Stable Long-Term Recurrent Video Super-Resolution
Supplementary Material

Benjamin Naoto Chiche [1,2], Arnaud Woiselle [1], Joana Frontera-Pons [2,3], Jean-Luc Starck [2]

[1] Safran Electronics & Defense, F-91344 Massy, France

[2] AIM, CEA, CNRS, Université Paris-Saclay, Université Paris Cité, F-91191 Gif-sur-Yvette, France

[3] DR2I, Institut Polytechnique des Sciences Avancées, F-94200 Ivry-sur-Seine, France

benjamin.chiche@safrangroup.com, arnaud.woiselle@safrangroup.com

joana.frontera-pons@cea.fr, https://orcid.org/0000-0003-2177-7794

1. Discussion

Given an existing pre-trained recurrent model, one way to operate it in a stable manner would be the following: chunking the long sequence into shorter, either overlapping or not overlapping sequences and using the model for each chunk. The recurrent features are reset to zero between each chunk. In our study, we do not consider this method for the following reasons. If the chunks do not overlap, this approach results in a severe visual flickering artifact between each pair of chunks, due to the reinitialization of recurrent features that makes the model go through a new burn-in period. If the chunks overlap, on the one hand, this still gives visual flickering artifacts due to discontinuities of recurrent features. On the other hand, the method becomes computationally redundant and inefficient, suppressing one of the main advantages of using a recurrent model. Computation time is doubled at overlapped regions and memory consumption is doubled.

2. Code, data and other materials

Please find our code and instructions in the following GitHub repository: https://github.com/bjmch/MRVSR. The repository also contains links to download the proposed dataset, network weights and videos reconstructed from different networks.