

A Distractor-Aware Memory for Visual Object Tracking with SAM2

Jovana Videnovic*, Alan Lukezic*, Matej Kristan

Faculty of Computer and Information Science, University of Ljubljana, Slovenia

jv8043@student.uni-lj.si, {alan.lukezic, matej.kristan}@fri.uni-lj.si

Abstract

Memory-based trackers are video object segmentation methods that form the target model by concatenating recently tracked frames into a memory buffer and localize the target by attending the current image to the buffered frames. While already achieving top performance on many benchmarks, it was the recent release of SAM2 that placed memory-based trackers into focus of the visual object tracking community. Nevertheless, modern trackers still struggle in the presence of distractors. We argue that a more sophisticated memory model is required, and propose a new distractor-aware memory model for SAM2 and an introspection-based update strategy that jointly addresses the segmentation accuracy as well as tracking robustness. The resulting tracker is denoted as SAM2.1++. We also propose a new distractor-distilled DiDi dataset to study the distractor problem better. SAM2.1++ outperforms SAM2.1 and related SAM memory extensions on seven benchmarks and sets a solid new state-of-the-art on six of them. The code and the new dataset will be available on <https://github.com/jovanavidenovic/DAM4SAM>.

1. Introduction

General visual object tracking is a classical computer vision problem that considers the localization of an arbitrary target in the video, given a single supervised training example in the first frame. The major source of tracking failures are so-called distractors, i.e., image regions that are difficult to distinguish from the tracked object, given the available target model (see Figure 1). These can be nearby objects similar to the tracked target (*external distractors*) or similar regions on the object when tracking only a part of the object (*internal distractors*). When the target leaves and re-enters the field of view, the external distractors become particularly challenging.

Various approaches have been proposed to reduce the vi-

* The authors contributed equally.

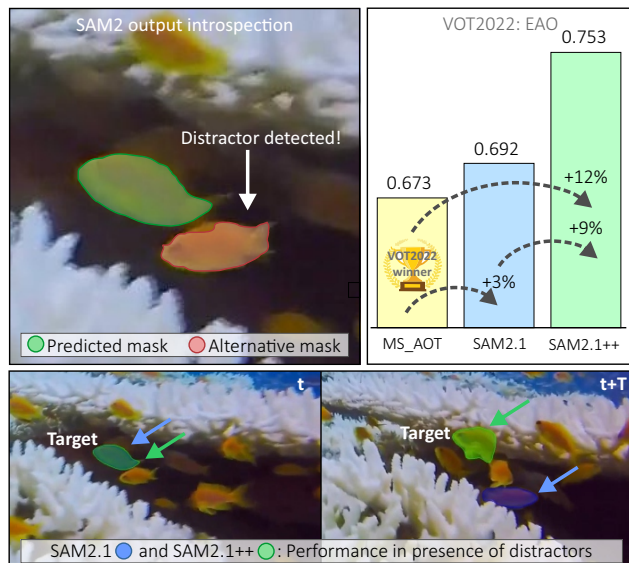


Figure 1. SAM2.1++ distractor-aware memory (DAM) update is triggered by the divergence between the predicted and the alternative masks (top left). This resolves the visual ambiguity and increases tracking robustness (bottom). DAM leads to a significant performance boost, setting a new sota on VOT2022 (top-right).

sual ambiguity caused by distractors. These include learning discriminative features [2, 5, 6, 10, 43] or explicitly modeling the foreground-background by dedicated modules [9, 28, 29, 50]. An emerging paradigm, already positioned at the top of the major benchmarks [23–25], are memory-based frameworks, which localize the target by pixel association with the past tracked frames [8, 47, 51].

The memory-based methods construct the target model by concatenating sequences of entire images with the segmented target, thus implicitly encoding the present distractors. [51] argue that the visual redundancy in the large memory leads to reduced localization capability due to the nature of cross-attention. They show that limiting the memory to the most recent frames and temporally time-stamping them in fact improves tracking. This paradigm was further verified by the recent tracking foun-

dation model SAM2 [36], which sets a solid state-of-the-art across several video segmentation and tracking benchmarks [13, 18, 25, 35, 41].

We argue that while the recent target appearances in the memory are required for accurate segmentation, another type of memory is required to distinguish the target from challenging distractors. To support this claim, we propose a new distractor-aware memory (DAM) and update mechanism for SAM2. The new memory is divided by its tracking functionality into two parts: the recent appearances memory (RAM) and distractor-resolving memory (DRM). While RAM contains the recent target appearances sampled at regular intervals, DRM contains anchor frames that help discriminate the target from critical distractors. A novel DRM updating mechanism is proposed that exploits the information of SAM2 output, which has been ignored so far by the tracking research.

In addition, we observe that standard benchmarks contain many sequences, which are no longer considered challenging by modern standards. The high performance on these sequences overwhelms the total score, saturates the benchmarks, and does not properly expose the tracking advances. To address this, we semi-automatically distill several benchmarks into a distractor-distilled tracking dataset (DiDi).

In summary, our main contribution is the new distractor-aware memory DAM for SAM2 and the updating strategy, resulting in SAM2++. To the best of our knowledge, this is the first memory formulation that divides and updates the memory with respect to its function in tracking. Our secondary contribution is the new DiDi dataset that more clearly exposes tracking advances in the presence of distractors. With no additional training, SAM2.1++ substantially outperforms SAM2.1 in robustness on several standard bounding box and segmentation tracking benchmarks, including the new DiDi dataset, and sets a new state-of-the-art in visual object tracking.

2. Related Work

Transformers are currently the dominant methodology in visual object tracking [23] and can be largely categorized into classification- and regression-based [5], corner prediction-based [43], and sequence-learning-based trackers [6].

The recent top-performing tracking frameworks are inspired by the video object segmentation methods based on memory networks [8, 9, 36, 47, 51]. These methods embed predictions from past frames into memory, therefore extending contextual information beyond just the initial or the previous frame. The attention mechanism is typically used to link frame representations stored in the memory with features extracted from the current frame. In the initial methods like [47], the arriving frames were continually added to the memory. This led to theoretically unbounded increase

in computational complexity and GPU memory.

This issue was addressed in [8, 9] by using multiple memory storages and efficient compression schemes to capture different temporal contexts, enhancing performance on long-term videos. Alternatively, [51] proposed to restrict the memory to the most recent frames with temporal stamping, which led to improved localization. The principle of the restricted memory is followed by the SAM2 [36] foundation model, which stores last 6 frames and the initial frame in the memory. Recently, SAM2Long [14] proposed a training-free method to enhance the performance of SAM2 [36] on long-term sequences by determining the optimal trajectory from multiple segmentation pathways using a constrained tree search.

Most of the existing tracking methods do not explicitly address tracking in the presence of distractors, even though distractors are a major source of tracking failures. Discriminative (deep) correlation filters [12, 33] are theoretically suitable to handle distractors but are in practice outperformed by the modern transformer-based trackers. However, there have been some recent attempts to address distractors. KeepTrack [32] casts the problem as a multi-target tracking setup, where it identifies target candidates and potential distractors, which are then associated with previously propagated identities using a learned association network. However, the method relies on accurate detection and cannot address internal distractors in practice. In [29] target localization accuracy and robustness are treated as two distinct tasks, which is demonstrated to be beneficial in situations with distractors. Despite explicit distractor handling mechanism, these methods lead to complicated architectures and cannot fully exploit the learning potential of modern frameworks. In contrast, memory-based methods [8, 9, 36, 51] have the capacity to implicitly handle distractors in an elegant way, since they store entire images and apply a learnable localization by segmentation. However, the existing memory management methods are not designed to effectively handle the distractors.

3. Distractor-aware memory for SAM2

This section describes the new DAM memory model for SAM2. Section 3.1 briefly outlines the SAM2 architecture, while the new model is described in Section 3.2.

3.1. SAM2 preliminaries

SAM2 extends the Segment Anything Model (SAM) [20], originally developed for interactive class-agnostic image segmentation, to video segmentation. It consists of four main components: (i) image encoder, (ii) prompt encoder, (iii) memory bank, and (iv) mask decoder.

The image encoder applies ViT Hiera¹ backbone [37]

¹ ViT-L version is used in all our experiments.

to embed the input image. Interactive inputs (e.g., positive/negative clicks) are absorbed by the prompt encoder and used for output mask refinement, however, note that these are not applicable in the general object tracking setup. The memory bank consists of the encoded initialization frame with a user-provided segmentation mask and six recent frames with segmentation masks generated by the tracking output. Temporal encodings are applied to the six recent frames to encode the frame order, while such encoding is not applied to the initialization frame to indicate its unique property of being a single supervised training example and thus serves as a sort of target prior model.

The memory bank transfers pixel-wise labels onto the current image by attending the features in the current frame to all memory frames, producing memory-conditioned features. The features are then decoded by the mask decoder, which predicts three output masks along with the IoU prediction for each. The mask with the highest IoU is chosen as the tracking output.

SAM2 applies a variant of the memory management proposed in [51]. The initialization frame is always kept in the memory, while the six recent frames are updated at every new frame by a first-in-first-out protocol. The memory and the management mechanism are visualized in Figure 2, while the reader is referred to [36] for other details.

3.2. Distractor-aware memory – DAM

Related works [8, 36, 47, 51] have clearly demonstrated the importance of the most recent frames, which are required to address target appearance changes and ensure accurate segmentation. However, a different type of frames is required to prevent drifting in the presence of critical distractors and for reliable target re-detection.

We propose to compose the memory with respect to its function during tracking into (i) *recent appearance memory* (RAM) and (ii) *distractor resolving memory* (DRM). RAM and DRM together form the distractor-aware memory (DAM), visualized in Figure 2. The function of RAM is to ensure segmentation accuracy in the considered frame, thus we design it akin to the current SAM2 [36] memory. It is composed of a FIFO buffer with $\frac{1}{2}N_{\text{DAM}} = 3$ slots, contains the most recent target appearances, and applies temporal encoding to identify temporally more relevant frames for the task.

On the other hand, DRM serves for ensuring tracking robustness and re-detection. It should contain accurately segmented frames with critical recent distractors, including the initialization frame. It is thus composed of a slot reserved for the initialization frame and a FIFO buffer with $\frac{1}{2}N_{\text{DAM}} = 3$ anchor frames updated during tracking². Since the purpose of DRM is to encode critical information

for resolving distractors, it does not apply temporal encoding. Note that the pre-trained SAM2 already contains the building blocks to implement the proposed memory structure.

3.2.1. RAM management protocol

A crucial element of the memory-based methods is the memory management protocol. To efficiently exploit the available memory slots, the memory should not be updated at every frame, since the consecutive frames are highly correlated. In fact, [51] argue that visual redundancy in memory should be avoided in attention-based localization. RAM is thus updated every $\Delta = 5$ frames and includes the most recent frame since it is the most relevant for accurate target segmentation in the considered frame.

SAM2 [36] updates the memory at every frame, including when the target is absent. However, even for a very short occlusion, the memory will quickly fill up with frames without the target, which reduces the target appearance diversity in the model, leading to a reduced segmentation accuracy upon target re-appearance. Furthermore, failing to re-detect the target leads to incorrectly updating the memory by an empty mask, which may cause error accumulation and ultimately re-detection failure. We thus propose to not update RAM when the target is not present, i.e., when the predicted target mask is empty.

3.2.2. DRM management protocol

DRM inherits the initial update rules from RAM, i.e., update only when the target is present and at least $\Delta = 5$ frames have passed since the last update. It considers an additional rule to identify anchor frames containing critical distractors. In particular, drifting to a distractor may be avoided by including a past temporally nearby frame with this distractor accurately segmented as the background. Recall that SAM2 predicts three output masks and selects the one with the highest predicted IoU (Section 3.1), which means we can consider it as a multi-hypothesis prediction model. Our preliminary study showed that in the frames before the failure occurs, SAM2 in fact detects such distractors in the alternative two predicted output masks (see Figure 1). We thus propose a simple anchor frame detection mechanism based on determining hypothesis divergence between the output and alternative masks.

A bounding box is fitted to the output mask and to the union of the output mask and the largest connected component in the alternative mask. If the ratio between the area of the two bounding boxes drops below $\theta_{\text{anc}} = 0.7$, the current frame is considered as a potential candidate to update DRM. Note that updating with a grossly incorrectly segmented target would lead to memory corruption and eventual tracking failure. We thus trigger the DRM update only during sufficiently stable tracking periods, i.e., if the predicted IoU score from SAM2 exceeds a threshold $\theta_{\text{IoU}} = 0.8$ and if the

²At tracker initialization, RAM occupies all N_{DAM} slots and drops to $\frac{1}{2}N_{\text{DAM}}$ as DRM entries arrive, to fully exploit the available capacity.




Figure 2. Overview of the SAM2 memory and the proposed Distractor-Aware Memory (DAM), which splits the model into *Recent Appearance Memory* (RAM) and *Distractor Resolving Memory* (DRM) and updates them by a new memory management protocol.

mask area is within $\theta_{\text{area}} = 20\%$ of the median area in the last $\theta_M = 10$ frames. Note that SAM2.1++ is insensitive to the exact value of these parameters.

4. A distractor-distilled dataset

While benchmarks played a major role in the recent visual object tracking breakthroughs, we note that many of these contain sequences, which are no longer considered challenging by modern standards. In fact, most of the modern trackers obtain high performance on these sequences, which overwhelms the total score and under-represents the improvements on challenging situations. To facilitate the tracking performance analysis of the designs proposed in this paper, we semi-automatically distill several benchmarks into a distractor-distilled tracking dataset (DiDi).

We considered validation and test sequences of the major tracking benchmarks, which are known for high-quality annotation, i.e., GoT-10k [19], LaSOT [15], UTB180 [1], VOT-ST2020 and VOT-LT2020 [22], and VOT-ST2022 and VOT-LT2022 [23]. This gave us a pool of 808 sequences. A sequence was selected for the DiDi dataset if at least one-third of the frames passed the distractor presence criterion described next.

A frame was classified as containing non-negligible distractors if it contained a large enough region visually similar to the target. This criterion should be independent from tracker localization method, yet should reflect the power of the modern backbones. We thus encoded the image by DINO2 [34] and for each pixel in the feature space computed the distractor score as the average cosine distance to the features within the ground truth target region. We then computed the ratio between the number of pixels outside and inside the target region that exceeded the average of the scores computed within the target region. If the ratio exceeded 0.5, we considered the frame as containing non-negligible distractors.

Using the aforementioned protocol, we finally obtained 180 sequences with an average sequence length of 1.5k frames (274,882 frames in total). Each sequence contains a single target annotated by an axis-aligned bounding box. In addition, we manually segmented the initial frames to allow

the initialization of segmentation-based trackers. Figure 3 shows frames from the proposed DiDi dataset. Please see the supplementary material for additional information.

5. Experiments

The proposed DAM for SAM2 memory model is rigorously analyzed. Section 5.1 reports a battery of experiments to justify the design choices. Section 5.2 compares the SAM2.1 extension with DAM memory with the state-of-the-art on the DiDi dataset. Detailed analysis on the challenging VOT tracking-by-segmentation benchmarks is reported in Section 5.3, while a comparison on the standard bounding-box tracking benchmarks is reported in Section 5.4.

5.1. Architecture justification

The design choices of the proposed distractor-aware memory and the management protocol from Section 3.2 are validated on the DiDi dataset from Section 4. We compute the VOTS [24] measures since they simultaneously account for short-term as well as long-term tracking performance. Performance is summarized by the tracking quality Q score and two auxiliary measures: robustness (i.e., portion of successfully tracked frames) and accuracy (i.e., the average IoU between the prediction and ground truth during successful tracking). Results are shown in Table 1 and in AR plot in Figure 4.

We first verify the argument from Section 3.2.1, that updating with frames without a target present causes memory degradation. We thus extend SAM2.1, with updating only when the predicted mask is not empty (denoted as SAM2.1_{PRES}). SAM2.1_{PRES} increases the tracking quality Q by 2.5%, primarily by improved robustness, which justifies our claim.

We next verify the assumption (also claimed in [51]) that frequent updates reduce tracking robustness due to highly correlated information stored in memory. We reduce the update rate in SAM2.1_{PRES} to every 5th frame (SAM2.1 _{$\Delta=5$}). This negligibly improves Q , but does increase the robustness by 1.2%, which supports the claim. We did not observe further performance improvement with increasing Δ .



Figure 3. Example frames from the DiDi dataset showing challenging distractors. Targets are denoted by green bounding boxes.

Table 1. SAM2.1++ architecture justification on DiDi dataset.

	Quality	Accuracy	Robustness
SAM2.1	0.649	0.720	0.887
SAM2.1 _{PRES}	0.665	0.723	0.903
SAM2.1 _{$\Delta=5$}	0.667	0.718	0.914
SAM2.1 _{DRM1}	0.672	0.710	0.932
SAM2.1 _{DRM2}	0.644	0.691	0.913
SAM2.1++	0.694	0.727	0.944
DRM _{tenc}	0.669	0.711	0.925
RAM _{last}	0.685	0.724	0.932

Finally, we focus on the distractor-resolving memory (DRM) part of our proposed memory in Section 3.2.1, that we claim is responsible for the tracking robustness in the presence of distractors. Recall that DRM is updated when a distractor is detected and under the condition that tracking is reliable – we first test the influence of these two conditions independently. We thus extend SAM2.1 _{$\Delta=5$} with the new DAM memory and update the DRM part only during reliable tracking periods (SAM2.1_{DRM1}). The tracking accuracy improves a little, with robustness increase of 2%. Alternatively, we change the rule to update only when the distractor is detected (SAM2.1_{DRM2}). While robustness remains unchanged w.r.t. SAM2.1 _{$\Delta=5$} , the accuracy in fact drops. This is expected since distractor detection may be triggered also by the error in the target segmentation, which gets amplified by the update. To verify this, we next apply both our proposed update DRM rules, arriving to the proposed DAM with SAM2.1 (SAM2.1++ for short). Compared to SAM2.1 _{$\Delta=5$} , we observe substantial improvement in tracking quality Q (4%), primarily due to 3.3% boost in robustness as well as 1.3% boost in accuracy, taking the top-right position in the AR plot (Figure 4) among all variants. This conclusively verifies that DRM should be updated with distractor detected only if tracking is sufficiently reliable.

In Section 3.2.1 we claim that DRM part of the memory should not be time-stamped, since the frame influence on the distractor resolving in the current frame should not be biased by the temporal proximity and should serve as a time-less prior. To test this claim, we modify SAM2.1++ by

using temporal encodings in DRM (except for the initialization frame) – we denote it as DRM_{tenc}. The tracking quality drops by 3.6%, which confirms the claim. We further inspect the updating regime in RAM, which always includes the most recent frame, but updates the memory slots every 5th frame. SAM2.1++ is modified to update all RAM slots at every 5th frame (RAM_{last}). This results in a slight tracking quality decrease (1.3%), which indicates that including the most recent frame in RAM is indeed beneficial, but not critical.




Figure 4. Accuracy-robustness plot on DiDi for the ablated versions of SAM2.1++. The tracking quality is given at each label.

5.2. SoTa comparison on DiDi

SAM2.1++ is compared on the DiDi dataset with the recent state-of-the-art trackers TransT [5], SeqTrack [6], AQA-Track [40] and AOT [47] as well as trackers with explicit distractor handling mechanisms: KeepTrack [32], Cutie [9] and ODTrack [50]. For completion, we include the most recent, yet unpublished, tracker with improved long-term memory update for SAM2, named SAM2.1Long [14].

Results in Table 2 reveal the advantages of the track-

ers with an explicit distractor-handling mechanism over the other trackers. Consider two similarly complex recent trackers SeqTrack and ODTrack, which are both based on the ViT-L backbone. On classical benchmarks like LaSoT, LaSoT_{ext} and GoT10k, ODTrack outperforms SeqTrack by 2%, 6% and 4%, respectively (see Table 6). However, the performance gap increases to 15% on DiDi (Table 2), which confirms that distractors are indeed a major challenge for modern trackers and that DiDi has a unique ability to emphasize the tracking capability under these conditions and to expose tracking design weaknesses.

Focusing on the evaluation of the proposed tracker, SAM2.1++ outperforms all trackers – the standard state-of-the-art trackers as well as trackers with explicit distractor handling mechanism. In particular, SAM2.1++ outperforms state-of-the-art distractor-aware trackers ODTrack and Cutie, by 14% and 21%, respectively.

We compare the proposed SAM2.1++ with the concurrent, unpublished work SAMURAI [45]. SAMURAI is also built on top of SAM2.1 [36], focuses on handling distractors and improves memory management by incorporating motion cues into memory selection and mask refinement process. In this respect, the work is closely related to ours. Results show that SAM2.1++ outperforms SAMURAI in tracking quality by 2% on DiDi, mostly due to the higher robustness (i.e., SAM2.1++ tracks longer than SAMURAI). This result demonstrates the superiority of our new DAM memory and the management protocol for distractor handling, which is also less complex than the concurrent counterpart proposed in SAMURAI.

Compared to another unpublished tracker with the alternative memory design SAM2.1Long [14], SAM2.1++ outperforms it by a healthy 7% tracking quality boost, indicating superiority of our proposed memory. The results reveal that the major source of the performance boost is the SAM2.1++ tracking robustness, which means it fails less often and thus much better handles distractors. In fact, a close inspection of the results shows that SAM2.1Long performs on par with the baseline SAM2.1, which indicates that the long-term memory update mechanism presented in [14] does not improve performance in the presence of distractors. Finally, comparing SAM2.1++ to the baseline SAM2.1 reveals 7% boost in tracking quality, again, attributed mostly due to the improved robustness (6%).

These results validate the benefits of the proposed DAM memory and its management protocol in handling challenging distractors. Qualitative results of tracking and segmentation with SAM2.1++ on DiDi (Figure 5) further demonstrate the remarkable tracking capability in the presence of challenging distractors.

Table 2. State-of-the-art comparison on DiDi dataset.

	Quality	Accuracy	Robustness
SAMURAI [45]	0.680 ②	0.722 ③	0.930 ②
SAM2.1Long [14]	0.646	0.719	0.883
ODTrack [50]	0.608	0.740 ①	0.809
Cutie [9]	0.575	0.704	0.776
AOT [47]	0.541	0.622	0.852
AQATrack [40]	0.535	0.693	0.753
SeqTrack [6]	0.529	0.714	0.718
KeepTrack [32]	0.502	0.646	0.748
TransT [5]	0.465	0.669	0.678
SAM2.1 [36]	0.649 ③	0.720	0.887 ③
SAM2.1++	0.694 ①	0.727 ②	0.944 ①




Figure 5. SAM2.1++ qualitative results on the DiDi dataset with predicted masks shown in green, and tracked objects denoted by arrows. Per-frame overlaps are shown above the figures to indicate failure-free tracking over the entire sequence.

5.3. SoTa comparison on VOT benchmarks

The VOT initiative [21] is the major tracking initiative, providing challenging datasets for their yearly challenges. In contrast to most of the tracking benchmarks, the targets are annotated by segmentation masks, which allows more accurate evaluation of segmentation trackers, compared to the classic bounding-box benchmarks. In this paper, we include two recent single-target challenges: VOT2020 [22] and VOT2022 [23], as well as the most recent multi-target challenge VOTS2024 [25].

Table 3. State-of-the-art comparison on the VOT2020 benchmark. The challenge winner is marked by ⚡.

	EAO	Accuracy	Robustness
ODTrack [50]	0.605 ③	0.761	0.902 ③
MixViT-L+AR [11]	0.584	0.755	0.890
SeqTrack-L [6]	0.561	-	-
MixFormer-L [10]	0.555	0.762 ③	0.855
RPT ⚡ [31]	0.530	0.700	0.869
OceanPlus [49]	0.491	0.685	0.842
AlphaRef [44]	0.482	0.754	0.777
AFOD [7]	0.472	0.713	0.795
SAM2.1 [36]	0.681 ②	0.778 ②	0.941 ②
SAM2.1++	0.729 ①	0.799 ①	0.961 ①

The VOT2020 benchmark [22] consists of 60 challenging sequences, while trackers are run using anchor-based protocol [22] to maximally utilize each sequence. Tracking performance is measured by the accuracy and robustness, summarized by the primary measure called the expected average overlap (EAO).

Results on VOT2020 are shown in Table 3. The proposed SAM2.1++ outperforms all compared trackers. In particular, it outperforms the recently published MixViT [11] by 25% EAO, improving both, accuracy and robustness. SAM2.1++ outperforms also the VOT2020 challenge winner RPT [31] by a large margin (37.5% in EAO). Comparing SAM2.1++ to the original SAM2.1, the EAO is boosted by 7%, while accuracy and robustness are improved by 2.7% and 2.1%, respectively.

The VOT2022 benchmark [23] uses a refreshed dataset with 62 sequences (simplest sequences removed, and more challenging added). Table 4 includes the challenge top-performers, including the winner MS_AOT [47], as well as the recent published state-of-the-art trackers: DiffusionTrack [30], MixFormer [10], OSTRack [48] and D3Sv2 [28]. The proposed SAM2.1++ outperforms the VOT2022 winner MS_AOT by a significant margin of 12% in EAO. Note that the performance improvement is a consequence of both improved accuracy (2%) and robustness (3%) compared to the MS_AOT. In addition to achieving state-of-the-art performance, SAM2.1++ outperforms also the baseline SAM2.1 by a healthy margin of 9% EAO.

The results on both, VOT2020 and VOT2022 clearly show that SAM2.1++ outperforms all trackers, including the challenges top-performers and the recently published trackers, setting new state-of-the-art on these benchmarks. Despite its simplicity, the proposed *training-free* memory management is the key element for achieving excellent tracking performance.

VOTS2024 benchmark. For a complete evaluation

Table 4. State-of-the-art comparison on the VOT2022 benchmark. The challenge winner is marked by ⚡.

	EAO	Accuracy	Robustness
MS_AOT ⚡ [47]	0.673 ③	0.781	0.944 ③
DiffusionTrack [30]	0.634	-	-
DAMTMask [23]	0.624	0.796 ③	0.891
MixFormerM [10]	0.589	0.799 ②	0.878
OSTrackSTS [48]	0.581	0.775	0.867
Linker [42]	0.559	0.772	0.861
SRATransTS [23]	0.547	0.743	0.866
TransT_M [5]	0.542	0.743	0.865
GDFormer [23]	0.538	0.744	0.861
TransLL [23]	0.530	0.735	0.861
LWL_B2S [3]	0.516	0.736	0.831
D3Sv2 [28]	0.497	0.713	0.827
SAM2.1 [36]	0.692 ②	0.779	0.946 ②
SAM2.1++	0.753 ①	0.800 ①	0.969 ①

on VOT, we report the performance also on the most recent VOTS2024 benchmark. In contrast to VOT2020 and VOT2022, the VOTS2024 benchmark introduces a new, larger dataset, tracking multiple objects in the same scene (with ground truth sequestered at an evaluation server) and a new performance measure, designed to address short-term, long-term³, single and multi-target tracking scenarios. The VOTS2024 is currently considered as the most challenging tracking benchmark.

Results are reported in Table 5. It is worth pointing out that the top performers are mostly unpublished (not peer-reviewed) trackers, tuned for the competition, and often complex ad-hoc combinations of multiple methods. For example, the challenge winner S3-Track combines visual and (mono)depth features, uses several huge backbones, and is much more complex than SAM2.1++. Despite this, using the same parameters as in other experiments, SAM2.1++ achieves a solid second place among the challenging VOTS2024 competition. In particular, it outperforms the challenge-tuned versions of the recently published trackers LORAT [27], Cutie [9], and the VOT2022 winner AOT [47]. Furthermore, the proposed memory management mechanism in SAM2.1++ contributes to 8% performance boost in tracking quality compared to the baseline SAM2.1, mostly due to 9% higher robustness.

5.4. SoTa comparison on bounding box benchmarks

For a complete evaluation, we compare SAM2.1++ on the following three standard bounding box tracking datasets: LaSoT [15], LaSoT_{ext} [16] and GoT10k [19]. Since frames are annotated by bounding boxes and SAM2 requires a seg-

³The target may leave the field of view and re-appear later in the video.

Table 5. State-of-the-art comparison on the VOT2024 benchmark. The challenge winner is marked by \mathfrak{Z} .

	Quality	Accuracy	Robustness
S3-Track \mathfrak{Z} [26]	0.722 ①	0.784	0.889 ①
DMAOT_SAM [46]	0.653	0.794 ①	0.780
HQ-DMAOT [46]	0.639	0.754	0.790
DMAOT [46]	0.636	0.751	0.795 ③
LY-SAM [25]	0.631	0.765	0.776
Cutie-SAM [9]	0.607	0.756	0.730
AOT [47]	0.550	0.698	0.767
LORAT [27]	0.536	0.725	0.784
SAM2.1 [36]	0.661 ③	0.791 ③	0.790
SAM2.1++	0.711 ②	0.793 ②	0.864 ②

mentation mask provided in the first frame, we use the same SAM2 model to estimate the initialization mask. The min-max operation is applied on the predicted masks to obtain the axis-aligned bounding boxes required for the evaluation. Tracking performance is computed using area under the success rate curve [39] (AUC) in LaSoT [15] and LaSoT_{ext} [16] and the average overlap [19] (AO) in Got10k.

LaSoT [15] is a large-scale tracking dataset with 1,400 video sequences, with 280 evaluation sequences and the rest are used for training. The sequences are equally split into 70 categories, where each category is represented by 20 sequences (16 for training and 4 for evaluation). The dataset consists of various scenarios covering short-term and long-term tracking. Results are shown in Table 6. The proposed SAM2.1++ outperforms the baseline SAM2.1 [36] by 7.3%, which indicates that the proposed memory management is important also in a bounding box tracking setup. Furthermore, SAM2.1++ performs on par with the top-performing tracker LORAT [27], which was tuned on LaSoT training set, i.e., on the categories included in the evaluation set. It is worth noting that LORAT [27] has approximately 50% more training parameters than SAM2.1++, making the model significantly more complex.

LaSoT_{ext} [16] is an extension of the LaSoT [15] dataset, by 150 test sequences, divided into 15 new categories, which are not present in the training dataset. The results in Table 6 show that SAM2.1++ outperforms the baseline version by a comfortable margin of 7% in AUC. In addition, it outperforms the second-best tracker LORAT [27] by 7.6%. This indicates that SAM2.1++ generalizes well across various object categories while existing trackers suffer a much larger performance drop.

GoT10k [19] is another widely used large-scale tracking dataset, composed of approximately 10k video sequences, from which 180 sequences are used for tracking evaluation. We observed that top-performing trackers on

GoT10k test set achieve excellent tracking performance, e.g., more than 78% of average overlap, which leaves only small room for potential improvements. However, a solid 3.7% boost in tracking performance is observed when comparing SAM2.1++ to the top-performers LORAT [27] and ODTrack [50]. A close inspection of the SAM2.1++ results reveals that more than 99% of frames are successfully tracked (i.e., with a non-zero overlap), which indicates that GoT10k [19] difficulty level is indeed diminishing for modern trackers.

Table 6. State-of-the-art comparison on three standard bounding-box benchmarks.

	LaSoT (AUC)	LaSoT _{ext} (AUC)	GoT10k (AO)
MixViT [11]	72.4	-	75.7
LORAT [27]	75.1 ①	56.6 ③	78.2 ③
ODTrack [50]	74.0 ②	53.9	78.2 ③
DiffusionTrack [30]	72.3	-	74.7
DropTrack [38]	71.8	52.7	75.9
SeqTrack [6]	72.5 ③	50.7	74.8
MixFormer [10]	70.1	-	71.2
GRM-256 [17]	69.9	-	73.4
ROMTrack [4]	71.4	51.3	74.2
OSTrack [48]	71.1	50.5	73.7
KeepTrack [32]	67.1	48.2	-
TOMP [33]	68.5	-	-
SAM2.1 [36]	70.0	56.9 ②	80.7 ②
SAM2.1++	75.1 ①	60.9 ①	81.1 ①

6. Conclusion

We proposed a new distractor-aware memory model DAM and a management regime for memory-based trackers. The new model divides the memory by its tracking functionality into the recent appearances memory (RAM) and a distractor-resolving memory (DRM), responsible for the tracking accuracy and robustness, respectively. Efficient update rules are proposed, which also utilize the tracker output to detect frames with critical distractors useful for updating DRM. In addition, a distractor-distilled dataset DiDi is proposed to facilitate studying tracking in challenging scenarios.

The proposed DAM memory is implemented with SAM2.1 [36], forming SAM2.1++. Extensive analysis confirms the design decisions. Without any retraining and using fixed parameters, SAM2.1++ sets a solid state-of-the-art on six benchmarks with a moderate (20%) speed reduction compared to SAM2.1 (i.e., 11 fps vs 13.3 fps). This makes a compelling case for arguably simpler localization architectures of memory-based frameworks compared to the current

tracking state-of-the-art. Furthermore, the results suggest more research should focus on efficient memory designs, with possibly learnable management policies. We believe these directions hold a strong potential for further performance boosts in future work.

Acknowledgements. This work was supported by Slovenian research agency program P2-0214 and projects J2-2506, L2-3169, Z2-4459 and COMET, and by supercomputing network SLING (ARNES, EuroHPC Vega - IZUM).

References

- [1] Basit Alawode, Yuhang Guo, Mehnaz Ummar, Naoufel Werghi, Jorge Dias, Ajmal Mian, and Sajid Javed. Utb180: A high-quality benchmark for underwater tracking. In *Proceedings of the Asian Conference on Computer Vision*, pages 3326–3342, 2022. 4, 13
- [2] Luca Bertinetto, Jack Valmadre, João F Henriques, Andrea Vedaldi, and Philip H S Torr. Fully-convolutional siamese networks for object tracking. In *European Conference on Computer Vision Workshops*, pages 850–865, 2016. 1
- [3] Goutam Bhat, Felix Järemo Lawin, Martin Danelljan, Andreas Robinson, Michael Felsberg, Luc Van Gool, and Radu Timofte. Learning what to learn for video object segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 777–794. Springer, 2020. 7
- [4] Yidong Cai, Jie Liu, Jie Tang, and Gangshan Wu. Robust object modeling for visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9589–9600, 2023. 8
- [5] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *Comp. Vis. Patt. Recognition*, pages 8126–8135, 2021. 1, 2, 5, 6, 7
- [6] Xin Chen, Houwen Peng, Dong Wang, Huchuan Lu, and Han Hu. Seqtrack: Sequence to sequence learning for visual object tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14572–14581, 2023. 1, 2, 5, 6, 7, 8
- [7] Yiwei Chen, Jingtao Xu, Jiaqian Yu, Qiang Wang, ByungIn Yoo, and Jae-Joon Han. Afod: Adaptive focused discriminative segmentation tracker. In *European Conference on Computer Vision*, pages 666–682. Springer, 2020. 7
- [8] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European Conference on Computer Vision*, pages 640–658. Springer, 2022. 1, 2, 3
- [9] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Joon-Young Lee, and Alexander Schwing. Putting the object back into video object segmentation. In *Comp. Vis. Patt. Recognition*, pages 3151–3161, 2024. 1, 2, 5, 6, 7, 8
- [10] Yutao Cui, Cheng Jiang, Limin Wang, and Gangshan Wu. Mixformer: End-to-end tracking with iterative mixed attention. In *Comp. Vis. Patt. Recognition*, pages 13608–13618, 2022. 1, 7, 8
- [11] Yutao Cui, Cheng Jiang, Gangshan Wu, and Limin Wang. Mixformer: End-to-end tracking with iterative mixed attention. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(6): 4129–4146, 2024. 7, 8
- [12] Martin Danelljan, Goutam Bhat, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Int. Conf. Computer Vision*, pages 6181–6190, 2019. 2
- [13] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, Philip HS Torr, and Song Bai. Mose: A new dataset for video object segmentation in complex scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20224–20234, 2023. 2
- [14] Shuangrui Ding, Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Yuwei Guo, Dahua Lin, and Jiaqi Wang. Sam2long: Enhancing sam 2 for long video segmentation with a training-free memory tree. *arXiv preprint arXiv:2410.16268*, 2024. 2, 5, 6
- [15] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *Comp. Vis. Patt. Recognition*, 2019. 4, 7, 8, 13
- [16] Heng Fan, Hexin Bai, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Harshit, Mingzhen Huang, Juehuan Liu, et al. Lasot: A high-quality large-scale single object tracking benchmark. *International Journal of Computer Vision*, 129: 439–461, 2021. 7, 8
- [17] Shenyuan Gao, Chunluan Zhou, and Jun Zhang. Generalized relation modeling for transformer tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18686–18695, 2023. 8
- [18] Lingyi Hong, Zhongying Liu, Wenchao Chen, Chenzhi Tan, Yuang Feng, Xinyu Zhou, Pinxue Guo, Jinglun Li, Zhaoyu Chen, Shuyong Gao, et al. Lvos: A benchmark for large-scale long-term video object segmentation. *arXiv preprint arXiv:2404.19326*, 2024. 2
- [19] Lianghua Huang, Xin Zhao, and Kaiqi Huang. GOT-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019. 4, 7, 8, 13
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollar, and Ross Girshick. Segment anything. In *Int. Conf. Computer Vision*, pages 4015–4026, 2023. 2
- [21] M. Kristan, J. Matas, A. Leonardis, T. Vojir, R. Pflugfelder, G. Fernandez, G. Nebehay, F. Porikli, and L. Cehovin. A novel performance evaluation methodology for single-target trackers. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2016. 6
- [22] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drobní, et al. The eighth visual object tracking VOT2020 challenge results. In *European Conference on Computer Vision Workshops*, 2020. 4, 6, 7, 13
- [23] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Hyung Jin

- Chang, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, et al. The tenth visual object tracking vot2022 challenge results. In *European Conference on Computer Vision*, pages 431–460. Springer, 2022. 1, 2, 4, 6, 7, 12, 13
- [24] Matej Kristan, Jiří Matas, Martin Danelljan, Michael Felsberg, Hyung Jin Chang, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drbohlav, Zhongqun Zhang, Khanh-Tung Tran, et al. The first visual object tracking segmentation vots2023 challenge results. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1796–1818, 2023. 4
- [25] Matej Kristan, Jiri Matas, Pavel Tokmakov, Michael Felsberg, Luka Čehovin Zajc, Alan Lukežič, Khanh-Tung Tran, Xuan-Son Vu, Johanna Bjorklund, Hyung Jin Chang, and Gustavo Fernandez. The second visual object tracking segmentation vots2024 challenge results, 2024. 1, 2, 6, 8
- [26] Xin Li, Deshui Miao, Zhenyu He, Yaowei Wang, Huchuan Lu, and Ming-Hsuan Yang. Learning spatial-semantic features for robust video object segmentation. *arXiv preprint arXiv:2407.07760*, 2024. 8
- [27] Liting Lin, Heng Fan, Zhipeng Zhang, Yaowei Wang, Yong Xu, and Haibin Ling. Tracking meets lora: Faster training, larger model, stronger performance. In *ECCV*, 2024. 7, 8
- [28] Alan Lukežič, Jiří Matas, and Matej Kristan. A discriminative single-shot segmentation network for visual object tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9742–9755, 2021. 1, 7
- [29] Alan Lukežič, Žiga Trojer, Jiří Matas, and Matej Kristan. A new dataset and a distractor-aware architecture for transparent object tracking. *International Journal of Computer Vision*, pages 1–14, 2024. 1, 2
- [30] Run Luo, Zikai Song, Lintao Ma, Jinlin Wei, Wei Yang, and Min Yang. Diffusiontrack: Diffusion model for multi-object tracking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3991–3999, 2024. 7, 8
- [31] Ziang Ma, Linyuan Wang, Haitao Zhang, Wei Lu, and Jun Yin. Rpt: Learning point set representation for siamese visual tracking. In *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 653–665. Springer, 2020. 7
- [32] Christoph Mayer, Martin Danelljan, Danda Pani Paudel, and Luc Van Gool. Learning target candidate association to keep track of what not to track. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 13444–13454, 2021. 2, 5, 6, 8
- [33] Christoph Mayer, Martin Danelljan, Goutam Bhat, Matthieu Paul, Danda Pani Paudel, Fisher Yu, and Luc Van Gool. Transforming model prediction for tracking. In *Comp. Vis. Patt. Recognition*, 2022. 2, 8
- [34] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *arXiv:2304.07193*, 2023. 4
- [35] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alexander Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv:1704.00675*, 2017. 2
- [36] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 2, 3, 6, 7, 8, 12
- [37] Chaitanya Ryali, Yuan-Ting Hu, Daniel Bolya, Chen Wei, Haoqi Fan, Po-Yao Huang, Vaibhav Aggarwal, Arkabandhu Chowdhury, Omid Poursaeed, Judy Hoffman, et al. Hiera: A hierarchical vision transformer without the bells-and-whistles. In *International Conference on Machine Learning*, pages 29441–29454. PMLR, 2023. 2
- [38] Qiangqiang Wu, Tianyu Yang, Ziquan Liu, Baoyuan Wu, Ying Shan, and Antoni B Chan. Dropmae: Masked autoencoders with spatial-attention dropout for tracking tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14561–14571, 2023. 8
- [39] Y. Wu, J. Lim, and M. H. Yang. Object tracking benchmark. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(9):1834–1848, 2015. 8
- [40] Jinxia Xie, Bineng Zhong, Zhiyi Mo, Shengping Zhang, Liangtao Shi, Shuxiang Song, and Rongrong Ji. Autoregressive queries for adaptive tracking with spatio-temporal transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19300–19309, 2024. 5, 6
- [41] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. 2
- [42] Zizheng Xun, Shangzhe Di, Yulu Gao, Zongheng Tang, Gang Wang, Si Liu, and Bo Li. Linker: Learning long short-term associations for robust visual tracking. *IEEE Transactions on Multimedia*, 26:6228–6237, 2024. 7
- [43] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10448–10457, 2021. 1, 2
- [44] Bin Yan, Xinyu Zhang, Dong Wang, Huchuan Lu, and Xiaoyun Yang. Alpha-refine: Boosting tracking performance by precise bounding box estimation. In *Comp. Vis. Patt. Recognition*, pages 5289–5298, 2021. 7
- [45] Cheng-Yen Yang, Hsiang-Wei Huang, Wenhao Chai, Zhongyu Jiang, and Jenq-Neng Hwang. Samurai: Adapting segment anything model for zero-shot visual tracking with motion-aware memory. *arXiv:2411.11922*, 2024. 6
- [46] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 8
- [47] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. *Ad-*

- vances in Neural Information Processing Systems*, 34:2491–2502, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [48] Botao Ye, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Joint feature learning and relation modeling for tracking: A one-stream framework. In *ECCV*, 2022. [7](#), [8](#)
- [49] Zhipeng Zhang, Houwen Peng, Jianlong Fu, Bing Li, and Weiming Hu. Ocean: Object-aware anchor-free tracking. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 771–787. Springer, 2020. [7](#)
- [50] Yaozong Zheng, Bineng Zhong, Qihua Liang, Zhiyi Mo, Shengping Zhang, and Xianxian Li. Odtrack: Online dense temporal token learning for visual tracking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7588–7596, 2024. [1](#), [5](#), [6](#), [7](#), [8](#)
- [51] Junbao Zhou, Ziqi Pang, and Yu-Xiong Wang. Rmem: Restricted memory banks improve video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18602–18611, 2024. [1](#), [2](#), [3](#), [4](#)

7. Appendix

7.1. Impact of the model size

The segment anything model 2 (SAM2) [36] was originally developed in four model sizes, denoted by tiny (T), small (S), base (B) and large (L). In Table 7 these four model sizes of unchanged SAM2.1 are compared with our SAM2.1++ version, presented in the paper on the new DiDi dataset. Results show a clear and consistent performance improvement across all four model sizes. In particular, the tracking quality improves by approximately 6% or 7%, depending on the model size, mostly due to the improved robustness. These results show that the proposed distractor-aware memory generalizes well over various model sizes, demonstrating that the model has not been tuned to the exact SAM2 model.

Table 7. Comparison of different model sizes on DiDi dataset. The ++ denotes the proposed tracker, while the T, S, B, L denote the tiny, small, base and large Hiera backbone sizes, respectively. *Params* denotes number of parameters, while *Acc.* and *Rob.* denote accuracy and robustness, respectively.

	Params	Quality	Acc.	Rob.
SAM2.1-T	39M	0.600	0.697	0.848
SAM2.1++-T	39M	0.642 $\uparrow 7\%$	0.695	0.907
SAM2.1-S	46M	0.630	0.718	0.866
SAM2.1++-S	46M	0.668 $\uparrow 6\%$	0.709	0.930
SAM2.1-B	81M	0.624	0.721	0.856
SAM2.1++-B	81M	0.664 $\uparrow 6\%$	0.709	0.930
SAM2.1-L	224M	0.649	0.720	0.887
SAM2.1++-L	224M	0.694 $\uparrow 7\%$	0.727	0.944

7.2. Impact of the model version

This section compares the performance improvements for the two individual SAM versions, i.e., SAM2 and SAM2.1. The SAM2.1 version improves the initial version in handling small and visually similar objects by introducing additional augmentation techniques in training. It also includes improved occlusion handling by training the model on longer frame sequences.

Results are shown in Table 8. To demonstrate the improvement of the 2.1 model version over version 2, we compare the SAM2 with the SAM2.1 on the DiDi dataset, which results in approximately 3.5% improvement in tracking quality. Next, we compare the original SAM2 and the version with our new memory model (i.e., SAM2++). The tracking performance improves by 7%, which is well beyond the performance improvement from SAM2 to SAM2.1 and supports the importance of a high-quality memory model and the memory management regime. A similar performance boost (close to 7%) is

observed between SAM2.1 and SAM2.1++, which implies complementarity of the new memory model with the baseline method performance improvements that come from better training. Similarly as in Section 7.1, we conclude that the proposed distractor-aware memory is robust to different model versions, demonstrating a consistent improvements in tracking performance on two SAM2 versions.

Table 8. Comparison of two SAM model versions: SAM2 and SAM2.1 on DiDi dataset.

	Quality	Accuracy	Robustness
SAM2	0.627	0.723	0.850
SAM2++	0.668 $\uparrow 7\%$	0.710	0.929
SAM2.1	0.649	0.720	0.887
SAM2.1++	0.694 $\uparrow 7\%$	0.727	0.944

7.3. Real-time performance

We further evaluate the proposed SAM2.1++ under real-time tracking constraints. We thus evaluate the tracker on VOT2022-RT [23] challenge⁴, which was specifically designed for real-time evaluation. Specifically, VOT challenges are run by VOT toolkits, which manage the real-time constraints. A frame is sent to the tracker, which needs to process it and report the target position at 20FPS frame rate. If the tracker is not able to process the frame in time, the prediction from the previous frame is used as the estimate for the current frame and next frame is sent to the tracker. Such a setup simulates actual real-time scenario, which is much more realistic than reporting just the average tracking speed. The tracking performance is measured using standard VOT2022 measures [23]: the primary measure expected average overlap (EAO), and two auxiliary measures, i.e., accuracy and robustness.

The results in Table 9 show that the proposed SAM2.1++ (L model size) outperforms all trackers that participated in VOT2022-RT challenge. In particular, it outperforms the challenge winner by 4% in EAO demonstrating excellent real-time performance. These results show that the proposed distractor-aware memory adds only a small computational overhead, yet bringing remarkable robustness capabilities and making it useful for real applications.

7.4. Sensitivity to threshold values

We analyze the sensitivity of the proposed SAM2.1++ to the exact value of the manually determined parameters. In particular, we focus on three thresholds defined in *Distractor-resolving memory* (Section 3.2.2). Experiments were conducted on the VOT2022 [23] dataset using the unsupervised

⁴SAM2.1 and SAM2.1++ were evaluated on the machine with the AMD EPYC 7763 64-Core 2.45 GHz CPU and Nvidia A100 40GB GPU.

Table 9. Real-time performance on VOT2022-RT challenge. The challenge winner is marked by ②.

	EAO	Accuracy	Robustness
MS_AOT ②	0.610 ③	0.751 ②	0.921 ③
OSTrackSTS	0.569	0.766 ①	0.860
SRATransTS	0.547	0.743 ③	0.866
SAM2.1	0.614 ②	0.722	0.922 ②
SAM2.1++	0.635 ①	0.717	0.942 ①

experiment to ensure fast execution and compute the average overlap (AO) as the performance measure.

We analyzed the influence of the threshold for the ratio between the target and alternative predicted masks, i.e. $\Theta_{anc} = 0.7$. This ratio is used for preemptive update upon distractor detection and is thus crucial for our distractor-aware memory. The results in Figure 6 show that tracking results are extremely stable for a wide range of $\Theta_{anc} \in [0.6, 0.9]$ demonstrating the robust design of the tracker.

Next, we tested the IoU score threshold ($\Theta_{IoU} = 0.8$), used to determine if the predicted mask is reliable for distractor testing. Results in Figure 6 show that tracking performance is extremely stable for a wide range of parameters ($\Theta_{IoU} \in [0.5, 0.8]$) scoring almost identical AO.

Finally, the mask area threshold was tested, i.e. $\Theta_{area} = 0.2$. This threshold is used to determine the tracking stability and is together with the Θ_{IoU} a necessary condition to trigger the DRM update. Results in Figure 6 show stable tracking performance across the wide range of thresholds and confirms that SAM2.1++ is not sensitive to the exact value of this parameter.

7.5. DiDi dataset statistics

In this section we provide additional information about the distractor-distilled dataset DiDi construction. In particular, number of sequences from each source dataset is given in the following:

- LaSoT [15]: 86 sequences
 - UTB-180 [1]: 56 sequences
 - VOT2022-ST [23]: 20 sequences
 - VOT2022-LT [23]: 7 sequences
 - VOT2020-LT [22]: 6 sequences
 - GoT10k [19]: 4 sequences
 - VOT2020-ST [22]: 1 sequences
-
- Total: 180 sequences (274,882 frames)

7.6. Qualitative analysis

Figure 7 presents a qualitative comparison between the baseline SAM2.1 and the proposed SAM2.1++ on four video sequences. In the first row a zebra is tracked with other zebras in its vicinity. When the zebra is partially oc-




Figure 6. Sensitivity of the SAM2.1++ to different values of three thresholds: Θ_{anc} , Θ_{IoU} and Θ_{area} . Experiments were done on VOT2022 using average overlap as the performance measure. The selected threshold value is marked by a green circle.

cluded, SAM2.1 drifts to the wrong zebra and starts to track it, while SAM2.1++ tracks only the visible part of the target during occlusion and stays on the selected zebra until the end of the sequence.

In the second row of Figure 7, the baseline SAM2.1 tracker successfully tracks the bus until the full occlusion and fails to re-detect it after the re-appearance. This failure occurs due to the too frequent memory updates when target is occluded and is successfully addressed with the proposed memory update in SAM2.1++.

The third row in Figure 7 shows tracking of a flamingo’s head. The baseline SAM2.1 tends to jump on the bird’s beak or extend to the whole body, since it prefers to segment the regions with so-called high objectness (i.e., regions with well-defined edges). The proposed SAM2.1++ successfully tracks the flamingo’s head even if the edge between the head and the neck is not clearly visible. In this case, part of the neck is segmented by an alternative mask and thus detected as a distractor. Updating the distractor resolving memory (DRM) using such *critical* frames results in a more stable and accurate tracking.

A similar effect is demonstrated in the fourth row of Figure 7, where a fish similar to the tracked fish occludes it and causes SAM2.1 to jump to it. On the other hand, SAM2.1++ successfully detects such critical frames, updates the DRM and avoids the tracking failure.



Figure 7. Qualitative comparison of the baseline SAM2.1 (blue) and the proposed SAM2.1++ (yellow). The symbol \approx denotes approximately identical outputs and \emptyset denotes an empty prediction (i.e, mask with all-zeros). Tracked object is denoted with a green arrow.