

SUPPLEMENTARY MATERIALS

FILM: Frame Interpolation for Large Motion

Fitsum Reda¹, Janne Kontkanen¹, Eric Tabellion¹, Deqing Sun¹,
Caroline Pantofaru¹, and Brian Curless^{1,2}

¹Google Research

²University of Washington

We provide additional implementation details and a supplementary video (visit <https://film-net.github.io>) for a quick overview of our paper, the motivations, as well as more visual results.

1 Implementation Details

1.1 Loss Combination Weights

Our style loss (\mathcal{L}_S) optimally combines \mathcal{L}_1 , \mathcal{L}_{VGG} , and \mathcal{L}_{Gram} . We use a piece-wise linear weight-schedule to select a weight for each loss at each iteration. Specifically, we assign weights of (1.0, 1.0, 0.0) for 1.5M iterations, and weights of (1.0, 0.25, 40.0) for the last 1.5M iterations. For the last 1.5M iterations, the loss weights are empirically selected such that each loss contributes equally to the combined Style loss.

1.2 SoftSplat Implementation

As described in the paper, we have implemented the SoftSplat[18] in Tensorflow 2, using the author’s tuned hyper-parameters, and we have been able to replicate their published benchmark scores. Note that, SoftSplat³ provides only a CuPy implementation of the softmax splatting operator.

In Figure 1, we show our faithful implementations on a DAVIS [34] image sample rendered in [18]. We found that renderings with our implementation to be quite comparable to the ones provided in the original paper.

2 Supplementary Video

Please watch the enclosed video or visit <https://film-net.github.io>. We have included motivations, illustrations of our methods, and more visual results and failure samples.

³ <https://github.com/sniklaus/softmax-splatting>



Fig. 1. Frame interpolation example from DAVIS-dataset [24] featuring a walking flamingo. From left to right: The input frames overlaid, SoftSplat- L_1 [18], SoftSplat- L_F [18], ABME [23], our FILM- L_1 , and our FILM- L_S . FILM- L_S produces crisp frame, while color distortions and transparencies can be seen in ABME and SoftSplat, respectively. SoftSplat renderings are generated from our faithful implementations, which we found to be quite comparable to the original paper [18].