
MMC Transformer: Multiscale Multigrid Comparator Transformer for Few-Shot Video Segmentation

Supplementary Material

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Our supplemental material is organized into six major sections. Section 1 docu-
2 ments an associated supplementary video. Section 2 provides additional exper-
3 imental design and implementation details that were not described in the main
4 submission. Section 3 discusses additional results in terms of an ablation study on
5 the number of memory entries used and the attention maps visualised along with a
6 frequency analysis that motivates the multigrid formulation. Section 4 documents
7 the attached code and how to reproduce YouTube-VIS FSVOS experiments. Fi-
8 nally, Sections 5 and 6 provide the licenses of the different assets and discusses the
9 societal impact of our downstream task.

10 1 Supplemental video

11 We include an accompanying video as part of the supplementary materials. In this video, we show
12 eight examples for few-shot video object segmentation on YouTube-VIS and three examples for
13 few-shot common action localisation on A2D. Additionally, we show the attention maps output from
14 the different memory entries for the finest scale across the entire input clip. The video is in MP4
15 format and is approximately eight minutes long. Layouts for each sampling pair are described in
16 detail followed by the example video samples. The codec used for the realization of the provided
17 video is H.264 (x264).

18 2 Experiment design details

19 In this section, we detail the experiment setup and implementation details that were not described in
20 the main submission.

21 **Evaluation.** In the few-shot video object segmentation task we use the YouTube-VIS [3] dataset that
22 is split into four folds, each with 30 training classes and 10 testing classes. In the few-shot common
23 action localisation we use Common A2D [9] that has 43 classes split into 33, five and five classes
24 for the training, validation and testing splits, resp. We choose as a baseline comparison algorithm a
25 previous approach that made use of multiscale processing relying on correlation tensors between the
26 support and query sets using features from different scales [5]; however, our baseline does not use a
27 transformer based approach. We follow previous evaluations by reporting the mean intersection over
28 union (mIoU) for one-way one-shot and five-shot evaluations [3, 9]. Evaluation is performed as the
29 average over five runs for YouTube-VIS FS-VOS, similar to [3]. For the evaluation on Common A2D
30 we report results from one run, following previous work [9].

Method	mIoU				
	1	2	3	4	Mean
$N = 2$	49.6	70.3	62.9	65.1	61.98
$N = 20$	51.5	70.6	63.0	64.6	62.4

Table 1: Ablation study on the number of memory entries in our multiscale memory learning mechanism. This study is performed on YouTube-VIS FS-VOS folds 1, 2, 3 and 4 with a five shot support set.

31 **Implementation details.** We follow standard few-shot segmentation and few-shot video segmentation
32 practice of assigning the novel class objects that exist in training images to background [8]. During
33 training on YouTube-VIS FS-VOS, we randomly sample five frames from the query set video and
34 randomly sample one image from another video for the support set. In Common A2D, we use 25
35 frame trimmed clips subsampled uniformly into five frames for the support and query sets during
36 training. During inference, in YouTube-VIS FS-VOS we infer over the entire video. While in
37 Common A2D during inference, we use 25 frame clips in a temporal sliding window over the entire
38 untrimmed query video. The features that are used to compute the correlation tensors in both tasks
39 are extracted as follows. We extract features at the end of each bottleneck before the ReLU activation
40 in the last three stages and concatenate the correlation tensors (with same spatial dimensions) along
41 the channel dimension; this operation results in a three scale hypercorrelation pyramid, $P = 3$. Our
42 MMC transformer is constructed as seven decoding layers with separate weights each corresponding
43 to one communication exchange across scale with bidirectional exchange. We go through coarse-
44 mid-fine-mid-coarse-mid-fine in that order, which results in the final enhanced feature maps at the
45 finest scale. Finally, we train our models on a single NVIDIA RTX A6000 GPU using batch size 19.

46 3 Additional empirical results

47 In this section, we provide additional empirical results to understand the multiscale memory learning
48 and confirm the motivation of our proposed multigrid formulation.

49 3.1 Memory ablation

50 We performed an ablation study on the number of memory entries, N , used in our multiscale memory
51 learning technique. In the main submission, all the experiments, including the ablation on the *Stacked*
52 *vs. Multigrid* formulation, were conducted using $N = 20$. In Table 1, we compare both $N = 2$ that
53 denotes the novel class and background *vs.* using an overcomplete set with more than two, $N = 20$.
54 It is seen that over all folds, and especially the first two folds, that using an overcomplete set improves
55 the separation between the background and novel class leading to better segmentation accuracy. The
56 overcomplete set captures different parts of the novel class and background and thereby improves the
57 segmentation, as described in the following section.

58 3.2 Memory attention maps

59 In Fig. 1, we visualise the attention maps from our multiscale memory learning on the finest scale
60 after the full bidirectional multigrid formulation is finalized on three randomly sampled frames from
61 YouTube-VIS video. These visualizations allow us to gain better insights on how our multiscale
62 memory learning enhances the output feature maps. Since we use $N = 20$ we show the attention
63 maps of the different memory entries along with the original image and the predicted segmentation
64 (highlighted in red). It is seen that our multiscale memory learning attends to different parts of the
65 novel class and the background. It also shows the temporal consistency of the attended parts, for
66 example memory entry 16 tends to attend to the novel class across all frames, while memory entry
67 three attends to parts of the background. Thus, it generally confirms the benefit from our multiscale
68 memory learning in enhancing the query feature maps and in separating the novel class from the
69 background for better segmentation accuracy.

70 Additionally, Fig. 2 shows the output attention maps from the mid levels during the bidirectional
71 information exchange for one randomly sampled frame. Notably, the bidirectional information
72 exchange with separate decoder layers entails that each layer captures a different aspect of the input

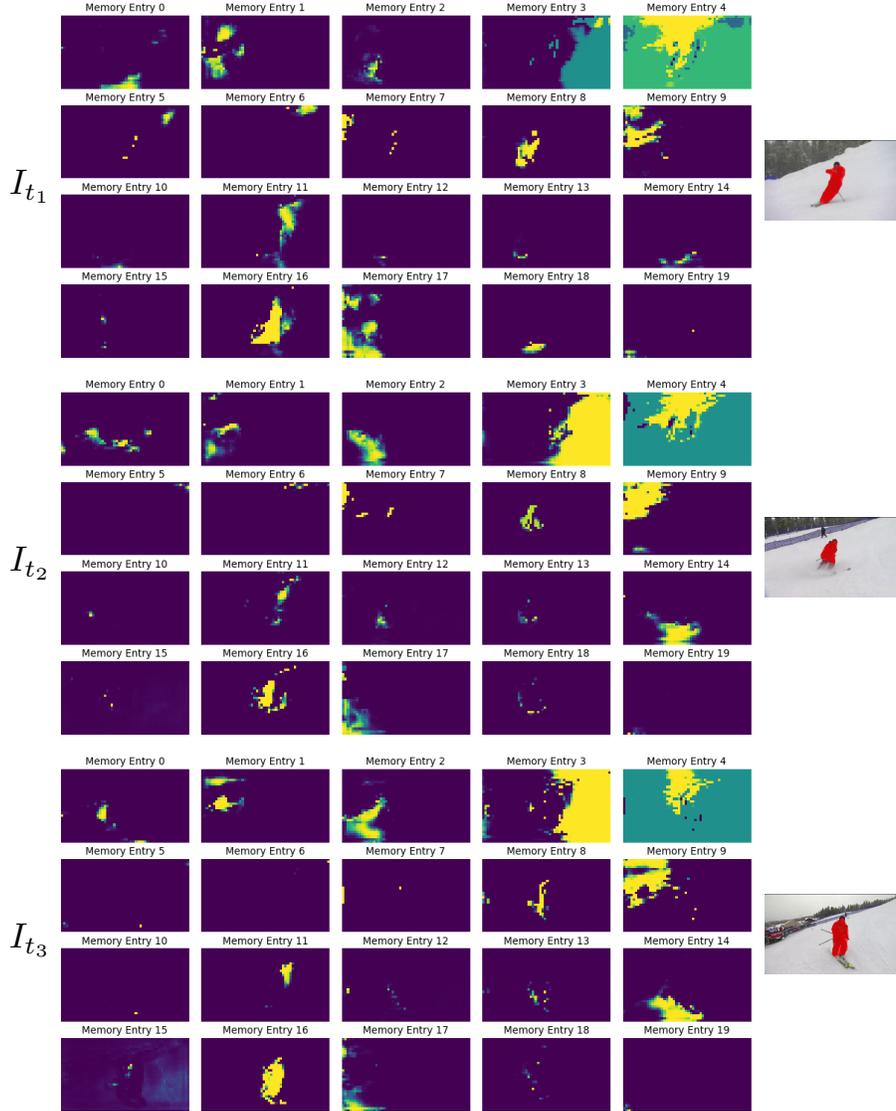


Figure 1: Visualisation of learned memory attention maps. **Mid:** Visualisation of the output attention maps in the finest scale after the full bidirectional multigrid formulation for the memory entries in the learned memory with $N = 20$ for three randomly sampled frames, $I_{t_1}, I_{t_2}, I_{t_3}$. Attention is visualised as a heatmap where yellow indicates higher attention. **Right:** Input image with the segmentation prediction highlighted in red.

73 even within the same scale. This fact suggests that the different decoder layers with separate learnable
 74 weights will capitalize on different components of the input imagery. Here, we focus on what is being
 75 captured in the mid level with each exchange of information in the coarse-mid-fine-mid-coarse-mid-
 76 fine scheme (*i.e.* three times communicating with the mid scale). It is seen that with each exchange
 77 of information focus is on a separate component of the input, where the early exchange, p_1 , focuses
 78 on the hard negatives around the novel class that can be confused with the background (*e.g.* memory
 79 entries five to seven). The second exchange, p_2 , captures better the novel classes (*e.g.* memory entry
 80 18), and finally the last encounter, p_3 , captures the highly confident parts of the novel class (*e.g.*
 81 memory entry 15).

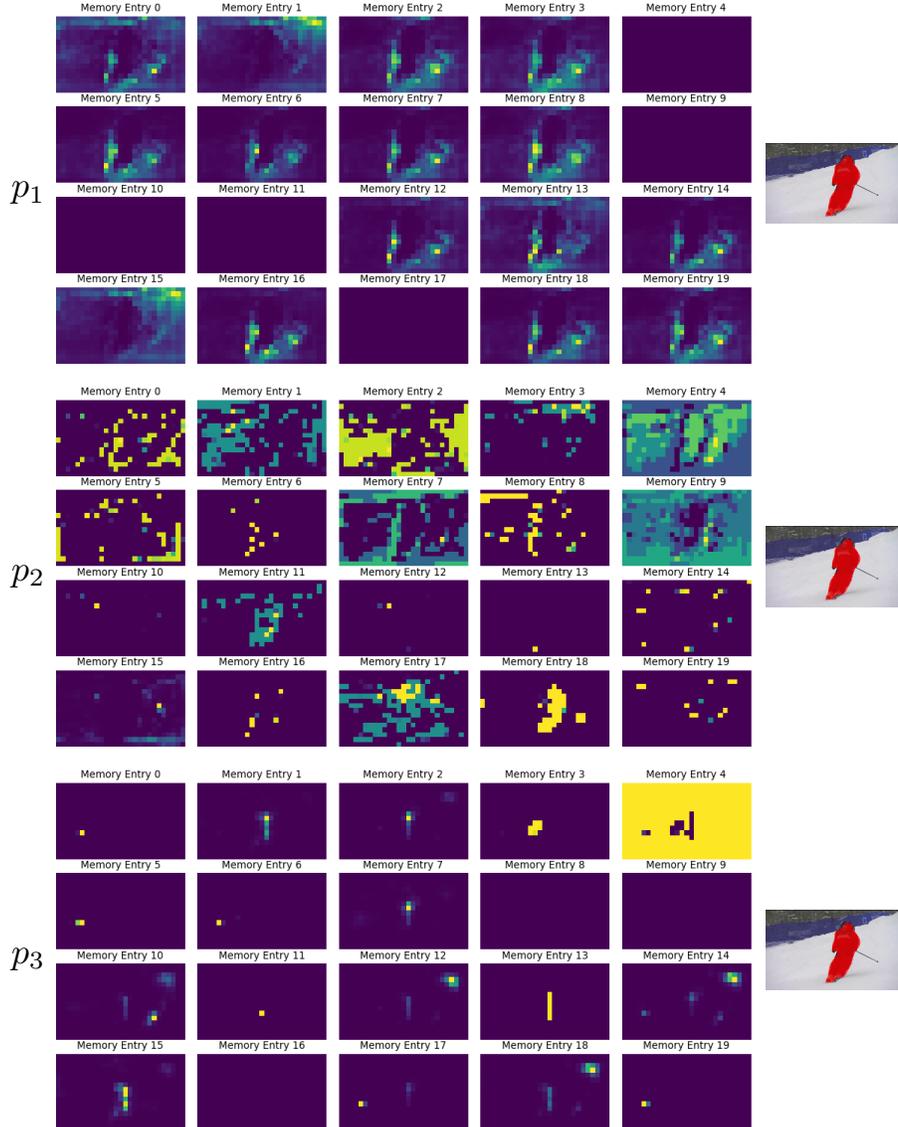


Figure 2: Visualisation of the output attention maps from our learned memory for the mid scale for each information exchange, $p_{1,2,3}$ in the bidirectional communication. Attention is visualised as a heatmap, where yellow indicates higher attention. **Right:** Input image with the segmentation prediction highlighted in red.

82 3.3 Frequency analysis

83 In the main submission, we hypothesize that one of the reasons behind our multigrid formulation
 84 benefit is related to the fact that different feature abstraction levels capture different frequency
 85 components, *cf.* [4]. The bidirectional multigrid formulation allows for a better information exchange
 86 across scales and hence better captures the different frequency components required for delineating
 87 the object boundaries. Inspired by previous work that analyzed transformers from a frequency domain
 88 perspective [1], we conduct a frequency analysis where we perform low pass filtering with a Gaussian
 89 kernel on the input images with different filter sizes and standard deviations, as shown in Fig. 3.
 90 We then evaluate our final predictions similar to the original setup on the input filtered images and
 91 average per filter size over the four folds on YouTube-VIS. In Fig. 3 across the early three filter sizes
 92 our approach can better cope with frequency band-passed input compared to the baseline. These

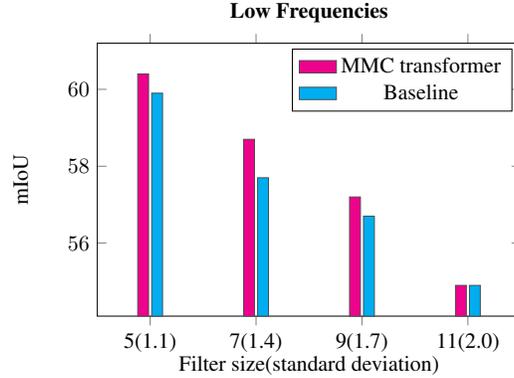


Figure 3: Evaluation of mIoU averaged over the four folds of our MMC transformer vs. the multiscale hypercorrelation squeeze network baseline [5] with input images that have gone through low pass filtering with different filter sizes and standard deviations.

93 results show that MMC transformer is more robust than the baseline to restricted frequency content,
 94 where here the restricted frequency content comes about through low pass filtering.

95 4 Code

96 We include our code as part of the supplementary material with a README file that can reproduce
 97 our experiments in the main submission, along with our trained weights to reproduce inference.

98 5 Assets and licenses

99 We use the YouTube-VIS¹ dataset which is under a Creative Commons Attribution 4.0 License that
 100 allows non commercial research use. We also used the A2D² dataset where the license states that the
 101 dataset may not be republished in any form without the written consent of the authors.

102 6 Societal impact

103 Few-shot video object segmentation, where the query set to be segmented is a video, is a crucial task
 104 that can help reduce the annotation cost required to label large-scale video datasets. It can serve a
 105 variety of applications in autonomous systems [2] and medical image processing [6] which require the
 106 model to learn from few labelled examples for novel classes that are beyond the closed set of training
 107 classes with abundant labels. It can also help bridge the gap between developing and developed
 108 countries, where the former lacks the resources necessary to annotate large-scale labelled datasets
 109 that are required in a variety of tasks that serves the community such as, the use of satellite imagery in
 110 agricultural monitoring and crop management [7]. We believe our work in general provides positive
 111 impact in empowering developing countries to establish labelled datasets that satisfy the needs of
 112 their own communities rather than following public benchmarks.

113 However, as with many artificial intelligence algorithms, video object segmentation can have negative
 114 societal impacts, *e.g.* through application to automatic target detection in military and surveillance
 115 systems. There are emerging movements to limit such applications, *e.g.* pledges on the part of
 116 researchers to band use of artificial intelligence in weaponry systems. We have participated in signing
 117 that pledge and are supporters of its enforcement through international laws. Nonetheless, we strongly
 118 believe these misuses are available in both few-shot and non few-shot methods and are not tied to
 119 the specific few-shot case. On the contrary, we argue that empowering developing countries towards
 120 decolonizing artificial intelligence can help go beyond centered power that currently lies within
 121 developed countries and big tech companies and is guided solely by their interests.

¹<https://youtube-vos.org/dataset/>

²<https://web.eecs.umich.edu/~jjcorso/r/a2d/>

References

- 122
- 123 [1] Jiawang Bai, Li Yuan, Shu-Tao Xia, Shuicheng Yan, Zhifeng Li, and Wei Liu. Improving vision
124 transformers by revisiting high-frequency components. *arXiv preprint arXiv:2204.00993*, 2022.
- 125 [2] Jun Cen, Peng Yun, Junhao Cai, Michael Yu Wang, and Ming Liu. Deep metric learning for
126 open world semantic segmentation. In *Proceedings of the IEEE International Conference on*
127 *Computer Vision*, pages 15333–15342, 2021.
- 128 [3] Haoxin Chen, Hanjie Wu, Nanxuan Zhao, Sucheng Ren, and Shengfeng He. Delving deep into
129 many-to-many attention for few-shot video object segmentation. In *Proceedings of the IEEE*
130 *Conference on Computer Vision and Pattern Recognition*, pages 14040–14049, 2021.
- 131 [4] Isma Hadji and Richard P Wildes. Why convolutional networks learn oriented bandpass filters:
132 Theory and empirical support. *arXiv preprint arXiv:2011.14665*, 2020.
- 133 [5] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation.
134 In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6941–6952,
135 2021.
- 136 [6] Tobias Ross, Annika Reinke, Peter M Full, Martin Wagner, Hannes Kenngott, Martin Apitz,
137 Hellena Hempe, Diana Mindroc Filimon, Patrick Scholz, Thuy Nuong Tran, et al. Robust medical
138 instrument segmentation challenge 2019. *arXiv preprint arXiv:2003.10299*, 2020.
- 139 [7] Joel Segarra, Maria Luisa Buchaillot, Jose Luis Araus, and Shawn C Kefauver. Remote sensing
140 for precision agriculture: Sentinel-2 improved features and applications. *Agronomy*, 10(5):641,
141 2020.
- 142 [8] Amirreza Shaban, Zhen Bansal, Shrayand Liu, Irfan Essa, and Byron Boots. One-shot learning
143 for semantic segmentation. In *Proceedings of the British Machine Vision Conference*, pages
144 167.1–167.13, 2017.
- 145 [9] Pengwan Yang, Pascal Mettes, and Cees GM Snoek. Few-shot transformation of common actions
146 into time and space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
147 *Pattern Recognition*, pages 16031–16040, 2021.