

Test-Time Training on Video Streams

Renhao Wang*, Yu Sun*, Arnub Tandon, Yossi Gandelsman,
Xinlei Chen, Alexei A. Efros, Xiaolong Wang

Editor: Samy Bengio

Abstract

Prior work has established Test-Time Training (TTT) as a general framework to further improve a trained model at test time. Before making a prediction on each test instance, the model is first trained on the same instance using a self-supervised task such as reconstruction. We extend TTT to the streaming setting, where multiple test instances – video frames in our case – arrive in temporal order. Our extension is online TTT: The current model is initialized from the previous model, then trained on the current frame and a small window of frames immediately before. Online TTT significantly outperforms the fixed-model baseline for four tasks, on three real-world datasets. The improvements are more than $2.2\times$ and $1.5\times$ for instance and panoptic segmentation. Surprisingly, online TTT also outperforms its offline variant that accesses strictly more information, training on all frames from the entire test video regardless of temporal order. This finding challenges those in prior work using synthetic videos. We formalize a notion of *locality* as the advantage of online over offline TTT, and analyze its role with ablations and a theory based on bias-variance trade-off.




Figure 1: In our streaming setting, the current model f_t makes a prediction on the current frame before it can see the next one. The prediction task here is segmentation. f_t is obtained through online TTT, initializing from the previous model f_{t-1} . Each video is treated as an independent unit. A sliding window of size k contains the current and previous frames as test-time training data for the self-supervised task.

*. Equal contribution. Correspondence to: yusun@berkeley.edu.

Renhao Wang, Yu Sun, Yossi Gandelsman, Alexei A. Efros are with UC Berkeley. Arnub Tandon is with Stanford University. Xinlei Chen is with Meta AI. Xiaolong Wang is with UC San Diego.

Project website with videos, dataset and code: <https://test-time-training.github.io/video>




Figure 2: Results for instance and panoptic segmentation on COCO Videos, and semantic segmentation on KITTI-STEP. Online TTT (green) performs the best, and only requires the realistic setting where the video frames arrive in a stream. Offline TTT (yellow) requires the rather unrealistic setting where all frames from the entire test video are available before making predictions. Still, online TTT outperforms offline by taking advantage of locality.

1. Introduction

Most models in machine learning today are fixed during deployment. As a consequence, a trained model must prepare to be robust to all possible futures. This can be hard because being ready for all futures limits the model’s capacity to be good at any particular one, even though only one future actually happens. The basic idea of Test-Time Training (TTT) is to continue training on the future once it arrives in the form of a test instance (Sun et al., 2020). Since each test instance is observed without a ground truth label, training is performed with self-supervision.

This paper investigates TTT on video streams, where each “future”, or test instance x_t , is a frame, and each video is treated as an independent unit. Naturally, x_t and x_{t+1} are visually similar. We focus on the intuition discussed above: Is training on the future once it actually happens better than training on all possible futures? More concretely, this question in the streaming setting asks about the role of *locality*: For making a prediction on a particular frame x_t , is it better to perform test-time training *offline* on all of x_1, \dots, x_T (the video ends at time T), or *online* on only x_t (and maybe a few previous frames)?

Our empirical evidence supports locality. The best performance is achieved through online TTT: For each x_t , only train on itself, and a small window of less than two seconds of frames immediately before t . We call this sliding window of frames the *explicit memory*. The optimal explicit memory needs to be short term – in plain language, some amount of **forgetting is actually beneficial**. This high-level finding challenges those from prior work in TTT (Wang et al., 2020; Volpi et al., 2022) and continual learning (Li and Hoiem, 2017; Lopez-Paz and Ranzato, 2017; Kirkpatrick et al., 2017), but is consistent with recent work in neuroscience (Gravitz, 2019).

For online TTT, parameters after training on x_t carry over as initialization for training on x_{t+1} . We call this the *implicit memory*. Because such an initialization is usually quite good to begin with, most of the benefit from online TTT is realized with only one gradient step per frame. Not surprisingly, the effectiveness of both explicit and implicit memory depends on *temporal smoothness* – that x_t and x_{t+1} are similar. In Section 6, we conduct ablations on both explicit and implicit memory, and develop a theory based on bias-variance trade-off under smoothness.

Experiments in this paper are also of practical interests, besides conceptual ones. Models for many computer vision tasks are trained with large datasets of still images, e.g. COCO, for segmentation, but deployed on video streams. The default is to naively run such models frame-by-frame, since temporal smoothing (i.e. averaging across a sliding window of predictions) offers little improvement. Online TTT significantly improves prediction quality on three real-world video datasets, for four tasks: semantic, instance and panoptic segmentation, and colorization. Figure 2 visualizes results for the first three tasks (since metrics for colorization are less reliable): **online TTT beats even the offline oracle.**

We also collect a new video dataset with dense annotations – COCO Videos. These videos are orders of magnitude longer than in other public datasets, and contain much harder scenes from diverse daily-life scenarios. Longer and more challenging videos better showcase the importance of locality, making it even harder to perform well on all futures at once. The improvements on COCO Videos are, respectively, more than $2.2\times$ and $1.5\times$ for instance and panoptic segmentation.

One of the most popular forms of self-supervision in computer vision is reconstruction: removing parts of the input image, then predicting the removed content (Vincent et al., 2008; Pathak et al., 2016; Bao et al., 2021; Xie et al., 2022). Recently, a class of deep learning models called masked autoencoders (MAE) (He et al., 2021), using reconstruction as the self-supervised task, has been highly influential. TTT-MAE (Gandelsman et al., 2022) adopts these models for test-time training using reconstruction. The main task in Gandelsman et al. (2022) is object recognition. Inspired by the empirical success of TTT-MAE, we use it as a subroutine inside online TTT, and extend it to other main tasks such as segmentation.

Prior work (Sun et al., 2020) experiments with online TTT (without explicit memory) in the streaming setting, but each x_t is drawn independently from the same test distribution. This test distribution is created by adding some synthetic corruption, e.g. Gaussian noise, to a test set of still images, e.g. ImageNet test set (Hendrycks and Dietterich, 2019). Therefore, all x_t s belong to the same “future”, and locality is meaningless: TTT on as many x_t s as possible achieves the best performance by learning to ignore the corruption. TTT on actual video streams is fundamentally different and much more natural.

More recently, Volpi et al. (2022) also experiments in the streaming setting. While the x_t s here are not independent, these videos are short clips, again simulated to contain synthetic corruptions, e.g. CityScapes with Artificial Weather. Therefore, like in Sun et al. (2020), each corruption moves all x_t s into almost the same “future”, which they call a domain. Since performance drop is caused by that shared corruption, it is best recovered by training on all x_t s. Their only dataset without corruptions (CityScapes) sees little improvement (1.4% relative to no TTT). There is no mentioning of locality, our basic concept of interest.

2. Related Work

2.1 Continual Learning

In the field of continual a.k.a. lifelong learning, a model learns a sequence of tasks in temporal order, and is asked to perform well on all of them (Van de Ven and Tolias, 2019; Hadsell et al., 2020). Here is the conventional setting: Each task is defined by a data distribution P_t , which produces a training set D_t^{tr} and a test set D_t^{te} . At each time t , the

model is evaluated on all the test sets $D_1^{\text{te}}, \dots, D_t^{\text{te}}$ of the past and present, and average performance is reported.

The basic solution is to simply train the model on all of $D_1^{\text{tr}}, \dots, D_t^{\text{tr}}$, which collectively have the same distribution as all the test sets. This is often referred to as the oracle with infinite memory (a.k.a. replay buffer) that remembers everything. However, due to memory constraints, the model at time t is only allowed to train on D_t^{tr} . More advanced solutions, therefore, focus on how to retain memory of past data only with model parameters (Santoro et al., 2016; Li and Hoiem, 2017; Lopez-Paz and Ranzato, 2017; Shin et al., 2017; Kirkpatrick et al., 2017; Gidaris and Komodakis, 2018).

Some of the literature extends beyond the conventional setting. Aljundi et al. (2019) uses continuous instead of discrete tasks across time. Purushwalkam et al. (2022) and Fini et al. (2022) perform self-supervised learning on unlabeled training sets, and evaluate the learned features on the test sets. Hoffman et al. (2014), Li and Hospedales (2020) and Panagiotakopoulos et al. (2022) use a labeled training set D_0^{tr} in addition to unlabeled training sets $D_1^{\text{tr}}, \dots, D_t^{\text{tr}}$, connecting with unsupervised domain adaptation. Díaz-Rodríguez et al. (2018) uses alternative metrics, e.g. forward transfer, to justify forgetting for reasons other than computational.

Much of continual learning is motivated by the hope to understand human memory and generalization through the lens of artificial intelligence (Hassabis et al., 2017; De Lange et al., 2021). Our work shares the same motivation, but focuses on test-time training, without distinct splits of training and test sets.

2.2 Test-Time Training

One of the earliest algorithms for training at test time is Bottou and Vapnik (1992): For each test input, train on its neighbors before making a prediction. This approach continues to be effective for support vector machines (SVM) (Zhang et al., 2006) and recently in large language models (Hardt and Sun, 2023). Bottou and Vapnik (1992), titled *Local Learning*, articulates locality as a basic concept in machine learning. Another line of work called transductive learning uses test data to add constraints to the margin of SVMs (Joachims, 2002; Collobert et al., 2006; Vapnik, 2013). The principle of transduction, as stated by Vapnik, also emphasizes locality (Gammerman et al., 1998; Vapnik and Kotz, 2006): “Try to get the answer that you really need but not a more general one.”

In computer vision, the idea of training at test time has been well explored for specific applications (Jain and Learned-Miller, 2011; Shocher et al., 2018; Nitzan et al., 2022; Xie et al., 2023), especially depth estimation (Tonioni et al., 2019a,b; Zhang et al., 2020; Zhong et al., 2018; Luo et al., 2020). Our paper extends TTT-MAE (Gandelsman et al., 2022), detailed in Section 3. TTT-MAE, in turn, is inspired Sun et al. (2020), which proposed the general framework for test-time training with self-supervision. The particular self-supervised task used in Sun et al. (2020) is rotation prediction (Gidaris et al., 2018). Many other papers have followed this framework since then (Hansen et al., 2020; Sun et al., 2021; Liu et al., 2021b; Yuan et al., 2023), including Volpi et al. (2022) on videos discussed in Section 1, and Azimi et al. (2022) which we discuss next.

In Azimi et al. (2022), each video is treated as a dataset of unordered frames instead of a stream. In particular, there is no concept of past vs. future frames. The same model is

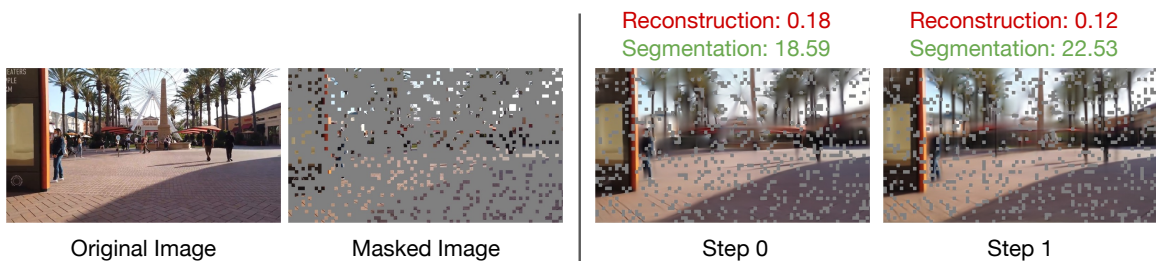


Figure 3: Training a masked autoencoder (MAE) to reconstruct each test image at test time. Reconstructed images on the right visualize the progress of gradient descent on this one-sample learning problem. For each test image, TTT-MAE first masks out majority of the patches. The masked image is given as input to the autoencoder, which then reconstructs those masked patches. The reconstruction loss is the pixel-wise mean squared error between the original and reconstructed patches. Loss on the main task – panoptic segmentation – also falls as reconstruction gets better.

used on the entire video. In contrast, our paper emphasizes locality. We have access to only the current and past frames, and our model keeps learning over time. In addition, all of our results are on real world videos, while Azimi et al. (2022) experiment on videos with artificial corruptions. These corruptions are also i.i.d. across frames.

Our paper is very much inspired by Mullapudi et al. (2018). To make video segmentation more efficient, their paper makes predictions frame-by-frame using a small student model. If the student is not confident, it queries an expensive teacher model, and then trains the student to fit the prediction from the teacher online. Thanks to temporal smoothness, the student can generalize confidently across many frames without querying the teacher, so learning and predicting combined is still faster than naively using the teacher at every frame. Our method only consists of one model, which learns from a self-supervised task instead of a teacher model. Rather than focusing on computational efficiency as in Mullapudi et al. (2018), the main goal of our paper is to improve inference quality. Behind their particular algorithm, however, we see the shared idea of locality, regardless of the form of supervision.

3. Background: TTT-MAE

Our paper extends the work of *Test-Time Training with Masked Autoencoders* (TTT-MAE) (Gandelsman et al., 2022), and uses TTT-MAE as the subroutine that updates the model for each frame. This section briefly describes TTT-MAE, as background for our extension. Figure 3 illustrates the process of TTT-MAE.

The general architecture for TTT with self-supervision (Sun et al., 2020) is Y-shaped with a stem and two heads: a prediction head g for the self-supervised task, another prediction head h for the main task, and a feature extractor f as the stem. The output features of f are shared between g and h as input. For TTT-MAE, the self-supervised task is masked image reconstruction (He et al., 2021). Following standard terminology for autoencoders, f is also called the encoder, and g the decoder.

Each input image x is first split into many non-overlapping patches. To produce the autoencoder input \tilde{x} , we mask out majority, e.g. 80%, of the patches in x at random. The self-supervised objective $\ell_s(g \circ f(\tilde{x}), x)$ compares the reconstructed patches from $g \circ f(\tilde{x})$ to the masked patches in x , and computes the pixel-wise mean squared error. For the main task, e.g. segmentation, all patches in the original x are given as input to $h \circ f$, during both training and testing.

3.1 Training-Time Training

There are three widely accepted ways to optimize the model components (f, g, h) at training time: joint training, probing, and fine-tuning. Fine-tuning is unsuitable for TTT, because it makes h rely too much on features that are used by the main task. Our paper uses joint training, described in Section 4. In contrast, Gandelsman et al. (2022) uses probing, which we describe next for completeness.

To prepare for probing, the common practice is to first train f and g with ℓ_s on the training set without ground truth. This preparation stage is also called self-supervised pre-training. Gandelsman et al. (2022) uses the encoder and decoder already pre-trained by He et al. (2021), denoted by f_0 and g_0 . During probing, the main task head h is then trained separately by optimizing for $\ell_m(h \circ f_0(x), y)$, on the training set with ground truth. f_0 is kept frozen. We denote h_0 as the main task head after probing. Since h_0 has been trained for the main task using features from f_0 as input, $h_0 \circ f_0$ can be directly applied on each test image as a baseline without TTT, keeping the parameters of f_0 and h_0 fixed.

3.2 Test-Time Training

At test time, TTT-MAE takes gradient steps on the following one-sample learning problem:

$$f', g' = \arg \min_{f, g} \ell_s(g \circ f(\tilde{x}'), x'), \quad (1)$$

then makes the final prediction $h_0 \circ f'(x')$. Crucially, the gradient-based optimization process always starts from f_0 and g_0 . When evaluating on a test set, Gandelsman et al. (2022) always discards f' and g' after making a prediction on each test input x' , and resets the weights back to f_0 and g_0 for the next test input. By test-time training on the test inputs independently, Gandelsman et al. (2022) does not assume that they can help each other.

In the original MAE design (He et al., 2021), g is very small relative to f , and only the visible patches, e.g. 20%, are processed by f . Therefore the overall computational cost of training for the self-supervised task is only a fraction, e.g. 25%, of training for the main task. In addition to speeding up training-time training for reconstruction, this reduces the extra test-time cost of TTT-MAE. Each gradient step at test time, counting both forward and backward, costs only half the time of forward prediction for the main task.

4. Test-Time Training on Video Streams

We consider each test video as a smoothly changing sequence of frames x_1, \dots, x_T ; time T is when the video ends. Each video is treated as an independent unit. In the streaming setting, an algorithm is evaluated on the video following its temporal order, like how a human would

consume it. At each time t , the algorithm should make a prediction on x_t after receiving it from the environment, before seeing any future frame. In addition to x_t , the past frames x_1, \dots, x_{t-1} are also available at time t , if the algorithm chooses to use them. Ground truth labels are never given to the algorithm on test videos.

Now we describe our algorithm for this streaming setting. At a high level, our algorithm simply amounts to a loop over the video frames, wrapped around TTT-MAE (Gandelsman et al., 2022). In practice, making it work involves many design choices.

4.1 Training-Time Training

At training time, if there was no self-supervision, then it is straightforward to optimize $h \circ f$ end-to-end for the main task only. The trained model produced by this process can already be applied on each x_t without TTT. In Table 1, we call this baseline *Main Task Only*. But such a model is not suitable for TTT, since then the self-supervised head g would have to be trained from scratch at test time.

To make g well-initialized for TTT, at training time we jointly optimize all three model components in a single stage, end-to-end, on both the self-supervised task and main task. This is called joint training. While joint training was also an option for prior work on TTT-MAE (Gandelsman et al., 2022), empirical experience at the time indicated that probing performed better (see Section 3). In this paper, however, we have successfully tuned joint training to be as effective as probing, and therefore default to joint training because it is simpler than the two-stage process of probing.

Following the notations in Section 3, the self-supervised task loss is denoted by ℓ_s , and the main task loss is ℓ_m . During joint training, we optimize those two losses together to produce a self-supervised task head g_0 , main task head h_0 , and feature extractor f_0 :

$$g_0, h_0, f_0 = \arg \min_{g, h, f} \frac{1}{n} \sum_{i=1}^n [\ell_m(h \circ f(x_i), y_i) + \ell_s(g \circ f(\tilde{x}_i), x_i)].$$

The summation is over the training set with n samples, each consisting of input x_i and label y_i . As discussed in Section 3, \tilde{x}_i is x_i transformed as input for the self-supervised task. In the case of MAE, \tilde{x}_i is obtained by masking 80% of the patches in x_i . Note that although the test instances come from video streams, training-time training uses labeled, still images, e.g. in the COCO training set, instead of unlabeled videos.

After joint training, the fixed model $h_0 \circ f_0$ can also be applied directly on each x_t without TTT, just like for *Main Task Only*. We call this new baseline *MAE Joint Training*. Empirically, these two baselines have roughly the same performance. Joint training does not hurt or help when only considering the fixed model after training-time training.

4.2 Test-Time Training

Another baseline is to blithely apply TTT-MAE by plugging each test frame x_t as x' into Equation 1, following the process in Section 3. We call this ablation *TTT-MAE No Memory* in Table 1 and Table 4. In this ablation, TTT for every x_t is initialized with h_0 and f_0 , by resetting the model parameters back to those after joint training. Like *Main Task Only* and

MAE Joint Training, this ablation misses the point of using a video. All three baselines treat each video as a collection of unordered, independent frames that might not contain any information about each other. None of the three can improve over time, no matter how long a video explores the same environment.

Improvement over time is only possible through some form of memory, by retaining information from the past frames x_1, \dots, x_{t-1} to help prediction on x_t . Because evaluation is performed at each timestep only on the current frame, our memory design should favor past data that are most relevant to the present. Fortunately, with the help of nature, the most recent frames usually happen to be the most relevant due to *temporal smoothness* – observations that are close in time tend to be similar. We design memory that favors recent frames in the following two ways.

Implicit memory. The most natural improvement is to simply not reset the model parameters between timesteps. That is, to initialize test-time training at timestep t with f_{t-1} and g_{t-1} , instead of f_0 and g_0 . This creates an implicit memory, since information carries over from the previous parameters, which was already optimized on previous frames. It also happens to be more biologically plausible: we humans do not constantly reset our minds. In prior work (Sun et al., 2020), TTT with implicit memory is called the “online” version, in contrast to the “standard” version with reset, for the baseline setting of independent images without temporal smoothness discussed in the paragraphs above.

Explicit memory. A more explicit way of remembering recent frames is to keep them in a sliding window. Let k denote the window size. At each timestep t , our method solves the following optimization problem instead of Equation 1:

$$f_t, g_t = \arg \min_{f, g} \frac{1}{k} \sum_{t'=t-k+1}^t \ell_s(g \circ f(\tilde{x}_{t'}), x_{t'}), \quad (2)$$

before predicting $h_0 \circ f_t(x_t)$. Optimization is performed with stochastic gradient descent: at each iteration, we sample a batch with replacement, uniformly from the same window. Masking is applied independently within and across batches. It turns out that only one iteration is sufficient for our final algorithm, because given temporal smoothness, implicit memory should already provide a good initialization for the optimization problem above.

4.3 Implementation Details

In principle, our method is applicable to any architecture. Our current implementation uses Mask2Former (Cheng et al., 2021), which has achieved state-of-the-art performance on many semantic, instance and panoptic segmentation benchmarks. Our Mask2Former uses a Swin-S (Liu et al., 2021c) backbone – in our case, this is also the shared encoder f . Everything following the backbone in the original architecture is taken as the main task head h , and our decoder g copies the architecture of h except the last layer that maps into pixel space for reconstruction. Joint training starts from their model checkpoint, which has already been trained for the main task. Only g is initialized from scratch.

Following He et al. (2021), we split each input into patches, and mask out 80% of them. However, unlike the Vision Transformers (Dosovitskiy et al., 2020) used in He et al. (2021), Swin Transformers use convolutions. Therefore, we must take the entire image as input (with the masked patches in black) instead of only the unmasked patches. Following Pathak

Setting	Method	COCO Videos		KITTI-STEP		
		Instance	Panoptic	Val.	Test	Time
Independent frames	Main Task Only	16.7	13.9	53.8	52.5	1.8
	MAE Joint Training	16.5	13.5	53.5	52.5	1.8
	TTT-MAE No Memory	35.4	20.1	53.6	52.5	3.8
Entire video available	Offline TTT-MAE All Frames	33.6	19.6	53.2	51.2	1.8
Frames in a stream	LN Adapt	16.5	14.7	53.8	52.5	2.0
	Tent	16.6	14.6	53.8	52.2	2.8
	Tent with Class Balance	16.7	14.8	53.8	52.5	3.7
	Self-Train	-	-	54.7	54.0	6.6
	Self-Train with Class Balance	-	-	54.1	53.6	6.9
	Online TTT-MAE (Ours)	37.6	21.7	55.4	54.3	4.1

Table 1: Metrics for instance, panoptic and semantic segmentation are, respectively, average precision (AP), panoptic quality (PQ), and mean IoU (%). Time is in seconds per frame, using a single A100 GPU, averaged over the KITTI-STEP test set. Time costs on COCO Videos are similar, thus omitted for clarity. The self-training baselines are not applicable for instance and panoptic segmentation because the model does not return a confidence per object instance.

et al. (2016), we use a fourth channel of binaries to indicate if the corresponding input pixels are masked. The model parameters for the fourth channel are initialized from scratch before joint training. If a completely transformer-based architecture for segmentation becomes available in the future, our method would like become even faster, by not encoding the masked patches (He et al., 2021; Gandelsman et al., 2022).

5. Results

We experiment with four applications on three real-world datasets: 1) semantic segmentation on KITTI-STEP – a public dataset of urban driving videos; 2) instance and panoptic segmentation on COCO Videos – a new dataset we annotated; 3) colorization on COCO Videos and a collection of black and white films. Please visit our project website at <https://video-ttt.github.io/> to watch videos of our results.

5.1 Additional Baselines

In Section 4, we have discussed three baselines: *Main Task Only*, *MAE Joint Training*, and *TTT-MAE No Memory*. We now discuss other baselines. Some of these baselines actually contain our own improvements.

Alternative architectures. The authors of Mask2Former did not evaluate it on KITTI-STEP. We benchmark Mask2Former on the KITTI-STEP validation set against two other popular models of comparable size: SegFormer B4 (Xie et al., 2021) (64.1M), and DeepLabV3+/RN101 (Chen et al., 2017) (62.7M), which is used by Volpi et al. (2022). Their

mean IoUs are, respectively, 42.0% and 53.1%. Given that *Main Task Only* in Table 1 has 53.8%, we can verify that our pre-trained model (69M) is indeed the state-of-the-art on KITTI-STEP. For COCO segmentation, the authors of Mask2Former have already compared with alternative architectures (Cheng et al., 2021), so we do not repeat their experiments.

Majority vote with augmentation. We also experiment with test-time augmentation of the input, applying the default data augmentation recipe in the codebase for 100 predictions per frame, then taking the majority vote across predictions as the final output. This improves *Main Task Only* by 1.2% mean IoU on the KITTI-STEP validation set. Combining the same technique with our method yields roughly the same improvement, indicating that they are again orthogonal. For clarity, we do not use majority vote elsewhere in this paper.

Temporal smoothing. We implement temporal smoothing by averaging the predictions across a sliding window, in the same fashion as our explicit memory. The window size is selected to optimize performance after smoothing on the KITTI-STEP validation set. This improves *Main Task Only* by only 0.4% mean IoU. Applying temporal smoothing to our method also yields 0.3% improvement. This indicates that our method is orthogonal to temporal smoothing. For clarity, we do not use temporal smoothing elsewhere in this paper.

Alternative techniques for TTT. Self-supervision with MAE is only one particular technique for test-time training. Subsection 4.2 describes a general loop, and any technique that does not use ground truth labels can be used as a subroutine to update the model inside the loop. We experiment with three promising techniques according to prior work: self-training (Volpi et al., 2022), layer norm (LN) adaptation (Schneider et al., 2020), and Tent (Wang et al., 2020). For self-training, our implementation significantly improves on the version in Volpi et al. (2022). Please refer to Appendix A for an in depth discussion of these three techniques.

Class balancing. Volpi et al. (2022) proposes a heuristic that can be used in conjunction with implicit memory: Record the number of predicted classes, for the initial model $h \circ f_0$ and the current model $h \circ f_t$. Reset the model parameters when the difference is large enough, in which case the predictions of the current model have likely collapsed. To compare with Volpi et al. (2022), we evaluate this heuristic on self-training and Tent. This heuristic cannot be applied to LN Adapt, which does not actually modify the trainable parameters in the model.

5.2 Semantic Segmentation on KITTI-STEP

KITTI-STEP (Weber et al., 2021) contains 9 validation videos and 12 test videos of urban driving scenes.¹ At the rate of 10 frames-per-second, these videos are the longest – up to 106 seconds – among public datasets with dense pixel-wise annotations. All hyper-parameters, even for COCO Videos, are selected on the KITTI-STEP validation set. Joint training is performed on CityScapes (Cordts et al., 2016), another driving dataset with exactly the same 19 categories as KITTI-STEP, but containing still images instead of videos.

1. KITTI-STEP was originally designed to benchmark instance-level tracking, and has a separate test set held-out by the organizers. The official website evaluates only tracking-related metrics on this test set. Therefore, we perform our own evaluation using the segmentation labels. Since we do not perform regular training on KITTI-STEP, we use the training set as test set.

Dataset	Length	Frames	Rate	Classes
CityScapes-VPS (Kim et al., 2020)	1.8	3000	17	19
DAVIS (Pont-Tuset et al., 2017)	3.5	3455	30	-
YouTube-VOS (Xu et al., 2018)	4.5	123,467	30	94
KITTI-STEP (Weber et al., 2021)	40	8,008	10	19
COCO Videos (Ours)	350	10,475	10	134

Table 2: Video datasets with annotations for segmentation. The columns are: average length per video in seconds, total number of frames in the entire dataset, rate in frames per second, and total number of classes. COCO Videos is larger than KITTI-STEP in total frames. The 3 videos in our new dataset are roughly an order of magnitude longer than those in KITTI-STEP, and much more diverse in terms of number of classes.

Table 1 presents our main results. Figure 7 in the appendix visualizes predictions on two frames. Please see project website for more visualizations. *Online TTT-MAE* in the streaming setting, using both implicit and explicit memory, performs the best. For semantic segmentation, such an improvement is usually considered highly significant in the community.

Baseline techniques that adapt the normalization layers alone do not help at all in these evaluations. This finding agrees with the evidence in Volpi et al. (2022): LN Adapt and Tent help significantly on datasets with synthetic corruptions, but do not help on real-world dataset (e.g. CityScapes).

Online TTT-MAE optimizes for only one iteration per frame, and turns out slower than the baselines without TTT by $2.3\times$. Comparing with Gandelsman et al. (2022), which optimizes for 20 iterations per frame (image), our method runs much faster, again, because implicit memory takes advantage of temporal smoothness to get a better initialization for every frame. Resetting parameters is wasteful on videos, because the adjacent frames are very similar.

5.3 COCO Videos

While KITTI-STEP already contains the longest annotated videos among publicly available datasets at the time of submission, they are still far too short for studying long-term phenomenon in locality. KITTI-STEP videos are also limited to driving scenarios, a small subset of the diverse scenarios in our daily lives. These limitations motivate us to collect and annotate our own dataset of videos.

We collected 3 videos, each about 5 minutes, annotated by professionals, in the same format as for COCO instance and panoptic segmentation (Lin et al., 2014). The benchmark metrics are also the same as in COCO: average precision (AP) for instance and panoptic quality (PQ) for panoptic. To put things into perspective, each of the 3 videos alone contains more frames, at the same rate, than all of the videos combined in the KITTI-STEP validation set. We compare this new dataset with other publicly available ones in Table 2.

Method	FID ↓	IS ↑	LPIPS ↑	PSNR ↑	SSIM ↑
Zhang et al. (Zhang et al., 2016)	62.39	5.00 ± 0.19	0.180	22.27	0.924
Main Task Only (Cheng et al., 2021)	59.96	5.23 ± 0.12	0.216	20.42	0.881
Online TTT-MAE (Ours)	56.47	5.31 ± 0.18	0.237	22.97	0.901

Table 3: Quantitative results for video colorization on COCO Videos. Arrows pointing up indicate higher the better, and pointing down indicate lower the better.

All videos are egocentric, similar to the visual experience of a human walking around. In particular, they do not follow any tracked object like in Oxford Long-Term Tracking (Valmadre et al., 2018) or ImageNet-Vid (Shankar et al., 2021). Objects leave and enter the camera view all the time. Unlike KITTI-STEP and CityScapes that focus on self-driving scenes, our videos are both indoor and outdoor.

We start with the publicly available Mask2Former model pre-trained on still images in the COCO training set. Analogous to our procedure for KITTI-STEP, joint training for TTT-MAE is also on COCO images, and our 3 videos are only used for evaluation. Mask2Former is the state-of-the-art on the COCO validation set, with 44.9 AP for instance and 53.6 PQ for panoptic segmentation. But its performance in Table 1 drops to 16.7 AP and 13.9 PQ on COCO Videos. This highlights the challenging nature of COCO Videos, and the fragility of models trained on still images when evaluated on videos in the wild.

We use exactly the same hyper-parameters as tuned on the KITTI-STEP validation set, for all algorithms considered. That is, all of our results for COCO Videos were completed in a single run. As it turns out in Figure 4, using a larger window size would further improve performance. However, we believe such hyper-parameters for TTT should not be tuned on the test videos, so we stick to the window size selected on the KITTI-STEP validation set.

Table 1 presents our main results. Comparing to *Main Task Only*, our relative improvements for instance and panoptic segmentation are, respectively, more than $2.2\times$ and $1.5\times$. Improvements of this magnitude on the state-of-the-art is usually considered dramatic. The self-training baselines are not applicable here because for instance and panoptic segmentation, the model does not return a confidence per object instance.

Interestingly, *TTT-MAE No Memory* also produces notable improvements on both tasks, and even outperforms *Offline TTT-MAE All Frames*. Considering this result from the perspective of locality, *Offline TTT-MAE All Frames* is the most global method, since it tries to be good at all frames in each video. At the other end of the spectrum, *TTT-MAE No Memory* is the most local, since it only uses information from the current frame. For COCO Videos, local is better than global, if one has to pick an extreme.

5.4 Video Colorization

The goal of colorization is to add realistic RGB colors to gray-scale images (Lei and Chen, 2019; Zhang et al., 2019). Our goal here is to demonstrate the generality of our method, not to achieve the state-of-the-art.

Method	COCO Videos		KITTI-STEP	
	Instance	Panoptic	Val.	Test
TTT-MAE No Memory	35.4	20.1	53.6	52.5
Implicit Memory Only	36.1	20.7	54.3	54.4
Explicit Memory Only	35.7	20.2	53.6	52.5
Online TTT-MAE (Ours)	37.6	21.7	55.4	54.3

Table 4: Ablations on our two forms of memory. For ease of reference, the values in Table 1 for *TTT-MAE No Memory* and *Online TTT-MAE (Ours)* are reproduced here.

Following Zhang et al. (2016), we simply treat colorization as a supervised learning problem. We use the same architecture as for segmentation – Swin Transformer with two heads, pre-trained on ImageNet – to predict colors given gray-scale images. We do not use domain-specific techniques, e.g., perceptual losses, adversarial learning, or diffusion. Our bare-minimal baseline already achieves results comparable to those in Zhang et al. (2016). *Online TTT-MAE* uses exactly the same hyper-parameters as for segmentation. All of our colorization experiments were completed in a single run. Because colorizing COCO Videos is expensive, we only evaluate *Online TTT-MAE* and *Main Task Only*.

For the quantitative results in Table 3, we colorize COCO Videos by first processing the 3 videos into black and white. This enables us to compare with the original videos in RGB. For qualitative results, we also colorize the 10 original black-and-white Lumiere Brothers films from 1895, roughly 40 seconds each, at the rate of 10 frames per second. Figure 9 in Appendix B provides a snapshot of our qualitative results. Please also see Appendix B for a list of the films and their lengths.

Our method outperforms the baseline and Zhang et al. (2016) on all metrics except SSIM. It is a field consensus that PSNR and SSIM often misrepresent actual visual quality because colorization is inherently multi-modal (Zhang et al., 2016, 2017), but we still include them for completeness. Please see the project website for the complete set of the original and colorized videos. Our method visually improves the quality in all of them comparing to the baseline, especially in terms of consistency across frames.

6. Analysis on Locality

Now we come back to the two philosophies presented at the beginning of our introduction: training on all possible futures in advance vs. training on the future once it actually happens. In other words, training globally vs. locally.

6.1 Empirical Analysis

Table 4 contains ablations on our two forms of memory: implicit and explicit. Both forms of memory contribute to the improvement of *Online TTT-MAE* over *TTT-MAE No Memory* (Gandelsman et al., 2022). Beyond these basic ablations, we further ablate three aspects of our method.




Figure 4: Effect of window size k on performance. In simple terms, *Online TTT-MAE* prefers a very short-term memory. For all window sizes, the batch size, and therefore computational cost, is fixed. The plot for Kitti-STEP is on the validation set, where we selected the optimal hyper-parameter $k = 16$. For all three tasks, with a rate of 10 frames per second, 16 frames cover only 1.6 seconds. The optimal k on COCO Videos turns out to be different for both semantic and panoptic segmentation, but the results we report in Table 1 still use $k = 16$. The y-values for window size $k = 1$ match those for *Implicit Memory Only* in Table 4. Note that the x-axis here is in log scale, in order to highlight the effect of k over a large range. Performance is actually not sensitive to k in linear scale.

Offline TTT-MAE. This ablation, presented at the beginning of the paper as the yellow bars in Figure 2, trains a single model for each test video. It lives in a new setting, where all frames from the entire test video are available for training with the self-supervised task, e.g. MAE, before predictions are made on that video. Strictly more information is provided here than the streaming setting, where only current and past frames are available. The frames are shuffled into a training set, and gradient iterations are taken on batches sampled from this training set, in the same way as sampled from the sliding window in *Online TTT-MAE*. To give *Offline TTT-MAE All Frames* even more advantage, we report results from the best iteration on each test video, as measured by actual test performance, which would not be available in real world. For many videos, this best iteration number is around 1000.

Window size. The choice of whether to use explicit memory is not binary. On one end of the spectrum, window size $k = 1$ is the same as *Implicit Memory Only*. On the other end, $k = \infty$ comes close to *Offline TTT-MAE All Frames*, except that the future frames cannot be trained on for $k = \infty$. Figure 4 analyzes the effect of window size on performance. We observe that too little memory hurts, but so does too much. This observation makes intuitive sense: frames further in the past become less relevant on average for making a prediction on the current frame, even though they provide more data for TTT. Figure 5 illustrates this intuition.

Temporal smoothness. As discussed in Section 4, temporal smoothness is the key assumption that makes our two forms of memory effective. While this assumption is intuitively necessary, we can verify its importance by shuffling all the frames within each video, therefore destroying temporal smoothness, and observing how results change. By construction, all three methods under the setting of independent frames – *Main Task Only*, *MAE Joint Training*, and *TTT-MAE No Memory* – are not affected. The same goes with *Offline TTT-MAE All Frames*, which already shuffles the frames during offline training. For




Figure 5: An illustration of the principle of locality in video streams. Our goal is improve prediction on the current frame, shot inside a lecture hall. The frame at $t - 10$ was still inside this hall. Including this frame in our sliding window decreases variance for TTT. However, the frame at $t - 200$ was shot before entering the hall. Including it would significantly increase bias, because it is no longer relevant to the current frame.

Online TTT-MAE, however, shuffling hurts performance dramatically. Performance on the KITTI-STEP validation set becomes worse than *Main Task Only*.

6.2 Theoretical Analysis

To complement our empirical observation that locality can be beneficial, we now rigorously analyze the effect of our window size k for TTT using any self-supervised task.

Notations. We first define the following functions of the shared model parameters θ :

$$\nabla \ell_m^t(\theta) := \nabla_{\theta} \ell_m(x_t, y_t; \theta), \quad (3)$$

$$\nabla \ell_s^t(\theta) := \nabla_{\theta} \ell_s(x_t; \theta). \quad (4)$$

These notations have appeared in Section 3 and 4, where the main task loss ℓ_m is defined for object recognition or segmentation, and the self-supervised task loss ℓ_s is instantiated as pixel-wise mean squared error for image reconstruction; θ refers to parameters of the encoder f .

Problem statement. Taking gradient steps with $\nabla \ell_m^t$ directly optimizes the test loss, since y_t is the ground truth label of test input x_t . However, y_t is not available, so TTT optimizes the self-supervised loss ℓ_s instead. Among the available gradients, $\nabla \ell_s^t$ is the most relevant. But we also have the past inputs x_1, \dots, x_{t-1} . Should we use some, or even all of them?

Theorem. For every timestep t , consider TTT with gradient-based optimization using:

$$\frac{1}{k} \sum_{t'=t-k+1}^t \nabla \ell_s^{t'}, \quad (5)$$

where k is the window size. Let θ_0 denote the initial condition, and $\tilde{\theta}$ where optimization converges for TTT. Let θ^* denote the optimal solution of ℓ_m^t in the local neighborhood of θ_0 .

Then we have

$$\mathbb{E} \left[\ell_m(x_t, y_t; \tilde{\theta}) - \ell_m(x_t, y_t; \theta^*) \right] \leq \frac{1}{2\alpha} \left(k^2 \beta^2 \eta^2 + \frac{1}{k} \sigma^2 \right),$$

under the following three assumptions:

1. In a local neighborhood of θ^* , ℓ_m^t is α -strongly convex in θ , and β -smooth in x .
2. $\|x_{t+1} - x_t\| \leq \eta$.
3. $\nabla \ell_m^t = \nabla \ell_s^t + \delta_t$, where δ_t is a random variable with mean zero and variance σ^2 .

The proof is in Appendix C.

Remark on assumptions. Assumption 1, that neural networks are strongly convex around their local minima, is widely accepted in the learning theory community (Allen-Zhu et al., 2019; Zhong et al., 2017; Wang et al., 2021). Assumption 2 is simply temporal smoothness in L^2 norm; any other norm could also be used here as long as the norm in Assumption 1 for strong convexity is changed accordingly. Assumption 3, that the main task and self-supervised task have correlated gradients, comes from the theoretical analysis of Sun et al. (2020).

Bias-variance trade-off. Disregarding the constant factor of $1/\alpha$, the upper bound in Theorem 1 is the sum of two terms: $k^2 \beta^2 \eta^2$ and σ^2/k . The former is the bias term, growing with η . The latter is the variance term, growing with σ^2 . More memory, i.e., sliding window with larger size k , reduces variance, but increases bias. This is consistent with our intuition in Figure 5. Optimizing this upper bound w.r.t. k reveals the theoretical sweet spot


$$k = \left(\frac{\sigma^2}{\beta^2 \eta^2} \right)^{1/3}.$$

7. Discussion

In the end, we connect our work to other ideas in machine learning.

Unsupervised domain adaptation. The setting of *Offline TTT-MAE All Frames*, where the entire unlabeled test video is available at once, is very similar to unsupervised domain adaptation (UDA). Each test video can be viewed as a target domain, and offline MAE practically treats the frames as i.i.d. data drawn from a single distribution. The only difference with UDA is that the unlabeled video serves as both training and test data. In fact, the modified version of UDA above is sometimes called *test-time adaptation*. Our results suggest that this perspective of seeing each video as a target domain might be misleading for algorithm design, because it discourages locality.

Continual learning. Conventional wisdom in the continual learning community believes that forgetting is harmful. Specifically, the best accuracy is achieved by remembering everything with an infinite replay buffer, given unlimited computer memory. Our streaming setting is different from those commonly studied by the continual learning community, because it does not have distinct splits of training and test sets, as explained in Subsection 2.1.



	Initial state	Update rule	Output rule	Cost
Naive RNN	$s_0 = \text{vector}()$	$s_t = \sigma(\theta_{ss}s_{t-1} + \theta_{sx}x_t)$	$z_t = \theta_{zs}s_t + \theta_{zx}x_t$	$O(1)$
Self-attention	$s_0 = \text{list}()$	$s_t = s_{t-1}.\text{append}(k_t, v_t)$	$z_t = V_t \text{softmax}(K_t^T q_t)$	$O(t)$
Naive TTT	$W_0 = f.\text{params}()$	$W_t = W_{t-1} - \eta \nabla \ell(W_{t-1}; x_t)$	$z_t = f(x_t; W_t)$	$O(1)$

Figure 6: Figure from Sun et al. (2024).

However, our sliding window can be viewed as a replay buffer, and limiting its size can be viewed as a form of forgetting. In this context, our results suggest that forgetting can actually be beneficial.

TTT on nearest neighbors. Here is an alternative heuristic for TTT: For each test instance, retrieve its nearest neighbors from a training set, and fine-tune the model on those neighbors before making a prediction on the test instance. This simple but effective heuristic has been explored in Bottou and Vapnik (1992) and Hardt and Sun (2023), as discussed in Subsection 2.2. Given temporal smoothness – that proximity in time translates to proximity in the retrieval metric, our sliding window can be seen as retrieving “neighbors” of the current frame. The only difference is that our “neighbors” are from the unlabeled test video instead of a labeled trainings set. This difference has two consequences. On one hand, we have to use self-supervision. On the other hand, our “neighbors” are still relevant (given temporal smoothness) even when the test instance is not represented by the training set.

In-context learning. In theory, each test video can be used as the context of a Transformer or RNN, both of which often exhibit the ability of “in-context learning” (Brown et al., 2020). As long as the model is autoregressive, the video is still processed as a stream. But in practice, this approach requires the model to be already trained on videos. Our approach, on the other hand, does not use videos as training data, as our very goal is to study the generalization from still images to videos.

Sequence modeling. Our method as shown in Figure 1 closely resembles an RNN as shown in Figure 6, if we think of W , the parameters of f , as the hidden state. From this perspective, online TTT can be regarded as compressing frames x_1, \dots, x_t into W_t , and gradient descent is simply a particular update rule. Following earlier versions of this paper, Sun et al. (2023) and Sun et al. (2024) program TTT into sequence modeling layers as an alternative to self-attention, and apply it to language modeling.

Acknowledgements

This project is supported in part by Oracle Cloud credits and related resources provided by the Oracle for Research program. Xiaolong Wang’s lab is supported, in part, by NSF CAREER Award IIS-2240014, Amazon Research Award, Adobe Data Science Research Award, and gifts from Qualcomm. We would like to thank Xueyang Yu and Yinghao Zhang for contributing to the published codebase. Yu Sun would like to thank his other PhD advisor, Moritz Hardt.

References

- Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11254–11263, 2019.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 242–252. PMLR, 09–15 Jun 2019.
- Yuki Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. *Advances in Neural Information Processing Systems*, 33:4660–4671, 2020.
- Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. *arXiv preprint arXiv:1911.05371*, 2019.
- Fatemeh Azimi, Sebastian Palacio, Federico Raue, Jörn Hees, Luca Bertinetto, and Andreas Dengel. Self-supervised test-time adaptation on video data. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3439–3448, 2022.
- Hangbo Bao, Li Dong, and Furu Wei. Beit: BERT pre-training of image transformers. *CoRR*, abs/2106.08254, 2021. URL <https://arxiv.org/abs/2106.08254>.
- Léon Bottou and Vladimir Vapnik. Local learning algorithms. *Neural computation*, 4(6): 888–900, 1992.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.

- Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. *arXiv preprint arXiv:2112.01527*, 2021.
- Ronan Collobert, Fabian Sinz, Jason Weston, Léon Bottou, and Thorsten Joachims. Large scale transductive svms. *Journal of Machine Learning Research*, 7(8), 2006.
- Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021.
- Natalia Díaz-Rodríguez, Vincenzo Lomonaco, David Filliat, and Davide Maltoni. Don't forget, there is more than forgetting: new metrics for continual learning. *arXiv preprint arXiv:1810.13166*, 2018.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Enrico Fini, Victor G Turrisi da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2022.
- A. Gammerman, V. Vovk, and V. Vapnik. Learning by transduction. In *In Uncertainty in Artificial Intelligence*, pages 148–155. Morgan Kaufmann, 1998.
- Yossi Gandelsman, Yu Sun, Xinlei Chen, and Alexei A. Efros. Test-time training with masked autoencoders. *Advances in Neural Information Processing Systems*, 2022.
- Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- Lauren Gravitz. The importance of forgetting. *Nature*, 571(July):S12–S14, 2019.
- Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020.

- Nicklas Hansen, Rishabh Jangir, Yu Sun, Guillem Alenyà, Pieter Abbeel, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. Self-supervised policy adaptation during deployment. *arXiv preprint arXiv:2007.04309*, 2020.
- Moritz Hardt and Yu Sun. Test-time training on nearest neighbors for large language models. *arXiv preprint arXiv:2305.18466*, 2023.
- Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258, 2017.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *CoRR*, abs/2111.06377, 2021.
- Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *CoRR*, abs/1903.12261, 2019. URL <http://arxiv.org/abs/1903.12261>.
- Judy Hoffman, Trevor Darrell, and Kate Saenko. Continuous manifold based adaptation for evolving visual domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 867–874, 2014.
- Vidit Jain and Erik Learned-Miller. Online domain adaptation of a pre-trained cascade of classifiers. In *CVPR 2011*, pages 577–584. IEEE, 2011.
- Thorsten Joachims. *Learning to classify text using support vector machines*, volume 668. Springer Science & Business Media, 2002.
- Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Video panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9859–9868, 2020.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Ananya Kumar, Tengyu Ma, and Percy Liang. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning*, pages 5468–5479. PMLR, 2020.
- Chenyang Lei and Qifeng Chen. Fully automatic video colorization with self-regularization and diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3753–3761, 2019.
- Da Li and Timothy Hospedales. Online meta-learning for multi-source and semi-supervised domain adaptation. In *European Conference on Computer Vision*, pages 382–403. Springer, 2020.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.

- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- Xiaofeng Liu, Bo Hu, Xiongchang Liu, Jun Lu, Jane You, and Lingsheng Kong. Energy-constrained self-training for unsupervised domain adaptation. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7515–7520. IEEE, 2021a.
- Yuejiang Liu, Parth Kothari, Bastien van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. Ttt++: When does self-supervised test-time training fail or thrive? *Advances in Neural Information Processing Systems*, 34, 2021b.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021c.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pages 6467–6476, 2017.
- Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (ToG)*, 39(4):71–1, 2020.
- Ke Mei, Chuang Zhu, Jiaqi Zou, and Shanghang Zhang. Instance adaptive self-training for unsupervised domain adaptation. In *European conference on computer vision*, pages 415–430. Springer, 2020.
- Ravi Teja Mullapudi, Steven Chen, Keyi Zhang, Deva Ramanan, and Kayvon Fatahalian. Online model distillation for efficient video inference. *arXiv preprint arXiv:1812.02699*, 2018.
- Yotam Nitzan, Kfir Aberman, Qiurui He, Orly Liba, Michal Yarom, Yossi Gandelsman, Inbar Mosseri, Yael Pritch, and Daniel Cohen-Or. Mystyle: A personalized generative prior. *arXiv preprint arXiv:2203.17272*, 2022.
- Theodoros Panagiotakopoulos, Pier Luigi Dovesi, Linus Härenstam-Nielsen, and Matteo Poggi. Online domain adaptation for semantic segmentation in ever-changing conditions. *arXiv preprint arXiv:2207.10667*, 2022.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- 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.
- Senthil Purushwalkam, Pedro Morgado, and Abhinav Gupta. The challenges of continuous self-supervised learning. *arXiv preprint arXiv:2203.12710*, 2022.

- Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. Semi-supervised self-training of object detection models. *arXiv*, 2005.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850, 2016.
- Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. *Advances in Neural Information Processing Systems*, 33:11539–11551, 2020.
- Vaishaal Shankar, Achal Dave, Rebecca Roelofs, Deva Ramanan, Benjamin Recht, and Ludwig Schmidt. Do image classifiers generalize across time? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9661–9669, 2021.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017.
- Assaf Shocher, Nadav Cohen, and Michal Irani. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3118–3126, 2018.
- Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020.
- Teo Spadotto, Marco Toldo, Umberto Michieli, and Pietro Zanuttigh. Unsupervised domain adaptation with multiple domain discriminators and adaptive self-training. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2845–2852. IEEE, 2021.
- Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning*, pages 9229–9248. PMLR, 2020.
- Yu Sun, Wyatt L Ubellacker, Wen-Loong Ma, Xiang Zhang, Changhao Wang, Noel V Csomay-Shanklin, Masayoshi Tomizuka, Koushil Sreenath, and Aaron D Ames. Online learning of unknown dynamics for model-based controllers in legged locomotion. *IEEE Robotics and Automation Letters*, 6(4):8442–8449, 2021.
- Yu Sun, Xinhao Li, Karan Dalal, Chloe Hsu, Sanmi Koyejo, Carlos Guestrin, Xiaolong Wang, Tatsunori Hashimoto, and Xinlei Chen. Learning to (learn at test time). *arXiv preprint arXiv:2310.13807*, 2023.
- Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong Wang, Sanmi Koyejo, et al. Learning to (learn at test time): Rnns with expressive hidden states. *arXiv preprint arXiv:2407.04620*, 2024.

- Alessio Tonioni, Oscar Rahnama, Thomas Joy, Luigi Di Stefano, Thalaiyasingam Ajanthan, and Philip HS Torr. Learning to adapt for stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9661–9670, 2019a.
- Alessio Tonioni, Fabio Tosi, Matteo Poggi, Stefano Mattoccia, and Luigi Di Stefano. Real-time self-adaptive deep stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 195–204, 2019b.
- Jack Valmadre, Luca Bertinetto, Joao F Henriques, Ran Tao, Andrea Vedaldi, Arnold WM Smeulders, Philip HS Torr, and Efstratios Gavves. Long-term tracking in the wild: A benchmark. In *Proceedings of the European conference on computer vision (ECCV)*, pages 670–685, 2018.
- Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- Vladimir Vapnik and S. Kotz. *Estimation of Dependences Based on Empirical Data: Empirical Inference Science (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006. ISBN 0387308652.
- Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, page 1096–1103, 2008.
- Riccardo Volpi, Pau De Jorge, Diane Larlus, and Gabriela Csurka. On the road to online adaptation for semantic image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19184–19195, 2022.
- Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020.
- Yifei Wang, Jonathan Lacotte, and Mert Pilanci. The hidden convex optimization landscape of regularized two-layer relu networks: an exact characterization of optimal solutions. In *International Conference on Learning Representations*, 2021.
- Mark Weber, Jun Xie, Maxwell Collins, Yukun Zhu, Paul Voigtlaender, Hartwig Adam, Bradley Green, Andreas Geiger, Bastian Leibe, Daniel Cremers, et al. Step: Segmenting and tracking every pixel. *arXiv preprint arXiv:2102.11859*, 2021.
- Binhui Xie, Shuang Li, Mingjia Li, Chi Harold Liu, Gao Huang, and Guoren Wang. Sepico: Semantic-guided pixel contrast for domain adaptive semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34, 2021.

- Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 585–601, 2018.
- Longhui Yuan, Binhui Xie, and Shuang Li. Robust test-time adaptation in dynamic scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15922–15932, 2023.
- Bo Zhang, Mingming He, Jing Liao, Pedro V Sander, Lu Yuan, Amine Bermak, and Dong Chen. Deep exemplar-based video colorization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8052–8061, 2019.
- Hao Zhang, Alexander C Berg, Michael Maire, and Jitendra Malik. Svm-knn: Discriminative nearest neighbor classification for visual category recognition. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2126–2136. IEEE, 2006.
- Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016.
- Richard Zhang, Jun-Yan Zhu, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. Real-time user-guided image colorization with learned deep priors. *arXiv*, 2017.
- Zhenyu Zhang, Stephane Lathuiliere, Elisa Ricci, Nicu Sebe, Yan Yan, and Jian Yang. Online depth learning against forgetting in monocular videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4494–4503, 2020.
- Kai Zhong, Zhao Song, Prateek Jain, Peter L Bartlett, and Inderjit S Dhillon. Recovery guarantees for one-hidden-layer neural networks. In *International conference on machine learning*, pages 4140–4149. PMLR, 2017.
- Yiran Zhong, Hongdong Li, and Yuchao Dai. Open-world stereo video matching with deep rnn. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 101–116, 2018.
- Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin Dogus Cubuk, and Quoc Le. Rethinking pre-training and self-training. *Advances in neural information processing systems*, 33:3833–3845, 2020.
- Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018.

Appendix A. Baseline Techniques for TTT

A.1 Self-Training

Self-training is a popular technique in semi-supervised learning (Radosavovic et al., 2018; Rosenberg et al., 2005; Zoph et al., 2020; Asano et al., 2019, 2020) and domain adaptation (Kumar et al., 2020; Zou et al., 2018; Mei et al., 2020; Liu et al., 2021a; Spadotto et al., 2021). It was also evaluated in Volpi et al. (2022) but produced inferior performance. We experiment with self-training both in its original form, and incorporating our own design improvements.

We assume that for each test image x , the prediction \hat{y} is also of the same shape in 2D. This assumption is satisfied in semantic segmentation and colorization. We also assume that F outputs a estimated confidence map \hat{c} of the same shape as \hat{y} . Specifically, for pixel $x[i, j]$, $\hat{y}[i, j]$ is the predicted class of this pixel, and $\hat{c}[i, j]$ is the estimated confidence of $\hat{y}[i, j]$.

Self-training repeats many iterations of the following:

- Start with an empty set of labels D for this iteration.
- Loop over every $[i, j]$ location, add pseudo-label $\hat{y}[i, j]$ to D if $\hat{c}[i, j] > \lambda$, for a fixed threshold λ .
- Train F to fit this iteration’s set D , as if the selected pseudo-labels are ground truth labels.

Our first design improvement is the confidence threshold λ . In Volpi et al. (2022), all predictions are used as pseudo-labels, regardless of confidence. Our experiments show that for low λ or $\lambda = 0$ as in Volpi et al. (2022), self-training is noisy and unstable, as expected.

However, for high λ , there is limited learning signal, e.g. little gradient, since f is already very confident about the pseudo-label. Our second design improvement, inspired by Sohn et al. (2020), is to make learning more challenging with an already confident prediction, by masking image patches in x . In Sohn et al. (2020), masking is applied sparingly on 2.5% of the pixels in average. We mask 80% of the pixels, inspired by He et al. (2021).

A.2 Layer Norm Adapt

Prior work (Schneider et al., 2020) shows that simply recalculating the batch normalization (BN) statistics works well for unsupervised domain adaptation. Volpi et al. (2022) applies this technique to video streams by accumulating the statistics with a forward pass on each frame once it is revealed. Since modern transformers use layer normalization (LN) instead, we apply the same technique to LN.

A.3 Tent

The normalization layers (BN and LN) also contain trainable parameters that modify the statistics. Optimizing those parameters requires a self-supervised objective. Tent (Wang et al., 2020) is an objective for learning only those parameters at test time, by minimizing the softmax entropy of the predicted distribution over classes. We update the LN statistics and parameters with Tent, in the same loop as our method, also using implicit and explicit memory. Hyper-parameters are searched on the KITTI-STEP validation set to be optimal for Tent.




Figure 7: Semantic segmentation predictions for adjacent frames from a video in KITTI-STEP. **Top:** Results using a fixed model baseline without TTT. Predictions are inconsistent between the two frames. The terrain on the right side of the road is incompletely segmented in both frames, and the terrain on the left is incorrectly classified as a wall on the first frame. **Bottom:** Results using *Online TTT-MAE*, by the same model, on the same frames as top. Predictions are now consistent and correct.




Figure 8: Panoptic segmentation predictions for adjacent frames from a video in our new COCO Videos dataset. **Top:** Results using a fixed model baseline without TTT. Predictions are inconsistent between the two frames. **Bottom:** Results using *Online TTT-MAE*, by the same model, on the same frames as top. Predictions are now consistent and correct. Please zoom in to see the instance labels.



Figure 9: Samples results for video colorization on the Lumiere Brothers films. **Top:** Using Zhang et al. (Zhang et al., 2016). **Middle:** Using our own baseline, Mask2Former with *Main Task Only*, which is already comparable, if not superior to Zhang et al. (2016). **Bottom:** After applying *Online TTT-MAE* on top of the baseline. Our colors are more vibrant and consistent within regions.

Appendix B. Colorization Dataset - Lumière Brothers Films

We provide results on the following 10 Lumiere Brothers films, all in the public domain:

1. Workers Leaving the Lumiere Factory (46 s)
2. The Gardener (49 s)
3. The Disembarkment of the Congress of Photographers in Lyon (48 s)
4. Horse Trick Riders (46 s)
5. Fishing for Goldfish (42 s)
6. Blacksmiths (49 s)
7. Baby's Meal (41 s)
8. Jumping Onto the Blanket (41 s)
9. Cordeliers Square in Lyon (44 s)
10. The Sea (38 s)

Appendix C. Proof of Theorem 1

We first prove the following lemma.

Lemma. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be α -strongly convex and continuously differentiable, and denote its optimal solution as x^* . Let

$$\tilde{f}(x) = f(x) + v^T x, \quad (6)$$

and denote its optimal solution as \tilde{x}^* . Then

$$f(\tilde{x}^*) - f(x^*) \leq \frac{1}{2\alpha} \|v\|^2. \quad (7)$$

Proof of lemma. It is a well known fact in convex optimization (Bubeck et al., 2015) that for f α -strongly convex and continuously differentiable,

$$\alpha(f(x) - f(x^*)) \leq \frac{1}{2} \|\nabla f(x)\|^2, \quad (8)$$

for all x . Since \tilde{x}^* is the optimal solution of \tilde{f} and \tilde{f} is also convex, we have $\nabla \tilde{f}(\tilde{x}^*) = 0$. But

$$\nabla \tilde{f}(x) = \nabla f(x) + v, \quad (9)$$

so we then have

$$\nabla f(\tilde{x}^*) = \nabla \tilde{f}(\tilde{x}^*) - v = -v. \quad (10)$$

Make $x = \tilde{x}^*$ in Equation 8, we finish the proof.

Proof of theorem. By Assumptions 1 and 2, we have

$$\|\nabla \ell_m^t(\theta) - \nabla \ell_m^{t-1}(\theta)\| \leq \beta \eta. \quad (11)$$

$$\frac{1}{k} \sum_{t'=t-k+1}^t \nabla \ell_s^{t'} = \frac{1}{k} \sum_{t'=t-k+1}^t \nabla \ell_m^{t'} + \frac{1}{k} \sum_{t'=t-k+1}^t \delta_{t'} \quad (12)$$

$$= \frac{1}{k} \sum_{t'=t-k+1}^t \left[\nabla \ell_m^t + \sum_{t''=t'}^{t-1} \left(\nabla \ell_m^{t''} - \nabla \ell_m^{t''+1} \right) \right] + \frac{1}{k} \sum_{t'=t-k+1}^t \delta_{t'} \quad (13)$$

$$= \nabla \ell_m^t + \frac{1}{k} \left[\sum_{t'=t-k+1}^t \sum_{t''=t'}^{t-1} \left(\nabla \ell_m^{t''} - \nabla \ell_m^{t''+1} \right) + \sum_{t'=t-k+1}^t \delta_{t'} \right] \quad (14)$$

To simplify notations, define

$$A = \sum_{t'=t-k+1}^t \sum_{t''=t'}^{t-1} \left(\nabla \ell_m^{t''} - \nabla \ell_m^{t''+1} \right), \quad (15)$$

$$B = \sum_{t'=t-k+1}^t \delta_{t'}. \quad (16)$$

So

$$\frac{1}{k} \sum_{t'=t-k+1}^t \nabla \ell_s^{t'} - \nabla \ell_m^t = (A + B)/k. \quad (17)$$

Because ℓ_m^t is convex in θ , we know that taking gradient steps with $\nabla \ell_m^t$ would eventually reach the local optima of ℓ_m^t . Because $\frac{1}{k} \sum_{t'=t-k+1}^t \nabla \ell_s^{t'}$ differs from $\nabla \ell_m^t$ by $(A + B)/k$, we know that taking gradient steps with the former reaches the local optima of $\ell_m^t + (A + B)\theta/2$. Now we can invoke our lemma. To do so, we first calculate

$$\mathbb{E} \left\| \frac{1}{k} \sum_{t'=t-k+1}^t \nabla \ell_s^{t'} - \nabla \ell_m^t \right\|^2 = \frac{1}{k^2} \mathbb{E} \|A + B\|^2 \quad (18)$$

$$= \frac{1}{k^2} (\|A\|^2 + \mathbb{E} \|B\|^2 + \mathbb{E} A^T B) \quad (19)$$

$$\leq \frac{1}{k^2} (k^4 \beta^2 \eta^2 + k \sigma^2) \quad (20)$$

$$= k^2 \beta^2 \eta^2 + \frac{1}{k} \sigma^2. \quad (21)$$

Then by our lemma, we have

$$\mathbb{E} \left[\ell_m(x_t, y_t; \tilde{\theta}) - \ell_m^* \right] \leq \frac{1}{2\alpha} \mathbb{E} \left\| \frac{1}{k} \sum_{t'=t-k+1}^t \nabla \ell_s^{t'} - \nabla \ell_m^t \right\|^2 \leq \frac{1}{2\alpha} \left(k^2 \beta^2 \eta^2 + \frac{1}{k} \sigma^2 \right). \quad (22)$$

This finishes the proof.