

Streaming Multiscale Deep Equilibrium Models

Can Ufuk Ertenli^[0000-0001-7795-3617], Emre Akbas^{*[0000-0002-3760-6722]}, and
Ramazan Gokberk Cinbis^{*[0000-0003-0962-7101]}

Middle East Technical University (METU), Department of Computer Engineering
{ufuk.ertenli,eakbas,gcinbis}@metu.edu.tr

Abstract. We present StreamDEQ, a method that infers frame-wise representations on videos with minimal per-frame computation. In contrast to conventional methods where compute time grows at least linearly with the network depth, we aim to update the representations in a continuous manner. For this purpose, we leverage the recently emerging implicit layer models, which infer the representation of an image by solving a fixed-point problem. Our main insight is to leverage the slowly changing nature of videos and use the previous frame representation as an initial condition on each frame. This scheme effectively recycles the recent inference computations and greatly reduces the needed processing time. Through extensive experimental analysis, we show that StreamDEQ is able to recover near-optimal representations in a few frames time, and maintain an up-to-date representation throughout the video duration. Our experiments on video semantic segmentation and video object detection show that StreamDEQ achieves on par accuracy with the baseline (standard MDEQ) while being more than $3\times$ faster. Code and additional results are available at <https://ufukertenli.github.io/streamdeq/>.

Keywords: Implicit layer models, video analysis and understanding, video object detection, video semantic segmentation

1 Introduction

Modern convolutional deep networks excel at numerous recognition tasks. It is commonly observed that deeper models tend to outperform their shallower counterparts [21,25,48], e.g. the prediction quality tends to increase with network depth using the architectures with residual connections [21]. Due to the sequential nature of the layer-wise calculations, however, increasing the network depth results in longer inference times. While the increase in inference duration can be acceptable for various offline recognition problems, it is typically of concern for many streaming video analysis tasks. For example, in perception modules of autonomous systems, it is not only necessary to keep up with the frame rate but

* Equal contribution for senior authorship.
To appear at ECCV 2022.

also desirable to minimize the computational burden of each recognition component to reduce the hardware requirements and/or save resources for additional tasks. Similar concerns arise in large-scale video analysis tasks, e.g. on video sharing platforms, a small increase in per-frame calculations can add up to great increments in total consumption.

Various techniques have been proposed to speed up the inference in deep networks. A widely studied idea is to apply a large model to selected *key-frames* and then either interpolate its features to the intermediate frames [53,51] or apply a smaller model to them [47,35]. However, such approaches come with several potential complications: (i) each time the larger model is applied, the model lags behind, the handling of which demands a complicated system design. (ii) Most methods require optical flow or motion estimates [53,51], which brings in an additional estimation problem and an additional point of failure. In addition, the time cost of the flow estimation naturally reduces the time budget for all dependent steps. (iii) Special techniques need to be developed to maintain the compatibility of the representations and/or confidence scores obtained across the key and intermediate frames. (iv) The training schemes tend to be complicated due to the need for training over video mini-batches. It is also noteworthy that several models, e.g. [29,45], rely on forward **and** backward flow estimates, making them less suitable for streaming recognition problems due to non-causal processing.

A related approach is to select a subset of each frame to process. These methods typically aim to identify the most informative regions in the input [38,1,14]. For static images, the region selection process can continue until the model becomes confident about its predictions. However, when applied to videos, such subset selection strategies share shortcomings similar to approaches relying on flow-based intra-frame prediction approximations. The inputs change over time, therefore, the models have to choose between relying on optical flow to warp the rest of the features or to omit them entirely, which may result in obsolete representations over time [51].

In this context, the recently introduced *implicit layer models*, pioneered by the work on Deep Equilibrium Model (DEQ) [2] and Multiscale Deep Equilibrium Model (MDEQ) [4], offer a fundamentally different alternative to deep neural networks. DEQ (and MDEQ) shows that by using the fixed-points of a network as the representation, one can gain the representation power of deep models, using a network with only a few layers. The potential of DEQ to eliminate long chains of computations over network layers, therefore, renders it an attractive candidate towards building efficient streaming recognition models.

However, while DEQ provides a way to learn deep representations using shallow networks, the test-time inference process involves iterative fixed-point estimation algorithms, such as root finding or fixed-point iterations. Since each iteration can be interpreted as an increment in the network depth, DEQ effectively constructs deep networks for inference, and therefore, can still suffer from the run-time costs as in explicit deep networks.




Fig. 1. Our method, StreamDEQ, exploits the temporal smoothness between successive frames and extracts features via a small number of solver iterations, starting from the previous frame’s representation as initial solution. StreamDEQ accumulates and transfers the extracted information continuously over successive frames; effectively sharing computations across video frames in a causal manner

Our main insight is the potential to speed up the inference process, *i.e.* fixed-point estimation, by exploiting the temporal smoothness across neighboring frames in videos.¹ More specifically, we observe that the fully estimated MDEQ representation of a frame can be used for obtaining the approximate representations of the following frames, using only a few inference iterations. We further develop the idea, and show that even without fully estimating the representation at any time step, the implicit layer representation can be kept up-to-date by running the inference iterations over the iteration steps *and* video time steps, in a continuous manner. The final scheme, starting from scratch, accumulates and transfers the extracted information throughout the video duration. We, therefore, refer to the proposed method as Streaming DEQ, or StreamDEQ for short.

The main difference between standard DEQ and StreamDEQ is illustrated in Figure 1. While DEQ typically requires a large number of inference iterations, StreamDEQ enables inference with only a few iterations per frame by leveraging the relevance of the most recent frame’s representation. On the start of a new stream, or after a major content change (e.g. a shot change), StreamDEQ quickly adapts to the video in a few frames, much like a person adapting her/his focus and attention when watching a new video. In the following frames, it continuously updates the representation to adapt to minor changes (e.g. objects moving, entering or exiting the scene).

Overall, StreamDEQ provides a simple and *lean* solution to the streaming recognition with implicit layer models, where a single model naturally performs cost-effective recognition, without relying on external inputs and heuristics, such as optical flow [22], post-processing methods (Seq-NMS [20] or tubelet re-scoring [30,29]). Our method also maintains the causality of the system, and executes in a continuous manner. We also note that it allows dynamic time budgeting; the duration of the inference process can be tuned on-the-fly by a

¹ We do not refer to a mathematical definition of smoothness, but rather emphasize that the changes between neighboring frames are small.

controller, depending on the instantaneous compute system load, which can be a desirable feature in real-world scenarios.

We verify the effectiveness of the proposed method through extensive experiments on video semantic segmentation and video object detection. Our experimental results show that StreamDEQ recovers near-optimal representations at much lower inference costs. More specifically, on the ImageNet-VID video object detection task, StreamDEQ converges to the mAP scores of 50.4 and 54.8 using only 4 and 8 inference iterations per frame, respectively. In comparison, the standard DEQ inference scheme yields only 8.2 and 32.6 mAP scores using 4 and 8 iterations, respectively. Similarly, on the Cityscapes semantic segmentation task, using StreamDEQ instead of the standard DEQ inference scheme improves the converged streaming mIoU score from 42.3 to 71.5 when 4 inference iterations are used per frame, and from 73.2 to 78.2 when 8 iterations are used per frame.

2 Related Work

Here, we summarize efficient video processing methods, video object detection and segmentation models. Furthermore, we discuss saliency based techniques for video processing. Finally, we give an overview of implicit models (DEQs).

Efficient Video Processing and Inference. There have been many efforts to improve video processing efficiency to reach real-time processing speeds. Most of these works take a system-oriented approach [10,39,31]. For example, Carreira et al. [10] develop an efficient parallelization scheme over multiple GPUs and process different parts of a model in separate GPUs to improve efficiency while sacrificing accuracy due to frame delays. Narayanan et al. [39] propose a novel scheduling mechanism that efficiently schedules and divides forward and backward passes over multiple GPUs. In another work, Li et al. [31] use a dynamic scheduler in which the model chooses to skip a frame when the delays build up to the point where it would be impossible to calculate the results of the next frame in the allotted time.

We also note that works on low-cost network designs, such as MobileNets [23,43] and low-resolution networks [34,49], are also relevant. Such efforts are valuable especially for replacing network components with more compute-friendly counterparts. However, the advantages of such techniques can also be limited due to natural trade-offs between speed and performance [50] as the lower-cost network components tend to have lower expressive power. Nevertheless, one can easily incorporate low-cost model design principles into DEQ or StreamDEQ models, thanks to the architecture-agnostic definition of implicit layer models. While such efforts may bring reductions in inference wall-clock time, they are outside the scope of our work.

Video Semantic Segmentation. Semantic segmentation is a costly, spatially dense prediction task. Its application to videos remains relatively limited. To

reduce the computational cost, most works rely on exploiting temporal relations between frames using methods such as feature warping [19,47,26,28], feature propagation [44,32], feature aggregation [24], and knowledge distillation [36].

Gadde et al. [19] propose warping features of the previous frame at different depths based on optical flow. Xu et al. [47] evaluate regions of the input frame and decide whether to warp the features with a cheap flow network or use the large segmentation model based on a confidence score. Huang et al. [26] keep a moving average over time by combining the segmentation maps from the current frame with the warped map from the previous frame. Jain et al. [28] warp high-quality features from the previous key-frame and fuse them with lower quality features calculated on the current frame to make predictions.

Shelhamer et al. [44] propose an adaptive method that schedules updates to the multi-level feature map so that features of layers with smaller changes are carried forward (without any transformation). Li et al. [32] introduce an adaptive key-frame scheduling method based on the deviation of low-level features compared to the previous key-frame and if the deviation is small, the features are propagated with spatially variant convolution.

Hu et al. [24] use a set of shallow networks, each calculating features of consecutive frames starting from scratch. Then, these features are aggregated at the current frame with an attention-based module. Liu et al. [36] propose to use an expensive network during training including optical flow and applies knowledge distillation on a student network to benefit from the high capacity of a teacher network while cutting computational costs thanks to a smaller and more efficient student network which the authors use during inference.

In contrast to all these approaches, the proposed StreamDEQ method directly leverages the similarities across video frames, without requiring any ad-hoc video handling strategies, as a way to adapt the implicit layer inference mechanism to efficient streaming video analysis.

Video Object Detection. Most modern video object detection methods exploit temporal information to improve the accuracy and/or efficiency. To this end, optical flow [53,52,29,45], feature aggregation [52,8,46,12] and post-processing techniques [20,30] are prominently used.

Zhu et al. [53] use optical flow estimates to warp features on selected key-frames to intermediate frames for increased efficiency. Zhu et al. [52] also propose FGFA that uses optical flow to warp features of nearby frames to the current frame and aggregates these features adaptively based on a weight calculated from feature similarity. Kang et al. [29] create links between objects through time (tubelets) from the predictions calculated with optical flow across a video linking objects through time and apply tubelet re-scoring to keep detections of high-confidence. Wang et al. [45] adds an instance level calibration module to FGFA [52] and combines them to generate better predictions.

Bertasius et al. [8] sample features from neighboring support frames via deformable convolution that learns object offsets between frames and aggregates these features over these frames. Wu et al. [46] focus on linking object proposals

in a video according to their semantic similarities. Chen et al. [12] propose a model that aggregates local and global information with a long-range memory.

Another common way to improve performance is to apply a post-processing method. For example, Han et al. [20] introduce Seq-NMS to exploit temporal consistency by constructing a temporal graph to link objects in adjacent frames. With a similar idea, Kang et al. [30] generate tubelets by combining single image detections through the video and use a tracker to re-scored the tubelets as a post-processing to improve temporal consistency.

Saliency Based Techniques. To reduce computational cost, another viable approach is to select important regions in an image and process only those small patches [38,1,14,37]. Video extensions of these models also exist [7,18,51].

Mnih et al. [38] and Ba et al. [1] model human eye movements by capturing *glimpses* from images with a recurrent structure and process those glimpses at each step. Cordonnier et al. [14] propose selecting most important regions to process by first processing a downsampled version of the image. Liu et al. [37] stops processing for regions with high-confidence predictions at an earlier stage.

Bazzani et al. [7] and Denil et al. [18] approach video processing in a human-like manner where the model *looks at* a different patch around the objects of interest at each frame and tracks them. Zhu et al. [51] takes a key-frame based approach. At each key-frame they process the full input and at intermediate frames, they update the feature maps partially based on temporal consistency.

Implicit Layer Models. Implicit layer models have seen a recent surge of interest and have been successful at numerous tasks. DEQs [2] are a recent addition to the implicit model family aimed at solving sequence modeling tasks. DEQs pose a fixed-point solving problem as a root finding problem and utilize an iterative root finding algorithm to find a solution. Multiscale Deep Equilibrium Models (MDEQ) [4] are the extension of the base DEQ to image-based models. Bai et al. [6] propose adding a Jacobian regularization term to improve model training.

In addition, Huang et al. [27] propose re-using the fixed-point across training iterations with the drawback of having to stay in full-batch mode for the training. Also, Bai et al. [5] suggests a new initialization scheme that is realized through a small network. Furthermore, inferring information from the last few iterations reduces the number of solver iterations required for convergence.

3 Proposed Method

We build our method on the Deep Equilibrium Model (DEQ). In this section, we first give an overview of DEQ and then present the details of our method.

3.1 DEQ Overview

Weight-tied networks are models where some or all layers share the same weights [3,16]. A DEQ is essentially a weight-tied network with only one shallow block. DEQ

leverages the fact that continuously applying the same layer to its output tends to guide the output to an equilibrium point, *i.e.* a fixed-point. Let \mathbf{x} represent the model’s input, \mathbf{z}^* the equilibrium point and f_θ the applied shallow block, then a DEQ can be described as

$$\lim_{i \rightarrow \infty} \mathbf{z}^{[i+1]} = \lim_{i \rightarrow \infty} f_\theta(\mathbf{z}^{[i]}; \mathbf{x}) \equiv f_\theta(\mathbf{z}^*; \mathbf{x}) = \mathbf{z}^*. \quad (1)$$

Since the *depth* of processing is obtained through repeatedly applying the same layer(s), these models are also called *implicit deep models*. DEQ’s fundamental difference from a standard weight-tied model is that the fixed-point is found by root finding algorithms in both forward and backward passes as in Eq. 2:

$$g_\theta(\mathbf{z}; \mathbf{x}) = f_\theta(\mathbf{z}; \mathbf{x}) - \mathbf{z} = 0 \implies \mathbf{z}^* = \text{RootFind}(g_\theta; \mathbf{x}). \quad (2)$$

DEQ uses the Broyden’s method [9] for this purpose. The accuracy of the solution depends on the number of Broyden iterations [5,6]. While more iterations yield better accuracy, they increase computation cost. DEQs have been successfully adapted to computer vision tasks, too, with the introduction of Multiscale Deep Equilibrium Models (MDEQ) [4]. MDEQ is a multiscale model where each scale is driven to equilibrium together with other scales again by using Broyden’s method. Iterations start with $\mathbf{z}^{[0]} = \mathbf{0}$ and continue N times to obtain the final solution, $\mathbf{z}^{[N]}$. N is set to 26 for ImageNet classification and 27 for Cityscapes semantic segmentation in MDEQ [4].

3.2 Streaming DEQ

Let \mathbf{X} be a $H \times W \times 3 \times T$ dimensional tensor representing a video where T is the temporal dimension. We represent the frame at time t with \mathbf{x}_t which is a $H \times W \times 3$ tensor. It should be noted that, we primarily target videos with temporal continuity, without too frequent shot changes. But we also study the effects of shot changes in Section 4.

To process a video, DEQ can be applied to each video frame \mathbf{x}_t to obtain \mathbf{z}_t , the representation of that frame. This amounts to running the Broyden solver for N iterations starting from $\mathbf{z}_t^{[0]} = \mathbf{0}$ for each frame.

However, we know a priori that transitions between subsequent video frames are typically smooth, *i.e.* $\mathbf{x}_{t+1} \sim \mathbf{x}_t$. From this observation, we hypothesize that the corresponding fixed-points, *i.e.* representations \mathbf{z}_t^* and \mathbf{z}_{t+1}^* , are likely to be similar. Therefore, the representation of the previous frame can be used effectively as a starting point for inferring the representation of the current frame. To validate this hypothesis, we run an analysis on the ImageNet-VID [42] dataset using the ImageNet pretrained MDEQ model. We assume that at each frame \mathbf{x}_t , we have access to the reference representation, \mathbf{z}_{t-1}^* , of the previous frame. Reference representations are obtained by running the MDEQ model until convergence ($N = 26$ iterations). At each frame, we use the reference representation of the previous frame as the starting point of the solver:

$$\mathbf{z}_t^{[0]} = \mathbf{z}_{t-1}^*, \quad (3)$$

and run the solver for various but small numbers of iterations, M . To analyze the amount of change in representations over time, we use an ImageNet-pretrained model (MDEQ-XL), since ImageNet representations are known to be useful in many transfer learning tasks. In Figure 2, we show the squared Euclidean distance between $\mathbf{z}_t^{[M]}$ and \mathbf{z}_t^* for various t values, when the solver is started with Eq. 3. Dashed lines correspond to the squared Euclidean distance between MDEQ-XL’s reference and M -iteration based representations.

From the results presented in Figure 2, we observe that when the solver is initialized with the preceding frame’s fixed-point, the inference process quickly converges towards the reference representation. We also observe that after starting from the reference representation of the previous frame and performing only 1 iteration on the current frame, the approximate representation is already more similar to the reference representation than starting from scratch and performing 8 iterations.

Next, we examine the case where the inference method is given access to the reference representations only at certain frames. To simulate this case, at each video clip, we compute the reference representation only at the first frame \mathbf{x}_0 , *i.e.* \mathbf{z}_0^* . In all following ones, we initialize the solver with the estimated representation of the preceding frame and run the solver for M iterations. That is,

$$\mathbf{z}_1^{[0]} = \mathbf{z}_0^* \text{ and } \mathbf{z}_t^{[0]} = \mathbf{z}_{t-1}^{[M]}. \quad (4)$$

We present the results of this scheme for $M \in \{1, 2, 4, 8\}$ in the left hand side of Figure 3. We observe that starting with the reference representation on the initial frame is still useful but for longer clips its effect diminishes. Still, this scheme helps us maintain a stable performance even after several frames. For example, starting with the reference representation and then applying $M = 2$ iterations per frame throughout the following 20 frames yields a representation that is closer to the reference representation of the final frame than the one given by baseline DEQ inference with 4 solver iterations. This result shows that the M -step inference scheme is able to keep up with the changes in the scene by starting from a good initial point.

While this scheme can provide efficient inference on novel frames, we would still need the reference representations of the initial frames, or key-frame(s), which would share the same problems with key-frame based video recognition approaches, *e.g.* [53,47,35]. To address this problem, we further develop the idea, and hypothesize that we can start from scratch (*i.e.* all zeros), do a limited number of iterations per frame, and pass the representation to the next frame




Fig. 2. Squared Euclidean approximation error as a function of inference steps, when the solver is initialized with the reference representation of the preceding frame




Fig. 3. Distance between the reference representations and StreamDEQ estimations for varying number of iterations, when StreamDEQ is initialized with reference representations (on the left) or just zeros (on the right) on the very first frame

as the starting point. That is,

$$\mathbf{z}_0^{[0]} = 0 \quad \text{and} \quad \mathbf{z}_t^{[0]} = \mathbf{z}_{t-1}^{[M]}. \quad (5)$$

We present the representation distance results for this final scheme in Figure 3 (right). The representation distances to the reference representations stabilize in 20 frames. Converged distance values (in 20 frames) are almost the same with those of the previous scheme (Eq.4). Additionally, the initial representations have relatively large distances but these differences get smaller as new frames arrive. We call this final scheme as StreamDEQ. This scheme avoids any heavy processing in any one of the frames, and completely avoids the concept of key-frames. The number of Broyden iterations can be tuned, which allows easy control over the time-vs-accuracy trade-off. Therefore, the inference iterations can be run as much as the time budget allows. An illustration of the StreamDEQ inference process is given in Figure 4.

4 Experimental Results

We evaluate our method on video semantic segmentation and video object detection. In the following, we provide technical details regarding training and inference setups, the datasets used, and present our experimental findings. We use the PyTorch [40] framework for all experiments.

4.1 Video Semantic Segmentation

Experimental Setup. We use the Cityscapes semantic segmentation dataset [15], which consists of 5K finely annotated and 20K coarsely annotated images. These finely annotated images are divided into train, val and test set, each containing 2975, 500, and 1525 images, respectively. They correspond to frames extracted

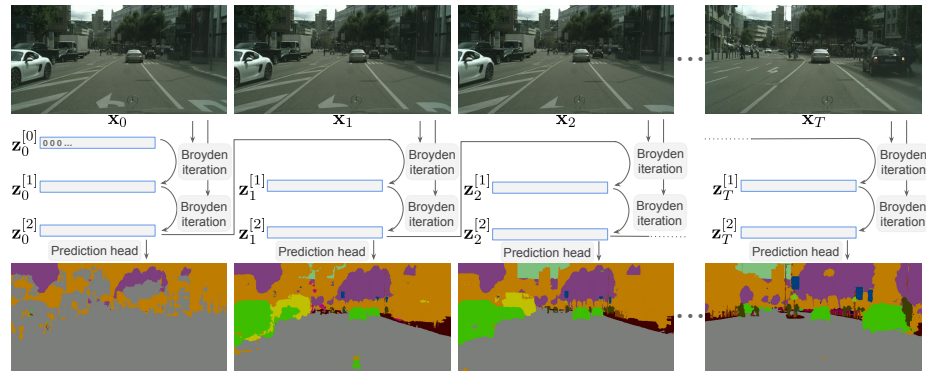


Fig. 4. StreamDEQ applied to a streaming video, performing two iterations per frame. The representation inference process is initialized with zeros in the very first frame ($\mathbf{z}_0^{[0]} = \mathbf{0}$), and with the most recent representation ($\mathbf{z}_t^{[0]} = \mathbf{z}_{t-1}^{[2]}$) in the rest of the stream. This scheme effectively recycles all recent computations for time-efficient inference on a new frame, and therefore, allows approximating a long inference chain (*i.e.* a deep network) by a few inference steps (*i.e.* a few layers) throughout the video stream

from video clips where each annotated image is the 20th frame of its respective clip. To evaluate over videos, we use these clips up to the 20th frame, which has fine annotations, and perform the evaluation on that frame.

We use the pretrained MDEQ-XL segmentation model from the MDEQ paper [4] and do not perform any additional training. We also do not make any changes to its evaluation setup or hyperparameters, perform the evaluation on Cityscapes `val` and report *mean intersection over union* (mIoU) results. For further details, we refer the reader to MDEQ [4].

Results. We present the results of StreamDEQ for two scenarios. The first scenario corresponds to Eq. 4, where we use the reference representations of the first frame to initialize the solver, and apply StreamDEQ then on. Results of this experiment in Figure 5 (left) show that as the offset of the evaluated frame increases, mIoU starts decreasing, which is expected because the further we move away from the first frame the more irrelevant its representation will become. However, mIoU then stabilizes at a value that is proportional to the number of Broyden iterations (the more iterations, the better the mIoU). This shows that StreamDEQ is able to extract better features over time. StreamDEQ’s performance with 8 iterations is still comparable with the baseline (MDEQ) with 27 iterations.

The second scenario corresponds to our final StreamDEQ proposal (*i.e.* Eq. 5), where we initialize the solver from scratch, *i.e.* with all zeros, and apply StreamDEQ. Results of this case are shown in Figure 5 (right). As the videos progress, it might be expected that the Broyden solver cannot keep up with the changing scenes. However, we observe that even after 20 frames, the accuracy




Fig. 5. StreamDEQ semantic segmentation results (in mIoU) on the Cityscapes dataset as a function of solver iterations when the first frame representation is initialized with the reference representation (left) versus zeros (right)

does not drop. Additionally, the impact of this method is clearer for lower numbers of iterations. For example, performing 1 iteration on every frame without our method would only yield an mIoU score of 2.2. However, StreamDEQ obtains a mIoU score of 44.9 in 10 frames. This is an improvement of over $20\times$. For 8 iterations, StreamDEQ is able to obtain 78.1 mIoU in 10 frames whereas the non-streaming baseline achieves only 73.2 mIoU. Moreover, the converged mIoU values (at larger frame offsets) are similar in Figure 5 (left) and Figure 5 (right). Therefore, we conclude that the initial point where we start the solver becomes less important as the video streams and the performance stabilizes at some value higher than the non-streaming case.

We also illustrate these results qualitatively in Figure 6. For 1 iteration, while the baseline cannot produce any meaningful segmentation, StreamDEQ starts capturing many segments correctly at the 4th frame. With 2 iterations, while the DEQ baseline still produces poor results, StreamDEQ starts to yield accurate predictions in early frames compared to the single iteration case. With 4 iterations, while both models provide rough but relevant predictions in the first frame, StreamDEQ predictions start to become clearly more accurate in the following frames; for example, tree trunks and the sky become visible only when StreamDEQ is applied.

We examine the effects of increasing the number of iterations on inference speed in Table 1. We note that our method does not introduce any computation overhead other than the time it takes to store the fixed-point representation of the previous frame. Therefore, we observe a linear increase in compute times as the number of iterations increases. StreamDEQ with 4 iterations achieves an mIoU score of 71.5 at 530 ms per image. MDEQ with 4 iterations can only achieve 42.3 mIoU. Additional inference results can be found in the videos provided on our project page.

Table 1. Inference time comparisons of the proposed method with differing number of iterations on Cityscapes and ImageNet-VID datasets

Model	# iterations	Cityscapes		ImageNet-VID	
		mIoU	FPS	mAP@50	FPS
StreamDEQ	1	45.5	4.3	9.1	10.3
StreamDEQ	2	57.9	2.9	39.5	9.2
StreamDEQ	4	71.5	1.9	50.4	6.2
StreamDEQ	8	78.2	1.1	54.8	3.5
MDEQ (Baseline)	27/26	79.7	0.3	55.0	1.2

Effect of Shot Changes. We also study the effects of shot changes for the video semantic segmentation task. For this purpose, we initialize the solver with the reference representations from a random frame from the Cityscapes or ImageNet-VID datasets and run StreamDEQ starting from the representations of that frame. Results of ImageNet-VID to Cityscapes shot change experiments can be found in Figure 7. In this case, even though at first, the obtained scores are lower, following a similar trajectory to the ones in Figure 5 (right), mIoU scores stabilize to a value close to our original experiment. We conclude that, even with occasional shot changes, our method is able to adapt to the new scene in a few frames. The Cityscapes-to-Cityscapes shot change experiments, with similar results, are provided in the supplementary material.

4.2 Video Object Detection

Experimental Setup. For the video object detection task, we evaluate our method on the ImageNet-VID dataset [42] utilizing the MMDetection [11] and MMTracking [13] frameworks. ImageNet-VID dataset consists of 3862 training and 555 validation videos from 30 classes that are a subset of the 200 classes of the ImageNet-DET dataset. The frames and annotations for each video are available at a rate of 25-30 FPS per video. Note that the ImageNet-DET dataset consists only of images rather than videos. We follow the widely used protocol [52,45,17,46,12] and train our model on the combination of ImageNet-VID and ImageNet-DET datasets using the 30 overlapping classes. We use a mini-batch size of 4, distributed to 4 NVIDIA A100 GPUs. We resize each image to have a shorter side of 600 pixels and train the model for a total of 12 epochs in 3 stages. We initialize the learning rate to 0.01 and reduce it by a factor of 10 at epochs 7 and 10. We test the model on ImageNet-VID val and report mAP@50 scores following the common practice.

We adopt Faster R-CNN [41] by replacing its ResNet backbone with the MDEQ-XL model. To incorporate multi-level representations, we also use a Feature Pyramid Network (FPN) [33] module after MDEQ-XL. Without any additional modifications, we directly use the model while keeping the model hyperparameters and remaining architecture details same as in Faster R-CNN models

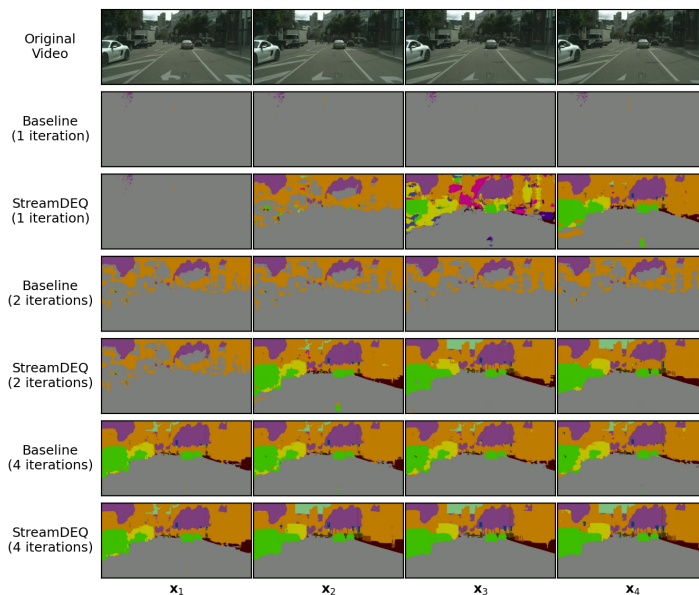


Fig. 6. Qualitative comparison of the baseline with StreamDEQ with different numbers of iterations on the Cityscapes dataset

for ResNet backbones. Exceptionally, we only modify the number of channels for the FPN module to match that of MDEQ-XL. We start training with the ImageNet pretrained MDEQ-XL model from MDEQ [4]. During training, we use 26 solver iterations for both forward and backward passes of the MDEQ, following the ImageNet classification experiments in MDEQ [4]. Unlike the video object detection models [52,17,12], we train our model in the causal single-frame setting, meaning we do not use any temporal information for improved training. Furthermore, we also incorporate the Jacobian regularization loss for MDEQs introduced by Bai et al. [6] to stabilize training.

Results. To the best of our knowledge, this is the first time that an implicit model (MDEQ) is used for an object detection task. We achieve 55.0 mAP@50 on ImageNet-VID val. We are aware that Faster R-CNN with ResNet-18 backbone yields 64.9 mAP@50, however, Faster R-CNN is highly optimized to perform well with ResNet backbones. Yet, we use this same setting with an MDEQ without any parameter optimization, as our focus is not on constructing a MDEQ-based state-of-the-art video object detector. We believe that there is room for improvement in detector design and tuning details, which we leave for future work.

Similar to the video segmentation task, we run StreamDEQ with different numbers of iterations. We present the results of this experiment in Figure 8. We observe the same trends with the segmentation task. Over time, detection performance increases and stabilizes at a value proportional to the number of Broyden




Fig. 7. mIoU results of StreamDEQ with shot changes from ImageNet-VID




Fig. 8. mAP@50 results of StreamDEQ for various number of iterations after initialization with zeros from the beginning of a clip on the ImageNet-VID dataset

iterations. The 1 and 2 iteration cases do not produce any detection results in the non-streaming mode, but if we perform 2 iterations with StreamDEQ, we improve the performance from 0 to 39.5 mAP@50 in 20 frames.

In addition, we also compare the inference speed of StreamDEQ with our baseline for different number of iterations in Table 1. In the 8 iteration case, we obtain a score of 54.8 with StreamDEQ which is only 0.2 lower than our baseline model, while being almost 3 times faster.

5 Conclusions & Future Work

In this paper, we proposed StreamDEQ, an efficient streaming video application of the multiscale implicit deep model, MDEQ. We showed that our model can start from scratch (*i.e.* all zeros) and efficiently update its representations to reach near-optimal representations as the video streams. We validated this claim on video semantic segmentation and video object detection tasks with thorough experiments. StreamDEQ presents a viable approach for both real-time video analysis and off-line large-scale methods. StreamDEQ is not specific to segmentation or object detection, and can be used as a drop-in replacement for most other structured prediction problems on streaming videos as long as the prediction task is performed by an implicit model. In addition, application to time series prediction and event camera based recognition tasks can be interesting future work directions.

Acknowledgments. The numerical calculations were partially performed at TUBITAK ULAKBIM, High Performance and Grid Computing Center (TRUBA) and METU Robotics and AI Technologies Research Center (ROMER) resources. Dr. Cinbis is supported by a Google Faculty Research Award. Dr. Akbas is supported by the BAGEP Award of the Science Academy, Turkey.

References

1. Ba, J., Mnih, V., Kavukcuoglu, K.: Multiple object recognition with visual attention. In: ICLR (2015)
2. Bai, S., Kolter, J.Z., Koltun, V.: Deep equilibrium models. In: Advances in Neural Information Processing Systems (NeurIPS) (2019)
3. Bai, S., Kolter, J.Z., Koltun, V.: Trellis networks for sequence modeling. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=HyeVtoRqtQ>
4. Bai, S., Koltun, V., Kolter, J.Z.: Multiscale deep equilibrium models. In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
5. Bai, S., Koltun, V., Kolter, J.Z.: Neural deep equilibrium solvers. In: International Conference on Learning Representations (2021)
6. Bai, S., Koltun, V., Kolter, J.Z.: Stabilizing equilibrium models by jacobian regularization. In: International Conference on Machine Learning (ICML) (2021)
7. Bazzani, L., de Freitas, N., Larochelle, H., Murino, V., Ting, J.A.: Learning attentional policies for tracking and recognition in video with deep networks. In: ICML (2011)
8. Bertasius, G., Torresani, L., Shi, J.: Object detection in video with spatiotemporal sampling networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 331–346 (2018)
9. Broyden, C.G.: A class of methods for solving nonlinear simultaneous equations. *Mathematics of computation* **19**(92), 577–593 (1965)
10. Carreira, J., Patraucean, V., Mazare, L., Zisserman, A., Osindero, S.: Massively parallel video networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 649–666 (2018)
11. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: MMDetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 (2019)
12. Chen, Y., Cao, Y., Hu, H., Wang, L.: Memory enhanced global-local aggregation for video object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10337–10346 (2020)
13. Contributors, M.: MMTracking: OpenMMLab video perception toolbox and benchmark. <https://github.com/open-mmlab/mtracking> (2020)
14. Cordonnier, J.B., Mahendran, A., Dosovitskiy, A., Weissenborn, D., Uszkoreit, J., Unterthiner, T.: Differentiable patch selection for image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2351–2360 (2021)
15. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
16. Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., Kaiser, L.: Universal transformers. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=HyzdRiR9Y7>
17. Deng, J., Pan, Y., Yao, T., Zhou, W., Li, H., Mei, T.: Relation distillation networks for video object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7023–7032 (2019)

18. Denil, M., Bazzani, L., Larochelle, H., de Freitas, N.: Learning where to attend with deep architectures for image tracking. *Neural computation* **24**(8), 2151–2184 (2012)
19. Gadde, R., Jampani, V., Gehler, P.V.: Semantic video cnns through representation warping. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 4453–4462 (2017)
20. Han, W., Khorrani, P., Paine, T.L., Ramachandran, P., Babaeizadeh, M., Shi, H., Li, J., Yan, S., Huang, T.S.: Seq-nms for video object detection. *arXiv preprint arXiv:1602.08465* (2016)
21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
22. Horn, B.K., Schunck, B.G.: Determining optical flow. *Artificial intelligence* **17**(1-3), 185–203 (1981)
23. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)
24. Hu, P., Caba, F., Wang, O., Lin, Z., Sclaroff, S., Perazzi, F.: Temporally distributed networks for fast video semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8818–8827 (2020)
25. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4700–4708 (2017)
26. Huang, P.Y., Hsu, W.T., Chiu, C.Y., Wu, T.F., Sun, M.: Efficient uncertainty estimation for semantic segmentation in videos. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 520–535 (2018)
27. Huang, Z., Bai, S., Kolter, J.Z.: (Implicit)²: Implicit layers for implicit representations. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2021)
28. Jain, S., Wang, X., Gonzalez, J.E.: Accel: A corrective fusion network for efficient semantic segmentation on video. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8866–8875 (2019)
29. Kang, K., Li, H., Yan, J., Zeng, X., Yang, B., Xiao, T., Zhang, C., Wang, Z., Wang, R., Wang, X., et al.: T-cnn: Tubelets with convolutional neural networks for object detection from videos. *IEEE Transactions on Circuits and Systems for Video Technology* **28**(10), 2896–2907 (2017)
30. Kang, K., Ouyang, W., Li, H., Wang, X.: Object detection from video tubelets with convolutional neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 817–825 (2016)
31. Li, M., Wang, Y.X., Ramanan, D.: Towards streaming perception. In: *European Conference on Computer Vision*. pp. 473–488. Springer (2020)
32. Li, Y., Shi, J., Lin, D.: Low-latency video semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5997–6005 (2018)
33. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2117–2125 (2017)
34. Liu, M., Zhu, M.: Mobile video object detection with temporally-aware feature maps. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5686–5695 (2018)

35. Liu, M., Zhu, M., White, M., Li, Y., Kalenichenko, D.: Looking fast and slow: Memory-guided mobile video object detection. arXiv preprint arXiv:1903.10172 (2019)
36. Liu, Y., Shen, C., Yu, C., Wang, J.: Efficient semantic video segmentation with per-frame inference. In: European Conference on Computer Vision. pp. 352–368. Springer (2020)
37. Liu, Z., Wang, H.J., Xu, Z., Darrell, T., Shelhamer, E.: Confidence adaptive any-time pixel-level recognition. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=kNKFOX1euC>
38. Mnih, V., Heess, N., Graves, A., Kavukcuoglu, K.: Recurrent models of visual attention. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems. vol. 27. Curran Associates, Inc. (2014), <https://proceedings.neurips.cc/paper/2014/file/09c6c3783b4a70054da74f2538ed47c6-Paper.pdf>
39. Narayanan, D., Harlap, A., Phanishayee, A., Seshadri, V., Devanur, N.R., Ganger, G.R., Gibbons, P.B., Zaharia, M.: Pipedream: generalized pipeline parallelism for dnn training. In: Proceedings of the 27th ACM Symposium on Operating Systems Principles. pp. 1–15 (2019)
40. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems 32, pp. 8024–8035. Curran Associates, Inc. (2019), <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
41. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. Advances in neural information processing systems **28** (2015)
42. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV) **115**(3), 211–252 (2015). <https://doi.org/10.1007/s11263-015-0816-y>
43. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018)
44. Shelhamer, E., Rakelly, K., Hoffman, J., Darrell, T.: Clockwork convnets for video semantic segmentation. In: European Conference on Computer Vision. pp. 852–868. Springer (2016)
45. Wang, S., Zhou, Y., Yan, J., Deng, Z.: Fully motion-aware network for video object detection. In: Proceedings of the European conference on computer vision (ECCV). pp. 542–557 (2018)
46. Wu, H., Chen, Y., Wang, N., Zhang, Z.: Sequence level semantics aggregation for video object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9217–9225 (2019)
47. Xu, Y.S., Fu, T.J., Yang, H.K., Lee, C.Y.: Dynamic video segmentation network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6556–6565 (2018)
48. Zagoruyko, S., Komodakis, N.: Wide residual networks. arXiv preprint arXiv:1605.07146 (2016)

49. Zhao, B., Zhao, B., Tang, L., Han, Y., Wang, W.: Deep spatial-temporal joint feature representation for video object detection. *Sensors* **18**(3), 774 (2018)
50. Zhu, H., Wei, H., Li, B., Yuan, X., Kehtarnavaz, N.: A review of video object detection: Datasets, metrics and methods. *Applied Sciences* **10**(21), 7834 (2020)
51. Zhu, X., Dai, J., Yuan, L., Wei, Y.: Towards high performance video object detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7210–7218 (2018)
52. Zhu, X., Wang, Y., Dai, J., Yuan, L., Wei, Y.: Flow-guided feature aggregation for video object detection. In: *Proceedings of the IEEE international conference on computer vision*. pp. 408–417 (2017)
53. Zhu, X., Xiong, Y., Dai, J., Yuan, L., Wei, Y.: Deep feature flow for video recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2349–2358 (2017)

S1 Effect of Shot Changes

In this section, we show the shot change experiment where the connected clips are both from the Cityscapes dataset. This can be considered as a more simple version of the experiment explained in the main paper since, in this case, the context of the clips are more similar and therefore, the representations of two different video frames are likely to be more similar compared to the cross-dataset experiments. We demonstrate the results of this experiment in Figure S1. We notice that, for shot changes in similar contexts, Cityscapes-to-Cityscapes, the mIoU scores on initial frames are higher than our previous experiment in Figure 5 (right) and also higher than the ImageNet-VID to Cityscapes shot change scenario in Figure 7 in the main paper. Additionally, after the first few frames, the mIoU scores again get close to the values presented in Figure 5 (right) and Figure 7 in the main paper. Overall the results confirm that StreamDEQ is able to adapt the representation to sudden changes in the video.




Fig. S1. mIoU results of StreamDEQ with shot changes from Cityscapes