

Video OWL-ViT: Temporally-consistent open-world localization in video

Georg Heigold
Google DeepMind

Matthias Minderer
Google DeepMind

Alexey Gritsenko
Google DeepMind

Alex Bewley
Google DeepMind

Daniel Keysers
Google DeepMind

Mario Lučić
Google DeepMind

Fisher Yu[†]
ETH Zurich

Thomas Kipf
Google DeepMind

Abstract

We present an architecture and a training recipe that adapts pretrained open-world image models to localization in videos. Understanding the open visual world (without being constrained by fixed label spaces) is crucial for many real-world vision tasks. Contrastive pre-training on large image-text datasets has recently led to significant improvements for image-level tasks. For more structured tasks involving object localization applying pre-trained models is more challenging. This is particularly true for video tasks, where task-specific data is limited. We show successful transfer of open-world models by building on the OWL-ViT open-vocabulary detection model and adapting it to video by adding a transformer decoder. The decoder propagates object representations recurrently through time by using the output tokens for one frame as the object queries for the next. Our model is end-to-end trainable on video data and enjoys improved temporal consistency compared to tracking-by-detection baselines, while retaining the open-world capabilities of the backbone detector. We evaluate our model on the challenging TAO-OW benchmark and demonstrate that open-world capabilities, learned from large-scale image-text pretraining, can be transferred successfully to open-world localization across diverse videos.

1. Introduction

A central goal in computer vision is to develop models that can understand diverse and novel scenarios in the visual world. While this has been difficult for methods developed on datasets with closed label spaces, web-scale image-text pretraining has recently led to dramatic improvements in

open-world performance on a range of *image-level* vision tasks [12, 26, 19].

However, challenges still remain for *object-level* tasks on images and especially videos. First, object-level tasks require predicting more complex output structures compared to image-level tasks, making transfer of pretrained models more challenging. Second, training data for structured tasks is limited due to the prohibitive labeling cost. Therefore, a key research question is how to transfer the open-vocabulary capabilities of image-text models to object-level tasks like object detection and tracking.

For object detection, works such as ViLD [12], Region-CLIP [40], OWL-ViT [26], F-VLM [19] etc. demonstrate that image-level open-vocabulary capabilities can be transferred to object detection with relatively little detection-specific training data. Most recent works achieve this by combining image-text pre-trained encoder backbones (e.g. CLIP [27]) with detection heads. By transferring semantic knowledge from the backbone, the resulting models are capable of detecting objects for which no localization annotations were present in the detection training data.

Here, we extend this approach to video. We build on OWL-ViT [26], which provides a simple open-world detection architecture in which light-weight box prediction and classification heads are trained on top of a CLIP backbone. We transfer the open-world capabilities of OWL-ViT to video understanding with minimal video-specific training data. The key idea behind our approach is to apply the open-world detector autoregressively to the frames of a video, propagating representations through time to track objects. To allow representations to bind consistently to the same object irrespective of its location, we depart from the encoder-only OWL-ViT architecture: We decouple object representations from the image grid by adding a transformer decoder, as is common practice in end-to-end closed-world detectors and trackers [6, 25, 36]. The decoder maps from image-centric encoder tokens to object-centric “slots”. In-

[†]Work done while at Google.

Correspondence: heigold@google.com, tkipf@google.com

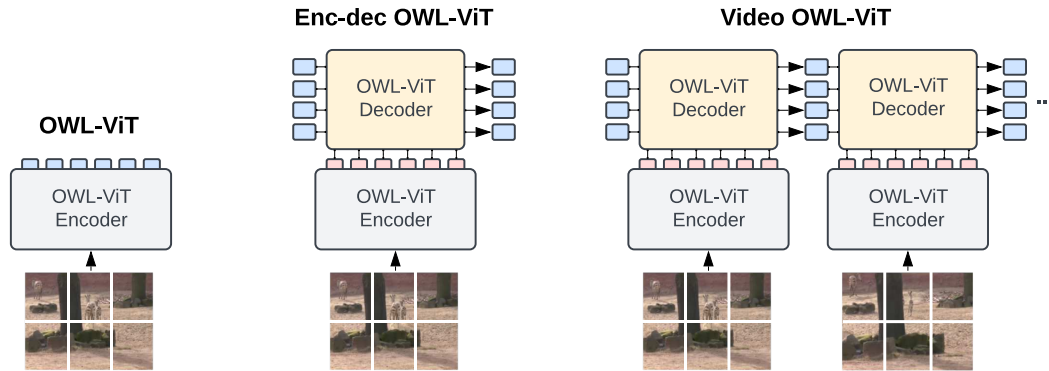


Figure 1: **Model overview.** Our starting point is **OWL-ViT** [26] (left), which uses an encoder-only Vision Transformer (ViT) [9] architecture for simple transfer from image-text pretraining to open-world detection: encoder tokens, arranged on the image grid, are used directly as object queries for detection. To transfer to temporal tasks without requiring frame-to-frame matching, we first develop a model variant inspired by DETR [6] that decouples object queries from the image grid (**Enc-dec OWL-ViT**, middle) by training a lightweight Transformer decoder on top of the ViT encoder while maintaining open-world detection capabilities. Finally, **Video OWL-ViT** (right) simply connects the output of Enc-dec OWL-ViT applied to one frame to the next frame by using the predicted object queries as queries for the OWL-ViT Decoder of the next time step, without the need for any matching.

formation can then be carried through time by using the object slot representations from one frame as decoder queries on the next frame.

The object-centric decoder queries allow the model to *learn* temporally consistent representations end-to-end from video data. This distinguishes our approach from previous open-world tracking models [24], which applied frozen detectors frame-by-frame and used heuristics to link detections through time.

We provide a recipe for incorporating a decoder into OWL-ViT and for fine-tuning the resulting model on video data without losing its open-world detection capabilities. We call the resulting model *Video OWL-ViT*. We demonstrate strong performance on a challenging open-world video localization and tracking task, TAO-OW [24], even for classes that were not seen during video training. We further demonstrate the zero-shot open-world generalization capabilities of Video OWL-ViT on a different dataset, YT-VIS [35], that was not used for training.

2. Related Work

Open-Vocabulary Object Recognition Methods using pretraining on large amounts of web data, most notably Contrastive language-image pretraining (CLIP [27]), have recently led to dramatic improvements in open-vocabulary performance on a range of vision tasks. Much research focuses on transferring these open-vocabulary capabilities to downstream tasks such as object detection.

The main challenge is to adapt the pretrained vision-language model to a downstream task in a way that retains

the semantic knowledge and open-vocabulary capabilities acquired during pretraining. Various approaches have been proposed in the case of object detection, such as distillation [12], freezing the backbone [19], or phrase grounding losses [21, 39]. OWL-ViT [26] proposes a simple recipe to directly transfer a vision-language model to detection with minimal modifications. We build on OWL-ViT due to its simplicity and end-to-end architecture.

Multiple Object Tracking (MOT) The prevailing paradigm for MOT is tracking by detection, in which methods first locate the objects of interest and then associate detections across frames in a separate step. SORT [3, 31] combined appearance cues with motion estimation and association optimization. Many works have targeted better motion estimation [2, 11, 14, 33], while the across-frame association step can also be learned [5, 28]. Recently, many works [25, 38] explore Transformer [30] encoder and decoder architectures to perform end-to-end learning of object tracking. Our work aims to extend this trend towards end-to-end open-world tracking.

Open-World Tracking Recently, open-world tracking has been introduced as an extension of the MOT task [24, 22]. In open-world tracking, the goal is to track not just known object categories, but all objects, including those from categories for which no annotated instances were seen during training. The ability to track unknown objects is critical for safety in applications such as autonomous driving. Since models have only recently become capable of

strong open-vocabulary detection, few open-world tracking models exist. Baselines for open-world tracking obtain per-frame proposals from open-world detectors and use heuristics to link proposals over time [24, 23]. In concurrent work, OVTrack [23], Li et al. similar to us propose data augmentation strategies to benefit from static images for learning temporal association, but they still rely on a tracking-by-detection heuristic to link detections over time. Here, we propose an end-to-end trainable architecture for open-world tracking.

3. Method

Our starting point is OWL-ViT [26] (Sec. 3.1), a simple yet effective Vision Transformer [9] model for open-world object detection in images. To generalize OWL-ViT to video tasks, we first develop an encoder-decoder variant of the model (Enc-dec OWL-ViT, Sec. 3.2) to decouple object queries from the image grid. This allows for a straightforward extension to video tasks, described in Sec. 3.3 (Video OWL-ViT). For an overview of our method, see Figure 1.

3.1. Background: OWL-ViT

We briefly review the OWL-ViT model, which we use as our detection backbone. OWL-ViT consists of a standard Vision Transformer [9] image encoder and an architecturally similar text encoder. The encoders are contrastively pretrained on a large datasets of image/text pairs [27]. After pretraining, the model is transferred to detection by adding lightweight classification and box regression heads that predict class embeddings and box coordinates directly from the image encoder output tokens. For open-vocabulary classification, similarities are computed as the inner product between class embeddings derived from image patches and text embeddings of label names (provided text prompts). These similarities act as classification logits, which are shifted and scaled (using learned parameters), and trained using a sigmoid focal loss on standard detection datasets. To compute the loss, the Hungarian algorithm is used to match predictions to ground-truth targets. Unless otherwise noted, we use the CLIP-based L/14 variant of OWL-ViT. For detection training, we use the same data (Objects365 [29] and Visual Genome [18]) and augmentations as in the original paper [26]. Next, we describe how we adapt the model to tracking.

3.2. Enc-Dec OWL-ViT

To enable temporally consistent representations that can track objects across frames, we must allow the model to decouple object representations from specific image tokens. We achieve this by inserting a Transformer decoder between the encoder and the object heads (i.e. readout heads for box and class prediction), similar to the original DETR architecture [6]. We use the decoder queries as “slots” that carry

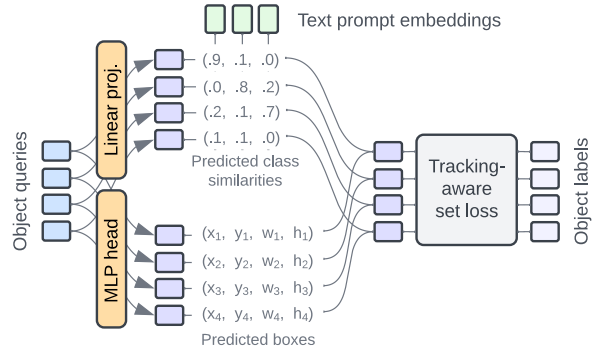


Figure 2: **Prediction heads and loss.** Video OWL-ViT predicts object bounding boxes and semantic embeddings from its set of object queries at every time step. Predicted semantic embeddings are compared against a set of text prompt embeddings to obtain predicted class similarity scores. For training, a tracking-aware matching-based set loss is used to compare predicted bounding boxes and class similarities against ground-truth object labels at every time step.

object representations recurrently from one timestep to the next. We refer to this architecture as Enc-dec OWL-ViT.

Given that OWL-ViT was originally designed as an encoder-only architecture, an important question is whether the open-vocabulary performance of the model can be maintained when adding a decoder. In Section 4.1, we show that the Enc-dec architecture retains most of the performance of the encoder-only architecture.

3.3. Video OWL-ViT

We adapt the Enc-dec OWL-ViT model to video simply by repeatedly applying the model to frames of a video, one frame at a time, while using the predicted object queries of the previous frame as query initialization for the next time step. We further introduce video-specific data augmentation to make adaptation to video more label-efficient.

Architecture The architecture of Video OWL-ViT follows that of our image-based Enc-dec OWL-ViT model variant: To process a video, we apply Enc-dec OWL-ViT iteratively over the individual video frames. On the first video frame, we initialize object queries using simple learned feature vectors. We then directly use the object queries predicted by the OWL-ViT decoder for this frame as decoder queries for the next time step. Model parameters are shared between time steps, i.e. the model is applied recurrently.

We further carry over the box prediction and classification heads from upstream image-based pre-training. For fine-tuning, the image encoder and the text encoder are frozen (for efficiency), and only the box prediction and classification heads, and the transformer decoder, are updated.

Finally, to obtain “objectness” scores for each instance at every frame, we take the maximum predicted classification logit across all classes.

Training and loss We train Video OWL-ViT using a tracking-aware set prediction loss similar to the tracking loss used in TrackFormer [25]. In contrast to the TrackFormer setup, we train on short video sequences (instead of only pairs of frames). During training, we use Hungarian matching to match predicted object features (boxes and class similarities, see Figure 2) to ground-truth object labels at every time step, starting from the first frame. Once a prediction is matched to a ground-truth track, it stays matched for the remainder of the training video clip, i.e. we only perform matching for previously unmatched objects.

Video OWL-ViT predicts both class similarities (compared to text prompt embeddings) and object bounding box coordinates (center and width/height) at every time step. These predictions are used both for computing the matching cost as well as the final loss. Like OWL-ViT, we use focal sigmoid cross-entropy [41] instead of the usual softmax cross-entropy, which is better suited for open-world detection, for both classification loss and matching cost. We use the same box prediction loss as in DETR [6], i.e. a weighted sum of L1 and generalized IoU loss terms.

Object queries that are unmatched or matched to an empty track (with no object present in the ground-truth track at this time step) are trained to predict low class similarity scores for all provided text prompts, i.e. all text prompts are treated as *negative examples* in the loss. For matched object queries, we train the model to predict high class similarity scores for all ground-truth text prompts describing the class of that object and a low class similarity score for all other prompts, i.e. ground-truth text prompts are treated as *positive examples* in the loss.

Note that, when applied to a single frame, our tracking-aware set prediction loss exactly matches the loss used in OWL-ViT [26].

Augmentations As we primarily rely on upstream image-based pre-training and assume limited availability of video data, we make use of several video-specific data augmentation techniques.

First, we create *pseudo videos* from images by aggressive scaling and cropping of an image, similar to prior work [25, 32, 22], but with a linear motion model to generate video clips longer than 2 frames. This simulates a slowly moving camera observing a static scene (see Figure 3, top row).

For video data, we perform temporally-consistent random re-scaling and cropping of entire video clips. To augment motion, we sub-sample the frame rate by randomly selecting frames while preserving temporal order.

We further find that extending the image mosaic augmentation used in OWL-ViT to video in the form of a *temporal video mosaic* proves beneficial for generalization. This video mosaic augmentation is similar to VideoMix proposed in [37] for video classification, and combines multiple video clips into a single clip by means of a scene cut (temporal video mosaic; Figure 3, bottom row).

Finally, for datasets like TAO that are only annotated at 1 FPS, we propagate annotations to non-annotated frames by linear interpolation to make use of all frames for training.

4. Experiments

Our experiments aim to answer the following main research questions:

- Does OWL-ViT maintain open-world detection performance when a decoder is added?
- How well does the resulting model transfer to video (*Video OWL-ViT*)?
- Do open-world capabilities from image pre-training carry over to tracking of unseen instances in video?

We further perform detailed ablations to investigate individual components of our model.

4.1. Encoder-Decoder Detection Evaluation

For our object detection backbone, we build on the CLIP-based OWL-ViT model [26] with a ViT-L/14 image encoder. To adapt this model for recurrent autoregressive application during tracking, we first investigate how detection performance is affected when adding a Transformer decoder between the image encoder and the prediction heads of the OWL-ViT architecture.

Starting from the architecture in the original paper [26], we pre-train the model on detection data (Objects365 and Visual Genome) for 70,000 steps with a batch size of 256 as described in the paper. We then add a Transformer decoder (architecture as in DETR [6]; 6 layers, 8 heads, 4096 MLP dim, 1024 QKV dim, 100 decoder queries). To evaluate the effect of adding a decoder on detection performance, we train the Enc-dec model for an additional 70,000 steps using the same training data and schedule. We keep the image size at 672×672 for all experiments.

We find that detection performance on unseen LVIS “rare” classes ($AP_{\text{rare}}^{\text{LVIS}}$) of the decoder model reaches 28.9, close to the 31.8 achieved by the encoder-only model. A small drop in performance is expected, given that the Enc-dec model outputs a significantly smaller number of object predictions (100 instead of 2304 for the encoder-only model). These results suggest that adding a decoder is a viable approach for adapting OWL-ViT to video.




Figure 3: **Data augmentation.** Top row: Example of camera sweep with original image (left) and pseudo video (right). Bottom row: Example of temporal mosaic (concatenation of two random clips).

4.2. Open-world Video Localization

Video training details As in the previous section, we start from an Enc-dec OWL-ViT model in which the encoder is pretrained on detection and the decoder is randomly initialized. We found that pre-training the whole Enc-dec model on detection provided no advantage, likely because the Transformer decoder fulfils a different role in video compared to image data: it has to predict the objects in the current frame, but also produce suitable queries for the next frame to enable tracking. To simplify the experimental pipeline, we thus start video training directly from the original pre-trained OWL-ViT model. Video OWL-ViT uses a set of 196 learnable object queries.

We train for 100k steps with a batch size of 32 first on pseudo-videos only, followed by another 100k steps on a mixture of pseudo-videos and real videos from TAO-OW. Pseudo-videos are obtained by augmenting LVIS and Objects365 images as described in Section 3.3. Image resolution is kept at 672×672 by resizing images while preserving aspect ratios and padding as needed.

Evaluation dataset We focus our evaluation on the recent TAO Open World (TAO-OW) [24] dataset: it is based on the Track Any Object (TAO) video dataset [8], but specifically tests for open-world detection and tracking performance by (1) restricting the object classes for which labels are provided during training, (2) providing out-of-distribution validation and test sets to evaluate on unseen object classes, and (3) introducing a metric that accounts for incomplete object annotations (e.g. due to filtering of known classes).

Since we are investigating the transfer of models pretrained on large-scale web data, we restrict the set of “seen” object classes only during the final *video* training stage, not during upstream image-level pre-training. Our only source of natural video training data is the TAO-OW training set, containing 500 videos.

As the test set annotations and evaluation server for TAO-OW are not yet publicly available, we evaluate our

model on the provided validation set, containing 988 videos, and perform model selection and hyperparameter tuning based on the training set.

Evaluation metric For evaluation and model comparison, we use the Open-World Tracking Accuracy (OWTA) metric introduced by Liu et al. [24], the standard metric for TAO-OW. OWTA is defined as the geometric mean of Detection Recall (D. Re.) and frame-to-frame Association Accuracy (A. Acc.), integrated over multiple localization thresholds. Importantly, D. Re. ignores false positive detections and is thus suitable for evaluating in the incompletely annotated open-world setting of TAO-OW. To avoid cheating the metric by producing a myriad of detection proposals, OWTA enforces a non-overlapping prediction constraint by requiring models to produce non-overlapping segmentation masks, so that every pixel is assigned to at most one detected object, or to the background.

To enforce this constraint, we train a separate segmentation head on top of the frozen OWL-ViT model (for details see appendix). Specifically, for each frame and pixel, among all instances whose mask overlaps with that pixel, we keep the instance with the highest score. Pixels belonging to other instances are removed from the segmentation mask. The remaining instance mask is then converted back to box coordinates for evaluation.

To account for the observation that predicted objectness scores can be miscalibrated for small objects and short tracks, we introduce two simple heuristics. To rank instances for overlap removal, we use the objectness score divided by the box area. This heuristic accounts for the observation that smaller objects tend to have lower objectness scores. To suppress false positives for short tracks, we mark parts of tracks as background if they have significantly lower objectness scores than the maximum objectness observed in the track. We mark any detection as background if its objectness score is lower than a pre-defined fraction of the maximum objectness score along the track.

Table 1: **TAO-OW** open world tracking. Baseline results from Liu et al. [24]. Rows labeled “w/o constraint” do not use the non-overlapping constraint during evaluation. Gray indicates results for classes that were seen during video training. AOA [10] performs video training on both known and unknown classes of TAO-OW and is thus not directly comparable (results marked in gray). All metrics in %. Best numbers highlighted in bold (excl. results in gray).

Model	Known			Unknown		
	OWTA↑	D. Re. ↑	A. Acc. ↑	OWTA↑	D. Re. ↑	A. Acc. ↑
SORT [3]	46.6	67.4	33.7	33.9	43.4	30.3
Tracktor [2]	57.9	80.2	42.6	22.8	54.0	10.0
OWTB [24]	60.2	77.2	47.4	39.2	46.9	34.5
OWL-ViT tracking-by-detection	37.7	36.1	40.1	31.0	32.0	31.8
Video OWL-ViT (Ours)	59.0	69.0	51.5	45.4	53.4	40.5
AOA (w/o constraint) [10]	52.8	72.5	39.1	49.7	74.7	33.4
SORT-TAO (w/o constraint) [8]	54.2	74.0	40.6	39.9	68.8	24.1
OWTB (w/o constraint) [24]	60.8	82.0	45.5	42.4	58.9	31.5
OWL-ViT tracking-by-det. (w/o constr.)	44.5	45.5	43.9	42.2	51.5	35.4
Video OWL-ViT (w/o constraint) (Ours)	56.6	73.2	44.6	47.3	62.3	37.2



Figure 4: Qualitative example for Video OWL-ViT detection and tracking of multiple instances on the **TAO-OW** validation set. This example includes different instances of the same kind (giraffes) and partial occlusion. Video OWL-ViT can recover from occlusion, despite being trained on short video clips.

Main results Our method makes two main contributions over heuristic tracking-by-detection baselines such as the Open World Tracking Baseline (OWTB) [24]: (1) Our model is end-to-end trainable on video and can therefore learn temporal consistency directly from data. (2) Our method transfers open-world semantic knowledge from image pretraining to tracking. Our results on TAO-OW show that both contributions translate to improved performance.

The advantage of end-to-end training is apparent from the association accuracy (A. Acc., Table 1), which measures the accuracy of associating detections across frames. Our model outperforms all baselines on A. Acc. on both known and unknown classes. We hypothesize that learning temporal associations from data, rather than matching single-frame detections heuristically, reduces tracking error accumulation. An end-to-end method like ours also promises to improve further when more video training data is available. We provide qualitative results in Figure 4 and in video format in the supplementary material.

The transfer of pretrained knowledge is apparent from the performance of Video OWL-ViT on unknown classes, i.e. classes for which no video training data is available. Video OWL-ViT outperforms OWTB on unknown classes on the OWTA metric by a wide margin, show-

ing that image-level open-world knowledge can be transferred to video with minimal video-specific training data. While we observe variance of OWTA scores on unknown classes between repeated training runs of approx. 1% (absolute) OWTA, Video OWL-ViT still reliably outperforms the baselines.

On known classes, Video OWL-ViT performs similarly to OWTB on the overall OWTA metric. However, while the baseline achieves its performance primarily through high detection recall (D. Re. in Table 1, i.e. single-frame performance) which compensates for its low association accuracy, the end-to-end Video OWL-ViT has more balanced detection and association performance.

We compare models both with and without the non-overlapping constraint of the OWTA metric. For evaluations without the constraint, we note that our model uses significantly fewer tracks, a total of 196, compared to the most competitive baseline (OWTB), which uses the top 1000 object proposals per frame, thus placing our model at a disadvantage in this constraint-free evaluation: a higher proposal budget makes it easier to score well on detection recall. Despite this disadvantage, Video OWL-ViT compares favorably against the baselines reported previously on this benchmark.

Table 2: Zero-shot transfer to **YT-VIS** open world tracking with Video OWL-ViT. All metrics in %.

Model	Known			Unknown		
	OWTA	D. Re.	A. Acc.	OWTA	D. Re.	A. Acc.
SORT [3]	43.0	40.5	48.7	45.0	48.1	44.9
OWL-ViT TbD	71.1	72.3	70.8	71.9	73.5	71.5
Video OWL-ViT	79.4	82.9	76.4	81.2	81.8	81.0

To isolate the benefit of the Video OWL-ViT decoder we also compare to a simple tracking-by-detection baseline based on appearance matching, depicted as *OWL-ViT tracking-by-detection* in Table 1. Specifically, our baseline performs optimal bipartite matching using the cosine similarity of the OWL-ViT embeddings from 300 proposals in the current and previous frame using the Jonker-Volgenant algorithm [16]. These matches are extended to form tracks across the entire sequence. For unknown classes, this baseline performs comparably to the OWTB without constraint, with a bias towards better association. Video OWL-ViT shows a substantial gain over this baseline, highlighting the benefit of end-to-end training enabled by our model. In terms of computational cost, we find that our temporal decoder adds approx. 6 ms to the to the inference time, which is similar to the bipartite matching in our tracking-by-detection baseline (4–11 ms).

Zero-shot transfer We study generalization capabilities of Video OWL-ViT by evaluating the model without any further training on the YT-VIS 2019 dataset [35], which contains annotations at significantly higher frame rates. The annotations for the official validation and test datasets of YT-VIS [35] are not published. For this reason we use validation and test splits taken from the training data with 200 videos each.¹ We report results on the test split. In this zero-shot transfer setup, we measure the same metrics as used for TAO-OW, but evaluate at a higher frame rate of 6 FPS. We distinguish between known and unknown classes based on which classes were available in the TAO-OW training set.

As can be seen in Table 2 (quantitative results) and Figure 6 (qualitative results), Video OWL-ViT shows strong transfer even to unseen classes in YT-VIS, despite evaluation at a significantly higher frame-rate. Compared to the *OWL-ViT tracking-by-detection (TbD)* and *SORT* [3] baselines, Video OWL-ViT generally shows better transfer to both known and unknown classes, and demonstrates improved association accuracy.

¹https://www.tensorflow.org/datasets/catalog/youtube_vis; Version: '480_640_only_frames_with_labels_train_split'




Figure 5: Visualization of box predictions on all images from the COCO 2017 val set, similar to Figure 7 in DETR [7]. Each dot indicates a box center. Color indicates size/shape as in [7], i.e. green represents small boxes while red/blue represents large horizontal/vertical boxes, respectively. Rows show the first three decoder slots (not cherry-picked). Columns show frames of synthetic videos created by sliding a cropped view over the image from left to right. Boxes with objectness > 0.2 are shown. On the first frame, like DETR, slots specialize to certain areas. However, on subsequent frames, they track appearance rather than remaining at the same image coordinates. The model learns to reserve some slots (e.g. Slot 2) for late-appearing objects.

Temporal association analysis To isolate temporal association (tracking) performance, we assess our method in the single-object-tracking (SOT) setting by initializing the track from the initial ground-truth bounding box on YT-VIS. We compare Video OWL-ViT to the tracking-by-detection (TbD) baseline in terms of 3D IoU ($\sum_t p_t \cap g_t / \sum_t p_t \cup g_t$), see Figure 7.




Figure 7: SOT.

Performance of the baseline decays faster than Video OWL-ViT performance, which indicates superior temporal association performance afforded by the end-to-end tracking architecture of Video OWL-ViT.

Location specificity of object queries To further illustrate the temporal association capability of Video OWL-ViT, we perform an analysis of the location specificity of object queries over time. In Figure 5, we visualize box predictions by object queries (slots).

To disentangle whether slots attend to image space or object appearance, we create synthetic videos by sliding a cropped view over the image. We find that slots move with the image (i.e. with appearance), rather than remaining fixed at a specific location. Additionally, the model learns to reserve some slots for late-appearing objects.



Figure 6: Zero-shot transfer to **YT-VIS** open world tracking with Video OWL-ViT. Examples for a known class (“dog”, left) and an unknown class (“parrot”, right). The 6-second videos are visualized by three equidistant frames each.

Table 3: Clip length for training on **TAO-OW** open world tracking (unknown classes).

Clip length	OWTA \uparrow	D. Re. \uparrow	A. Acc. \uparrow
2 frames	35.2	40.7	32.6
4 frames	45.3	53.2	40.5
6 frames	45.4	53.4	40.5

4.3. Understanding the Challenges of Open-World Video Modeling

In the previous sections, we have presented an end-to-end trainable open-world video localization model. We now analyze the error modes of our model and describe how we address them, with a particular focus on end-to-end learning of tracking dynamics from limited video data. Overall, we find that a realistic distribution of object dynamics in the training data, as well as careful modeling of object presence, are important factors affecting model performance.

Training clip length A common approach for video models [25] is to train on the shortest possible clip length (i.e. two frames), which is memory and compute efficient. However, training on longer clips allows for more realistic learning of object dynamics, including motion, appearance/disappearance, occlusion, recovery from tracking errors, and long-term dependencies. This may be especially important for a model that learns object dynamics directly from the data. We empirically confirm this in Table 3. We find that training on 4-frame clips significantly improves performance compared to training on 2-frame clips. Beyond 4 frames, performance quickly saturates.

Training on pseudo-videos TAO-OW contains only 500 videos in the training set, which poses the risk of overfitting on the small number of represented object classes. To maintain performance on object classes not represented in the video data, we leverage more abundant image data by generating pseudo-videos from still images (Section 3.3). The image data includes the training data from LVIS [13] (approx. 100k images) and Objects365 [29] (approx. 600k images). As shown in Table 4, training on pseudo-videos

Table 4: Comparison between fine-tuning OWL-ViT with pseudo-videos or real videos. Results on **TAO-OW** open world tracking; all metrics in %.

Supervision	Known OWTA	Unknown OWTA
Real videos only	54.6	33.9
Pseudo-videos only	56.9	45.9
Real + pseudo-videos	59.0	45.4

Table 5: Score calibration and temporal mosaic data augmentation improve performance on medium and short video tracks. Results on **TAO-OW** open world tracking; all metrics in %. The track length buckets are *short* (shorter than 3 seconds), *medium* (between 3 and 10 seconds), and *long* (longer than 10 seconds), with the following distributions: 11%/25%/64% (known) and 4%/15%/81% (unknown).

Score calibration	Temporal mosaic	Known OWTA			Unknown OWTA		
		Short	Med.	Long	Short	Med.	Long
×	×	15.9	20.2	58.6	12.5	16.0	45.7
✓	×	26.1	30.6	59.4	22.2	26.2	46.8
✓	✓	31.3	32.5	60.6	22.8	27.6	45.9

yields a substantial improvement in OWTA on unknown classes compared to training on real videos. Performance is also improved for known classes, which is notable given that pseudo-videos do not have realistic motion dynamics and underscores the importance of sufficient training data. Combining real and pseudo-videos further improves performance on known classes, but not on unknown classes. The fact that training on pseudo-videos alone performs similar on unknown classes compared to training on real videos illustrates how challenging the small amount of video data is for open-world performance.

Performance on short tracks We find that association accuracy is significantly lower for short than for long tracks (Table 5). One reason for lower performance on short tracks may be that the objectness score is not a well-calibrated indicator for deciding whether an instance is an object or background. Since short tracks contain more frames during which the object is not visible (i.e. “background”), they are

Table 6: Evaluation at different frame rates (FPS). Results on **TAO-OW** open world tracking; all metrics in %.

FPS	Known OWTA	Unknown OWTA
1	59.0	45.4
2	59.0	45.4
4	61.1	46.1
8	60.5	45.5

disproportionately affected by poor objectness calibration. To mitigate this effect, we use a simple heuristic to mark parts of tracks as “background” (i.e. no object) that significantly differ from the maximum objectness score o_{\max} across the track. We find that a simple per-track threshold of $0.3 \cdot o_{\max}$, below which we mark detections as background, suffices to significantly improve performance on shorter tracks (“Score calibration” in Table 5).

A second reason for poor short-track performance may be that the training data contains fewer short than long tracks. To address this imbalance, we create artificially shortened tracks by concatenating short clips from different videos into longer sequences (“Temporal mosaic” in Table 5).

We find that both score calibration and video mosaic data augmentation can significantly improve performance on short and medium-length tracks. The effect of score calibration is especially large, despite using a simple heuristic for recalibration. This suggests that a more sophisticated learnable and directly supervised presence indicator may lead to further improvements.

Inference frame rate While TAO is annotated at 1 FPS, the videos come at a higher frame rate. This allows us to operate the model at higher frame rates, using also intermediate frames to compute the predictions. For the metrics, only the predictions associated with annotated frames are kept. According to Liu et al. [24], using intermediate frames generally improves known accuracy, but harms unknown accuracy for the tracking-by-detection OWTB baseline. In contrast, Video OWL-ViT improves known accuracy without degrading unknown accuracy (Table 6), demonstrating the benefit of *learning* frame-to-frame association from data instead of relying on a matching heuristic. This helps close the known-accuracy gap between OWTB and Video OWL-ViT (at the expense of increased compute).

5. Conclusion

We introduced Video OWL-ViT, a simple end-to-end model for open-world localization and tracking in video. Video OWL-ViT builds on the open-world detection recipe of OWL-ViT and transfers an image-text pre-trained vi-

sion transformer model to video via fine-tuning and tracking-specific data augmentation. To enable temporally-consistent localization of objects, we add a decoder to OWL-ViT to decouple object queries from the input pixel grid and train using a tracking-aware set prediction loss.

Video OWL-ViT achieves performance competitive with tracking-by-detection baselines on the open-world TAO-OW benchmark, while presenting several advantages, such as matching-free tracking at test time and consistent performance even at higher frame rates for long videos. Our analyses of performance limitations of the model suggest that improving the amount and quality of video training data, and the modeling of object presence are promising future directions.

Acknowledgements

We would like to thank Alexey Dosovitskiy for providing detailed feedback on an earlier draft of the paper.

References

- [1] Rodrigo Benenson, Stefan Popov, and Vittorio Ferrari. Large-scale interactive object segmentation with human annotators. In *CVPR*, 2019. 11
- [2] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking without bells and whistles. In *ICCV*, 2019. 2, 6
- [3] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Uppcroft. Simple online and realtime tracking. In *ICIP*, 2016. 2, 6, 7
- [4] Vighnesh Birodkar, Zhichao Lu, Siyang Li, Vivek Rathod, and Jonathan Huang. The surprising impact of mask-head architecture on novel class segmentation. In *ICCV*, 2021. 11
- [5] Guillem Brasó and Laura Leal-Taixé. Learning a neural solver for multiple object tracking. In *CVPR*, 2020. 2
- [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 1, 2, 3, 4
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 7
- [8] Achal Dave, Tarasha Khurana, Pavel Tokmakov, Cordelia Schmid, and Deva Ramanan. TAO: A Large-Scale Benchmark for Tracking Any Object. In *ECCV*, 2020. 5, 6
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2, 3
- [10] Fei Du, Boao Xu, Jiasheng Tang, Yuqi Zhang, F. Wang, and Hao Li. 1st place solution to ECCV-TAO-2020: Detect and represent any object for tracking. *arXiv preprint arXiv:2101.08040*, 2021. 6
- [11] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *ICCV*, 2017. 2

- [12] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *ICLR*, 2022. 1, 2
- [13] Agrim Gupta, Piotr Dollár, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 8
- [14] David Held, Sebastian Thrun, and Silvio Savarese. Learning to track at 100 FPS with deep regression networks. In *ECCV*, 2016. 2
- [15] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016. 11
- [16] Roy Jonker and Ton Volgenant. A shortest augmenting path algorithm for dense and sparse linear assignment problems. In *DGOR/NSOR: Papers of the 16th Annual Meeting of DGOR in Cooperation with NSOR*, pages 622–622. Springer, 1988. 7
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 11
- [18] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123:32–73, 2017. 3
- [19] Weicheng Kuo, Yin Cui, Xiuye Gu, AJ Piergiovanni, and Anelia Angelova. Open-vocabulary object detection upon frozen vision and language models. In *ICLR*, 2023. 1, 2
- [20] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 128:1956–1981, 2020. 11
- [21] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In *CVPR*, 2022. 2
- [22] Siyuan Li, Martin Danelljan, Henghui Ding, Thomas E. Huang, and Fisher Yu. Tracking every thing in the wild. In *ECCV*, 2022. 2, 4
- [23] Siyuan Li, Tobias Fischer, Lei Ke, Henghui Ding, Martin Danelljan, and Fisher Yu. OVTrack: Open-Vocabulary Multiple Object Tracking. In *CVPR*, 2023. 3
- [24] Yang Liu, Idil Esen Zulfikar, Jonathon Luiten, Achal Dave, Aljosa Osep, Deva Ramanan, Bastian Leibe, and Laura Leal-Taixé. Opening up open world tracking. In *CVPR*, 2021. 2, 3, 5, 6, 9
- [25] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixé, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *CVPR*, 2022. 1, 2, 4, 8
- [26] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, Xiao Wang, Xiaohua Zhai, Thomas Kipf, and Neil Houlsby. Simple open-vocabulary object detection with vision transformers. In *ECCV*, 2022. 1, 2, 3, 4, 11
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 2, 3
- [28] Samuel Schulter, Paul Vernaza, Wongun Choi, and Manmohan Chandraker. Deep network flow for multi-object tracking. In *CVPR*, 2017. 2
- [29] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *ICCV*, 2019. 3, 8
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2
- [31] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *ICIP*, 2017. 2
- [32] Sanghyun Woo, Kwanyong Park, Seoung Wug Oh, In So Kweon, and Joon-Young Lee. Bridging images and videos: A simple learning framework for large vocabulary video object detection. In *ECCV*, 2022. 4
- [33] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *ECCV*, 2018. 2
- [34] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *ICML*, 2020. 11
- [35] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *ICCV*, 2019. 2, 7
- [36] Zhuyu Yao, Jiangbo Ai, Boxun Li, and Chi Zhang. Efficient DETR: Improving end-to-end object detector with dense prior. *arXiv preprint arXiv:2104.01318*, 2021. 1
- [37] Sangdoon Yun, Seong Joon Oh, Byeongho Heo, Dongyoon Han, and Jinhyung Kim. Videomix: Rethinking data augmentation for video classification. *arXiv preprint arXiv:2012.03457*, 2020. 4
- [38] Fangao Zeng, Bin Dong, Yuang Zhang, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. MOTR: End-to-end multiple-object tracking with transformer. In *ECCV*, 2022. 2
- [39] Haotian Zhang, Pengchuan Zhang, Xiaowei Hu, Yen-Chun Chen, Liunian Harold Li, Xiyang Dai, Lijuan Wang, Lu Yuan, Jenq-Neng Hwang, and Jianfeng Gao. GLIPv2: Unifying localization and vision-language understanding. In *NeurIPS*, 2022. 2
- [40] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, and Jianfeng Gao. RegionCLIP: Region-based language-image pretraining. In *CVPR*, 2022. 1
- [41] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: Deformable transformers for end-to-end object detection. In *ICLR*, 2021. 4