering architecture

To address this problem, we propose a simple yet novel transformer-based architecture, as shown in Figure 2. The architecture comprises a transformer encoder (Φ_{enc}) that encodes the image patches, and a transformer decoder (Φ_{dec}) that encodes the ordering. To obtain an attribution map, we simply compute the pairwise cosine similarity between features from the encoder and queries from the decoder. We can then perform a max pooling operation across patches of the same image to get the ordering prediction.

Transformer Encoder (Φ_{enc}). To process an unordered sequence of F images, *i.e.*, $\mathcal{X} = \{\mathcal{I}_0, \mathcal{I}_1, \dots, \mathcal{I}_{F-1}\}$, we start by dividing each frame $\mathcal{I} \in \mathbb{R}^{C \times H \times W}$ into 2D patches of size (P, P) , resulting in $N = HW/P^2$ patches per frame and FN patches in total. Following the vision transformer approach, we flatten each patch using a learnable projection layer to D dimensions and add 3D positional encoding (spatial and frame) to each patch.

It is important to note that the frame positional encoding does not contain absolute temporal information since the frames are unordered, but it allows the patches to identify whether they belong to the same frame. As a result, after patchifying the input sequence, it ends up with a tensor of $\mathbf{x} \in \mathbb{R}^{F \times N \times D}$, which is then fed into a transformer encoder. The key difference to the standard vision transformer is that we output all the features, *i.e.*, $\mathbf{x}_{\text{enc}} \in \mathbb{R}^{F \times N \times D}$ instead of using a [CLS] token. In summary, we can express this procedure as $\mathbf{x}_{\text{enc}} = \Phi_{\text{enc}}(\mathcal{I}_0, \mathcal{I}_1, \dots, \mathcal{I}_{F-1})$.

Transformer Decoder (Φ_{dec}). The transformer decoder is composed of Q learnable queries $\mathbf{q} \in \mathbb{R}^{Q \times D}$, with each corresponding to an ordering position $(0, 1, \dots, Q - 1)$. The task for the transformer decoder is to align the query vector with the encoder feature that demonstrates the correct temporal order. These queries iteratively attend the visual outputs from the encoder with cross-attention in the standard transformer decoder. We denote the output of the decoder as, $\mathbf{x}_{\text{dec}} \in \mathbb{R}^{Q \times D} = \Phi_{\text{dec}}(\mathbf{q}, \mathbf{x}_{\text{enc}})$. In practice, $Q = F$.

Cosine similarity matrix (\mathbf{S}). Recall that we now possess two sets of features: encoder features $\mathbf{x}_{\text{enc}} \in \mathbb{R}^{F \times N \times D}$ and decoder features $\mathbf{x}_{\text{dec}} \in \mathbb{R}^{Q \times D}$. We then compute the pairwise cosine similarity matrix $\mathbf{S} \in \mathbb{R}^{F \times N \times Q}$: $[\mathbf{S}]_{i,j} = \cos(\mathbf{x}_{\text{enc},i}, \mathbf{x}_{\text{dec},j}) \in [-1, 1]$ between each i of the $F \times N$ features in \mathbf{x}_{enc} and each j of the Q features in \mathbf{x}_{dec} , where $\cos(\cdot, \cdot)$ denotes the cosine similarity function.

Given the similarity matrix, we want to obtain (i) the ordering of the frames and (ii) the attribution map that indicates the spatial evidence within each frame that gives rise to ordering. The matrix $\mathbf{S} \in \mathbb{R}^{F \times N \times Q}$ consists of $F \times Q$ different spatial maps of size N , each indicating the attention between each pair of queries ($j \in Q$) and images ($i \in F$).

Order prediction. To obtain the order predictions, we perform spatial max-pooling over the patches of each frame (along the N dimension), to obtain $\hat{y} \in \mathbb{R}^{F \times Q} = \max_{i \in N} \mathbf{S}_i$. This max-pooling is designed to create an information bottleneck – the query has to attend to the correct token(s) within the correct image in order to predict the order correctly. The resulting matrix serves as a

predictor for the position of each query in the ordering. We then apply a softmax along the query axis of the matrix to get the probability scores for each query.

Attribution map. Among the $F \times Q$ different spatial maps, we are only interested in the ones that correspond to the correct ordering. For each query $j \in Q$, we select one map $i \in F$ that has the maximum activation, *i.e.* $i = \arg \max \hat{y}_j$ resulting in Q maps. We then rearrange and resize each map of N patches back to the original resolution, resulting in $\mathbf{S}_{\text{att}} \in \mathbb{R}^{Q \times H \times W}$. Notably, this localization can be achieved without the need for additional fine-tuning, supervision, or post-hoc attribution methods [1, 22].

3.3 Training and inference

Temporal loss. Given the ground-truth order $y \in \mathbb{Z}^Q : y_i \in \{0, 1, \dots, F - 1\}$, the model can be trained via binary cross-entropy loss. Specifically, we convert the ground-truth order into a binary permutation matrix. With some notation abuse, we still call this matrix $y \in \mathbb{Z}^{F \times Q}$. The forward loss is then simply the elementwise binary cross-entropy between the two matrices: $\mathcal{L}_f(y, \hat{y}) = \frac{1}{FQ} \sum_{i \in F} \sum_{j \in Q} \text{cross-entropy}(\hat{y}_{ij}, y_{ij})$.

In practice, we find that allowing reversibility in the loss aids with training, as many changes are reversible in nature without prior knowledge of the arrow of time [65] (the sequence could equally be ordered from first to last, or last to first). To allow this, we calculate the loss as the minimum of the loss for both forward and backward sequences, *i.e.* $\mathcal{L}_r = \min(\mathcal{L}_f(y, \hat{y}), \mathcal{L}_f(\text{reverse}(y), \hat{y}))$. This loss is zero when the model predicts the order correctly in either direction.

Inference. At inference time, we simply take the argmax along each query axis as the order prediction, that is, $y_{\text{order}} \in \mathbb{Z}^Q : y_{\text{order},j} = \arg \max_{i \in F} \hat{y}_{i,j}$. In other words, each query picks the image that contains the maximum activation for its query, as illustrated in the bottom-right corner of Figure 2.

3.4 Discussion

Generalization to different sequence lengths. Our architecture is designed to handle sequences of arbitrary, possibly unequal length during training and inference, without the need to re-design or train separate models for each sequence length. At training time, we assume there is a maximum number of images, thus initialize a total of F_{max} learnable queries in the transformer, *i.e.*, $\mathbf{q} \in \mathbb{R}^{F_{\text{max}} \times D}$. While the model handles a sequence of F images, with $F \leq F_{\text{max}}$, it only uses the first F queries as input to the decoder, ignoring the rest. This approach enables each query to represent its positioning $(0, 1, \dots, F - 1)$, making it generalizable to different lengths during both training and testing. However, the model will not generalize to lengths above F_{max} .

Avoiding trivial solutions. There are two factors that we need to account for: camera motion and video compression artifacts. Camera motion can be smooth or uniform over a short time gap, which can result in an uninteresting cue. To address this, we apply a small random cropping on each frame in settings where the time gap between frames is small (*i.e.* $< 1\text{s}$). This slight jittering helps to

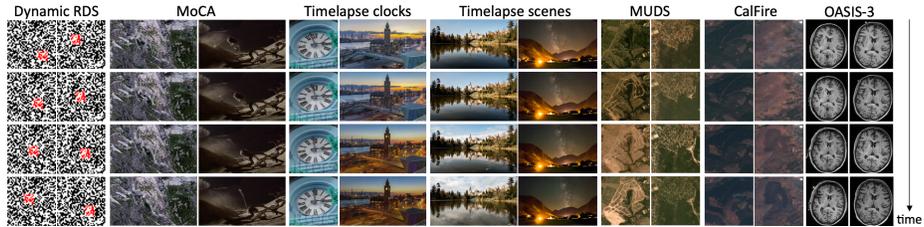


Fig. 3: Sequence datasets. From left to right: dynamic Random Dot Stereograms (RDS) (moving dots colored only for illustration), moving camouflaged animals (MoCA), timelapse clocks (cropped/full), timelapse scenes, MUDS, CalFire, OASIS-3.

prevent the model from learning trivial solutions. We note that this does not degrade the performance even if camera motion is absent. Another factor that can give rise to trivial solutions is inter-frame video compression artifacts. To address this, we follow conventional wisdom [30, 65] and use H.264 formatting for all videos, thus minimizing compression artifacts and preventing trivial solutions.

From localization to segmentation. While the attribution map is useful, some applications may benefit from going beyond just localization. Here, we propose three solutions. (i) We can directly obtain segmentation at patch-level granularity by thresholding the attribution map, together with minimal post-processing namely averaging across frames and removing small contours. (ii) We replace the linear projection layer on image patches with a pretrained DINOv2 to enhance the feature quality. (iii) Alternatively, we can use the highest activation points (*i.e.* centre pixels of patches) as prompt for the pretrained Segment Anything model (SAM) [34], to obtain pixel-level segmentation masks.

4 Experiments

4.1 Video datasets

To leverage the inherent temporal information in videos, we sample a sequence from a video, shuffle the frames, then train the model with the original ordering as groundtruth. This enables the attribution map to identify monotonic changes that contribute to the ordering while disregarding other changes. Our study explores sequences across multiple domains, each with distinct cues of interest, which we summarize in Table 1 and illustrate sample sequences in Figure 3.

Dynamic random dot stereograms are a type of image sequence that features a box of random dots moving smoothly over a background of random dots, originally used as a pair to demonstrate stereoscopic motion [48]. While the individual images may appear random, the box is visible when viewed in sequence. We generate a synthetic dataset of controllable dynamic random dot stereograms (Dynamic RDS) to test the model’s ability to detect subtle relative cues. Since this a synthetic dataset, we know the ground-truth of the box’s motion, so can compare this with the predictions.

Moving camouflaged animals (MoCA) [37] was constructed from videos of camouflaged animals. We use this dataset to investigate the use of short-term object locomotion as a cue, particularly where it is challenging to distinguish the object from the background. To accomplish this, we follow [67] and focus on the subset of 88 videos in which the animals are in motion. To evaluate on localisation, we assume that the change is object-level due to animal motion, and use the annotated object bounding box as the ground-truth for localization.

Timelapse analog clocks. Our study examines real-world scenes that feature both absolute cues (time on clocks) and relative cues (scene changes). To accomplish this, we utilize the Timelapse dataset [68] in two ways. Firstly, we use the entire dataset of 2,511 videos that features cropped clocks. Secondly, we create a subset consisting of 260 outdoor scenes with static cameras where the clock occupies only a small area of the scene.

Timelapse scenes. We gathered outdoor timelapse videos from WebVid-10M dataset [4] with ‘timelapse’ as the search query. The dataset comprises 180 static videos, and our aim is to investigate the cues that the model can extract from the scene to learn its order, as there are no specific absolute cues present.

4.2 Temporal image sequences

We also show effectiveness of our approach for image sequences that are collected over a period of time, including satellite images and longitudinal medical data.

Multi-temporal urban development dataset (MUDS) [62] is a dataset that contains 80 aerial satellite image sequences captured monthly over a two-year period. We aim to identify geographical changes that occur over time, including deforestation and urbanization, while ignoring other changes. As there are no existing datasets and benchmarks for this task, we hand-label 60 sub-sequences on the test set for segmentation masks where changes are monotonic, and call this evaluation set **Monotonic MUDS**. Samples are shown in Figure 4a. We use this dataset to evaluate localization and segmentation performance of the model trained on the MUDS dataset.

CalFire [44] is a satellite dataset tracking wildfires in California. It also contains other events, such as snowfall, new construction, changes in water level, that we aim to discover. We pre-process by removing scenes with significant cloud cover.

OASIS-3 [38] contains longitudinal MRI scans taking 1-4 years apart to study how brains age. We select sequences with 3 or more scans, resulting in 134 sequences. Following [32], we perform affine registration and use the centre slice.

4.3 Image ordering datasets

In addition, we showcase our general ordering capability by evaluating our method on standard benchmarks. We compare our ordering performance with related works [18, 26, 52, 53] on the task of sorting images of numbers in ascending order, as shown in Figure 4b. What the model has to learn here is different from the previous datasets, as the ordering is *absolute*, and not understanding changes.

dataset	Δt	cue	seq (trn/test)	EM	EW
Dynamic RDS	<1s	motion	∞	99.8	99.9
MoCA	<1s	motion	75/13	82.0	90.6
Clocks (cropped)	$\sim 1m$	clock	2011/500	62.5	73.0
Clocks (full)	$\sim 1h$	clock/scene	210/20	55.0	74.3
Timelapse scenes	$\sim 1h$	scene	130/50	61.8	79.4
MUDS	1mo	landscape	60/20	56.4	69.6
CalFire	1mo	landscape	800/276	76.6	87.5
OASIS-3	1y	brain	100/34	84.3	89.1

Table 1: Dataset attributes and ordering results. This table shows different datasets and their attributes, as well as the ordering results on the test (unseen) sequences on exact match (EM) and elementwise (EW) metrics.

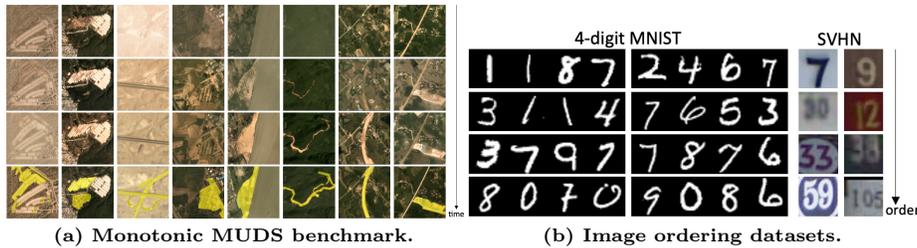


Fig. 4: (a) To evaluate localization and segmentation performance, we manually annotate the monotonically changing regions (shown in yellow) on the MUDS test set. Each sequence contains four frames, and the monotonic changes between the first and last frames are annotated. (b) 4-digit MNIST [39] (left) and SVHN [49] (right). The task is to order the images by the numbers they contain in increasing order (top to bottom).

Multi-digit MNIST [39] dataset is a modified version of the MNIST dataset in which four digits are concatenated to form a four-digit number. The goal is to order the image sequences in increasing order. To construct this dataset, we synthetically combine examples from the corresponding train and test sets of the MNIST dataset, resulting in a total of 50,000 training and 10,000 testing images. **Street view house numbers (SVHN)** [49], was collected from Google Street View and includes images of house numbers. Similarly, the task is to order these numbers in increasing order. The dataset consists of 33,402 images for training and 13,068 for testing. To ensure consistency with previous studies, we followed the data preprocessing and augmentation methods described in [24].

4.4 Evaluation metrics

Localization. We use a pointing-game evaluation method, this is to follow the convention of the localization literature in other domains, including audio-visual localization [2,3] and saliency methods [21,22], that is, if the pixel with maximum activation in the attribution map contains the change of interest, then it is

positive (1), otherwise it is negative (0). As our method only outputs patch-level attribution, we simply select the centre pixel of the patch as the highest activation. The overall accuracy is then calculated as the average over all sequences.

Segmentation. We use the standard metric for segmentation, *i.e.*, mean intersection over union (mIoU), where the mean is the average across all sequences.

Ordering. We use the evaluation metrics for sorting as outlined in [52, 53]. These metrics include exact match (EM) and elementwise (EW) accuracy. EM is considered correct if the entire sequence is ordered correctly, while EW considers the order accuracy per element. Following previous benchmarks we evaluate these at sequence lengths 5, 9 and 16. To test the generalization to different sequence lengths, we also evaluate the exact match accuracy at a fixed sequence length of 5 at test-time (EM5), regardless of the training sequence length.

4.5 Implementation details

We split each dataset into disjoint training and testing subsets, and then randomly sample frames from each video. We keep the time gap between sampled frames constant within each video, but vary this across videos to train a robust model. For image sequences (MUDS, CalFire, OASIS-3) where data collection is less regular, we relax these constraints and simply randomly sample between all frames within the sequence. We train each model separately for each dataset.

Architecture. For the encoder, we use a smaller version of TimeSFormer [8] with a divided space-time attention architecture, consisting of 256 dimensions, 4 heads, 6 layers, and 512 MLP dimensions. For the decoder, we use a standard transformer decoder with the same parameters, except for 64 dimensions and 3 layers. As a result, the model is lightweight, with only 4M parameters. We use Adam optimizer [33] with learning rate $1e-4$ in all experiments, and batch size 32 sequences with 4 frames per sequence for all video datasets, except 3 frames for OASIS-3 as this is the minimum MRI scan sessions per subject. For image ordering, we use batch size 100 with varying numbers of frames. All experiments are run on a single GPU. The code, datasets, and models will be released.

5 Results

5.1 Results on ordering video frames or image sequence

Video ordering results. The results for ordering as well as the main signals that the model can pick up, are shown by EM and EW in Table 1, with qualitative results in Figure 5. The model is able to successfully order sequences across different datasets, particularly in cases where changes are significant. It is expected that the scores are not perfect, as sampled data from videos is not guaranteed to contain ordering cues, as illustrated in the Supplementary.

Temporal image sequence ordering results. For satellite image sequences, the model is able to discover cues that are relevant to ordering, including road building and forest fires. This illustrates our model’s application on remote

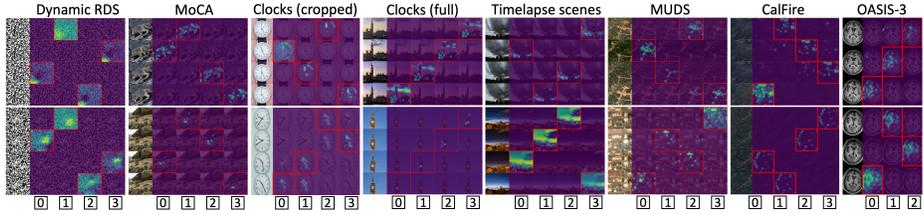


Fig. 5: Ordering and Localization results across various datasets, where the model is able to discover and localize various cues across different domains, including object motion, clocks, scenery, landscape and biological aging. The left column shows the input (unordered) images. Each column of the similarity matrix represents the model’s prediction of each individual order (0, 1, 2, 3), where the image in the red box is chosen, and the attention heat map within the box localizes the change.

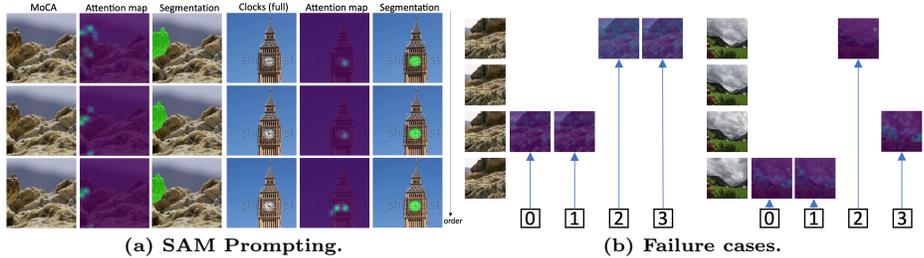


Fig. 6: (a) The attribution map is used to prompt SAM to obtain segmentation masks. (b) Unorderable sequences, one being too static and the other being too stochastic.

sensing imagery. In MRI scans, we explore the cues for age changes. Our results show that there are cues in the posterior part of the brain. This is consistent with the literature [12] that suggest that ventricular enlargement is a prominent feature, and causes the posterior horn to inflate in response to tissue loss. There are also some cues along the outline. This is in line with the literature [57, 61], which suggest that brain volume also decreases with old age.

Failure cases. A limitation of our model is that we do not force a one-to-one matching between queries and images, and this may result in some images being claimed by multiple different queries or by none at all, as seen in Figure 6b. This problem can easily be resolved by allowing each image to be predicted once. However, not being able to order also provides valuable information as not all real sequences can be ordered – for example, sequences where everything is static for a period, or very stochastic. Therefore, we treat invalid orderings as a means to provide information on whether particular frames can be ordered or not. We report further experimental analyses on the failure cases in the Supplementary.

5.2 Comparison with change detection methods

We compare against three previous approaches for the ability to detect monotonic temporal changes on Monotonic MUDS dataset, namely, a supervised Siamese

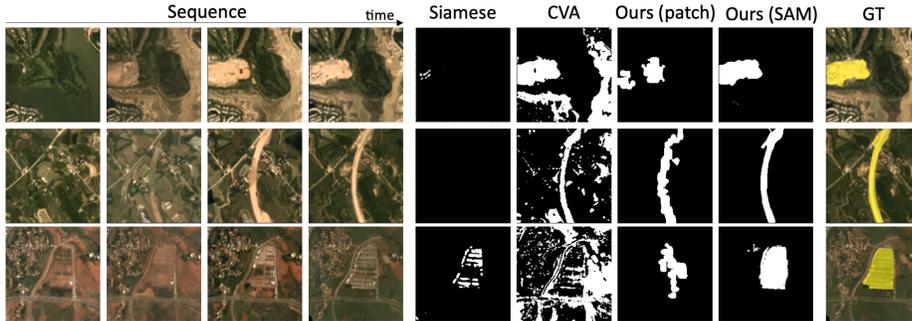


Fig. 7: Segmentation comparison with other methods: Siamese networks [27] and CVA [43]. Supervised methods ignore changes other than building, and pixel-based methods over-segments non-monotonic regions.

Method	Mono-MUDS		RDS	MoCA
	loc (acc) ↑	seg (mIoU) ↑	seg (mIoU) ↑	loc (acc) ↑
Siamese [27]	73.3	11.1	–	–
CVA [43]	71.7	34.6	22.3	69.6
DCVA [55]	70.0	35.5	–	–
Ours	83.3	37.9	34.2	75.0
Ours + DINOv2	80.0	41.3	–	–
Ours + SAM	83.3	45.1	–	–

Table 2: Localization and segmentation results via the pointing game accuracy for localization, and mIoU for segmentation. The methods for segmentation (patch and SAM) are described in Section 3.4.

Networks [27] trained on MUDS dataset for the task of *urban development tracking*, by highlighting the differences between buildings; Change Vector Analysis [43], which is a baseline for the task of *change detection*, highlighting all changes in an image pair without knowledge their nature, and Deep CVA [55], a learned version of CVA that has been trained specifically on images.

The performance comparison is given in Table 2, and illustrated in Figure 7. As can be seen, the urban development tracking method [27] under-segments changes that are correlated with urbanisation, as it is only trained to look at buildings. The change detection methods of [43, 55] over-segments changes that are non-monotonic as it has no concept of time. Both prior methods have shortcomings in detecting urban development: the former method misses changes like road building and deforestation, and the latter includes many erroneous regions such as the seasonal changes in vegetation and water level. Our model is able to highlight correctly the monotonic changes while being invariant to other changes. We further note that our model *discovers* such changes without any prior information on what to look for. Additionally, replacing the patch projection with DINOv2 and using the localization to prompt SAM both improve the results. We also evaluate on two other datasets: RDS and MoCA, where we have ground-truth for the moving objects in the video; and show that we obtain

methods	loc (acc) \uparrow	seg (mIoU) \uparrow	finetuning (F1) \uparrow
scratch	—	—	18.2
S&L [47]	23.3	20.9	15.8
OPN [40]	25.0	24.1	17.1
Ours (AoT proxy)	63.3	26.9	25.5
Ours	83.3	45.1	30.2

Table 3: Comparison with self-supervised methods on localization and segmentation on Mono-MUDS and on fine-tuning for building change detection on MUDS.

favourable results. Note that Siamese and DCVA are only trained on satellite images, hence do not generalize to other domains.

Quantitatively, in Table 2, we find that (i) our pointing-game localization and patch-level segmentation results outperform other methods despite operating only on patch-level (7×7) and not pixel-level granularity, and (ii) segmentation via SAM prompting further improves the results.

Qualitatively, the results of our localization experiments on various video datasets are presented in Figure 5. The model is capable of accurately identifying monotonic changes while remaining invariant to unrelated cues. Notably, these results were achieved on sequences unseen during training. We additionally show that our localization map works as a good prompt for the Segment Anything (SAM) model to obtain object-level changes, as in Figure 6a.

5.3 Comparison on self-supervised proxy tasks

Here, we compare to other self-supervised methods based on *time*. We include a discussion in the Supplementary on the subtle differences between the proxy tasks. We include results for training baselines from scratch on MUDS, and testing on Mono-MUDS in localizing and segmenting monotonic changes. The results are shown in Table 3, where we conduct three sets of experiments, as detailed below.

First, we observe that previous methods are extremely crude in localization; this is expected, as they are all based on conv5 feature (even AoT, via CAM). Given a 224p input, conv5 has a 13×13 feature map with 195p receptive field. Our architecture handles localization by design, and is hence more capable than other methods that use post-hoc attribution methods on top of standard backbones. We also note that AoT does not train, which is also expected as it ingests optical flow as input (and in satellite images using flow does not make sense (change \neq motion)), our method is more flexible in this regard.

Second, we ask if our method is still superior if the architectural gap is closed. To achieve this, we also compare the proxy task in AoT (time direction) with ours (ordering) under a fairer comparison (using our architecture and RGB input), and show this in the table under “Ours (AoT proxy)”. We conclude that both our architecture and proxy task leads to significant improvements.

Third, we investigate fine-tuning on a target task that requires both time and localization: change detection of buildings (like [27]). For each method, we pre-train the encoder on MUDS, freeze it, then train a lightweight head ($3 \times$

dataset	MNIST						SVHN					
frames	5		9		16		5		9		16	
DSort [52]	83.4 92.6	56.3 86.7	30.5 80.7	86.6	64.1 82.1	24.2 69.6	3.9 59.6	66.8				
DSv2 [53]	84.9 93.2	63.8 89.1	31.1 82.2	—	68.5 84.1	39.9 75.8	12.2 65.6	—				
Ptr-Net [63]	91.9 95.6	87.7 95.0	68.9 90.0	1.1	76.3 87.6	48.7 79.4	9.8 63.2	0.1				
Ours	93.9 96.7	87.9 95.2	72.2 91.2	92.9	77.3 88.2	53.9 81.0	19.4 67.9	67.6				

Table 4: Ordering on image datasets on two standard benchmarks (MNIST and SVHN) where the task is to order images of numbers in increasing order. Metrics are (EM|EW|EM5). EM and EW are evaluated at the sequence length the model has been trained on (5/9/16), whereas EM5 tests generalisation to test length 5.

deconvolutions) on top of the same dataset, and test on unseen sequences. To keep this fair, we keep the number of parameters roughly the same across methods. The results (Table 3, right column) show that our method learns good representations as compared to previous self-supervised methods in localization tasks.

5.4 Comparison with image ordering methods

We compare to two previous methods on image ordering benchmarks where the task is to arrange the images in increasing order. *Differentiable sorting networks* such as DiffSort [52] and its successor DSortv2 [53] employ a parameter-free sorting network to rank scalars. *Pointer Networks* [63] ranks features by using a recurrent encoder-decoder network with attention. We note that Ptr-Net is not initially designed for such a task, but for arranging a set of coordinates. We simply extend pointer networks by adding an image encoder and task the model to rank the image features from small to large. We then jointly train this encoder and the pointer network. For fair comparison, we use the same transformer encoder as our model, and use the pointer network as the decoder with similar size to our transformer decoder.

The quantitative results are presented in Table 4. Our results demonstrate that (i) we compare favourably in ordering performance – on both MNIST and SVHN, our method has the best performance of the four (ii) Ptr-Net does not automatically generalise to testing with different sequence lengths, while our method does (as reflected by the poor EM5 accuracy), and (iii) our method also has the added benefit of having an attribution map.

6 Conclusion

In this paper, we explore using time as a proxy loss for self-supervised training of models to discover and localize monotonic temporal changes in image sequences. Possible extensions include discovering more complex temporal changes (seasonal/periodic), or object state and attribute changes. It would also be interesting to investigate how the model scales with larger datasets and compute, and what new applications this task can enable. Overall, we hope this paper presents a valuable starting point for future research and applications in this area.

Acknowledgements

We thank Tengda Han, Ragav Sachdeva, and Aleksandar Shtedritski for suggestions and proofreading. This research is supported by the UK EPSRC CDT in AIMS (EP/S024050/1), and the UK EPSRC Programme Grant Visual AI (EP/T028572/1).

References

1. Abnar, S., Zuidema, W.: Quantifying attention flow in transformers. arXiv preprint arXiv:2005.00928 (2020)
2. Afouras, T., Owens, A., Chung, J.S., Zisserman, A.: Self-supervised learning of audio-visual objects from video. In: ECCV (2020)
3. Arandjelovic, R., Zisserman, A.: Objects that sound. In: ECCV (2018)
4. Bain, M., Nagrani, A., Varol, G., Zisserman, A.: Frozen in time: A joint video and image encoder for end-to-end retrieval. In: ICCV (2021)
5. Basha, T., Moses, Y., Avidan, S.: Photo sequencing. In: ECCV (2012)
6. Basha, T.D., Moses, Y., Avidan, S.: Space-time tradeoffs in photo sequencing. In: ICCV (2013)
7. Benaim, S., Ephrat, A., Lang, O., Mosseri, I., Freeman, W.T., Rubinstein, M., Irani, M., Dekel, T.: Speednet: Learning the speediness in videos. In: CVPR (2020)
8. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: ICML (2021)
9. Bian, Z., Jabri, A., Efros, A.A., Owens, A.: Learning pixel trajectories with multi-scale contrastive random walks. In: CVPR (2022)
10. Bideau, P., Learned-Miller, E.: A detailed rubric for motion segmentation. arXiv preprint arXiv:1610.10033 (2016)
11. Bideau, P., Learned-Miller, E.: It’s moving! a probabilistic model for causal motion segmentation in moving camera videos. In: ECCV (2016)
12. Blinkouskaya, Y., Weickenmeier, J.: Brain shape changes associated with cerebral atrophy in healthy aging and alzheimer’s disease. *Frontiers in Mechanical Engineering* (2021)
13. Brown, A., Xie, W., Kalogeiton, V., Zisserman, A.: Smooth-ap: Smoothing the path towards large-scale image retrieval. In: ECCV (2020)
14. Chen, H., Xie, W., Afouras, T., Nagrani, A., Vedaldi, A., Zisserman, A.: Localizing visual sounds the hard way. In: CVPR (2021)
15. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2021)
16. Chen, X., Qiu, X., Huang, X.: Neural sentence ordering. arXiv preprint arXiv:1607.06952 (2016)
17. Cui, B., Li, Y., Chen, M., Zhang, Z.: Deep attentive sentence ordering network. In: EMNLP (2018)
18. Cuturi, M., Teboul, O., Vert, J.P.: Differentiable ranking and sorting using optimal transport. In: NeurIPS (2019)
19. Fernando, B., Bilen, H., Gavves, E., Gould, S.: Self-supervised video representation learning with odd-one-out networks. In: CVPR (2017)
20. Fernando, B., Gavves, E., Oramas, J.M., Ghodrati, A., Tuytelaars, T.: Modeling video evolution for action recognition. In: CVPR (2015)

21. Fong, R., Patrick, M., Vedaldi, A.: Understanding deep networks via extremal perturbations and smooth masks. In: ICCV (2019)
22. Fong, R.C., Vedaldi, A.: Interpretable explanations of black boxes by meaningful perturbation. In: ICCV (2017)
23. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: ICCV (2019)
24. Goodfellow, I.J., Bulatov, Y., Ibarz, J., Arnoud, S., Shet, V.: Multi-digit number recognition from street view imagery using deep convolutional neural networks. arXiv preprint arXiv:1312.6082 (2013)
25. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent: A new approach to self-supervised learning. In: NeurIPS (2021)
26. Grover, A., Wang, E., Zweig, A., Ermon, S.: Stochastic optimization of sorting networks via continuous relaxations. arXiv preprint arXiv:1903.08850 (2019)
27. Hafner, S., Ban, Y., Nascetti, A.: Urban change detection using a dual-task siamese network and semi-supervised learning. In: IGARSS (2022)
28. Han, T., Xie, W., Zisserman, A.: Self-supervised co-training for video representation learning. In: NeurIPS (2020)
29. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: CVPR (2022)
30. Iashin, V., Xie, W., Rahtu, E., Zisserman, A.: Sparse in space and time: Audio-visual synchronisation with trainable selectors. In: BMVC (2022)
31. Jabri, A., Owens, A., Efros, A.A.: Space-time correspondence as a contrastive random walk. In: NeurIPS (2020)
32. Kim, H., Sabuncu, M.R.: Learning to compare longitudinal images. arXiv preprint arXiv:2304.02531 (2023)
33. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
34. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: ICCV (2023)
35. Lai, Z., Lu, E., Xie, W.: Mast: A memory-augmented self-supervised tracker. In: CVPR (2020)
36. Lamdouar, H., Xie, W., Zisserman, A.: Segmenting invisible moving objects. In: BMVC (2021)
37. Lamdouar, H., Yang, C., Xie, W., Zisserman, A.: Betrayed by motion: Camouflaged object discovery via motion segmentation. In: ACCV (2020)
38. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A.G., et al.: Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. MedRxiv (2019)
39. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proceedings of the IEEE (1998)
40. Lee, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Unsupervised representation learning by sorting sequences. In: ICCV (2017)
41. Liu, J., Ju, C., Xie, W., Zhang, Y.: Exploiting transformation invariance and equivariance for self-supervised sound localisation. In: ACM MM (2022)
42. Liu, P., Lyu, M., King, I., Xu, J.: Selfflow: Self-supervised learning of optical flow. In: CVPR (2019)

43. Malila, W.A.: Change vector analysis: An approach for detecting forest changes with landsat. In: LARS symposia (1980)
44. Mall, U., Hariharan, B., Bala, K.: Change event dataset for discovery from spatio-temporal remote sensing imagery. In: NeurIPS (2022)
45. Mall, U., Hariharan, B., Bala, K.: Change-aware sampling and contrastive learning for satellite images. In: CVPR (2023)
46. Meister, S., Hur, J., Roth, S.: Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In: AAAI (2018)
47. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: ECCV (2016)
48. Neff, R., Schwartz, S., Stork, D.G.: Electronics for generating simultaneous random-dot cyclopean and monocular stimuli. *Behavior Research Methods, Instruments, & Computers* (1985)
49. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning. In: NeurIPS (2011)
50. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: ECCV (2016)
51. Patriarche, J., Erickson, B.: A review of the automated detection of change in serial imaging studies of the brain. *Journal of digital imaging* (2004)
52. Petersen, F., Borgelt, C., Kuehne, H., Deussen, O.: Differentiable sorting networks for scalable sorting and ranking supervision. In: ICML (2021)
53. Petersen, F., Borgelt, C., Kuehne, H., Deussen, O.: Monotonic differentiable sorting networks. arXiv preprint arXiv:2203.09630 (2022)
54. Sachdeva, R., Zisserman, A.: The change you want to see. In: WACV (2023)
55. Saha, S., Bovolo, F., Bruzzone, L.: Unsupervised deep change vector analysis for multiple-change detection in vhr images. *IEEE Transactions on Geoscience and Remote Sensing* **57**(6), 3677–3693 (2019)
56. Sakurada, K., Okatani, T.: Change detection from a street image pair using cnn features and superpixel segmentation. In: BMVC (2015)
57. Scahill, R.I., Frost, C., Jenkins, R., Whitwell, J.L., Rossor, M.N., Fox, N.C.: A longitudinal study of brain volume changes in normal aging using serial registered magnetic resonance imaging. *Archives of neurology* **60**(7), 989–994 (2003)
58. Sevilla-Lara, L., Zha, S., Yan, Z., Goswami, V., Feiszli, M., Torresani, L.: Only time can tell: Discovering temporal data for temporal modeling. In: WACV (2021)
59. Shvetsova, N., Petersen, F., Kukleva, A., Schiele, B., Kuehne, H.: Learning by sorting: Self-supervised learning with group ordering constraints. ICCV (2023)
60. Stent, S., Gherardi, R., Stenger, B., Cipolla, R.: Detecting change for multi-view, long-term surface inspection. In: BMVC (2015)
61. Svennerholm, L., Boström, K., Jungbjer, B.: Changes in weight and compositions of major membrane components of human brain during the span of adult human life of swedes. *Acta neuropathologica* (1997)
62. Van Etten, A., Hogan, D., Manso, J.M., Shermeyer, J., Weir, N., Lewis, R.: The multi-temporal urban development spacenet dataset. In: CVPR (2021)
63. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer networks. In: NeurIPS (2015)
64. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: CVPR (2019)
65. Wei, D., Lim, J.J., Zisserman, A., Freeman, W.T.: Learning and using the arrow of time. In: CVPR (2018)
66. Xie, J., Xie, W., Zisserman, A.: Segmenting moving objects via an object-centric layered representation. In: NeurIPS (2022)

67. Yang, C., Lamdouar, H., Lu, E., Zisserman, A., Xie, W.: Self-supervised video object segmentation by motion grouping. In: ICCV (2021)
68. Yang, C., Xie, W., Zisserman, A.: It's about time: Analog clock reading in the wild. In: CVPR (2022)
69. Zarrabi, N., Avidan, S., Moses, Y.: Crowdcam: Dynamic region segmentation. arXiv preprint arXiv:1811.11455 (2018)
70. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)
71. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: CVPR (2017)
72. Zhukov, D., Alayrac, J.B., Laptev, I., Sivic, J.: Learning actionness via long-range temporal order verification. In: ECCV (2020)