



MarioNette: Self-Supervised Sprite Learning

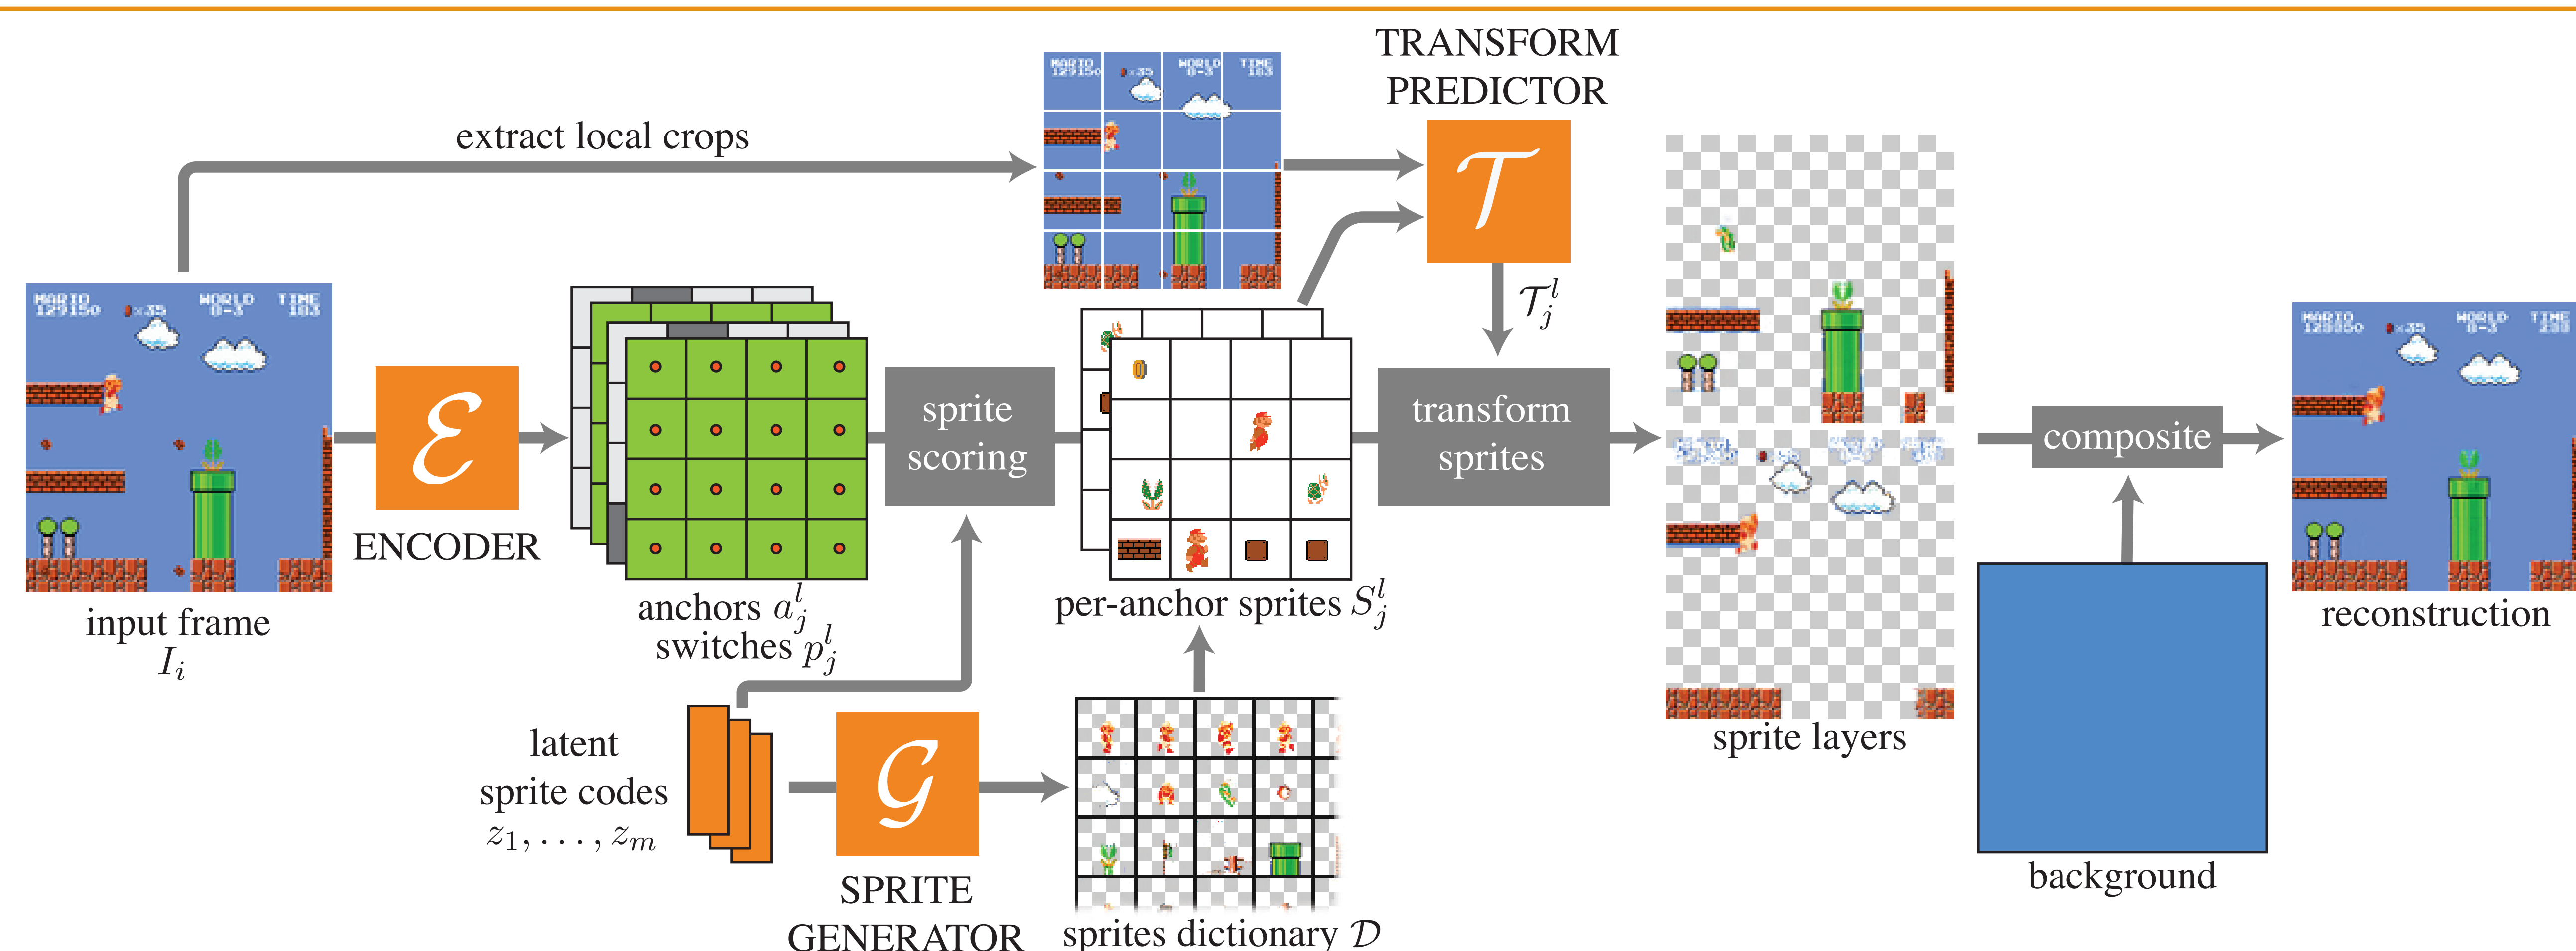
Dmitriy Smirnov¹ Michaël Gharbi² Matthew Fisher² Vitor Guizilini³ Alexei A. Efros⁴ Justin Solomon¹
¹Massachusetts Institute of Technology ²Adobe Research ³Toyota Research Institute ⁴UC Berkeley



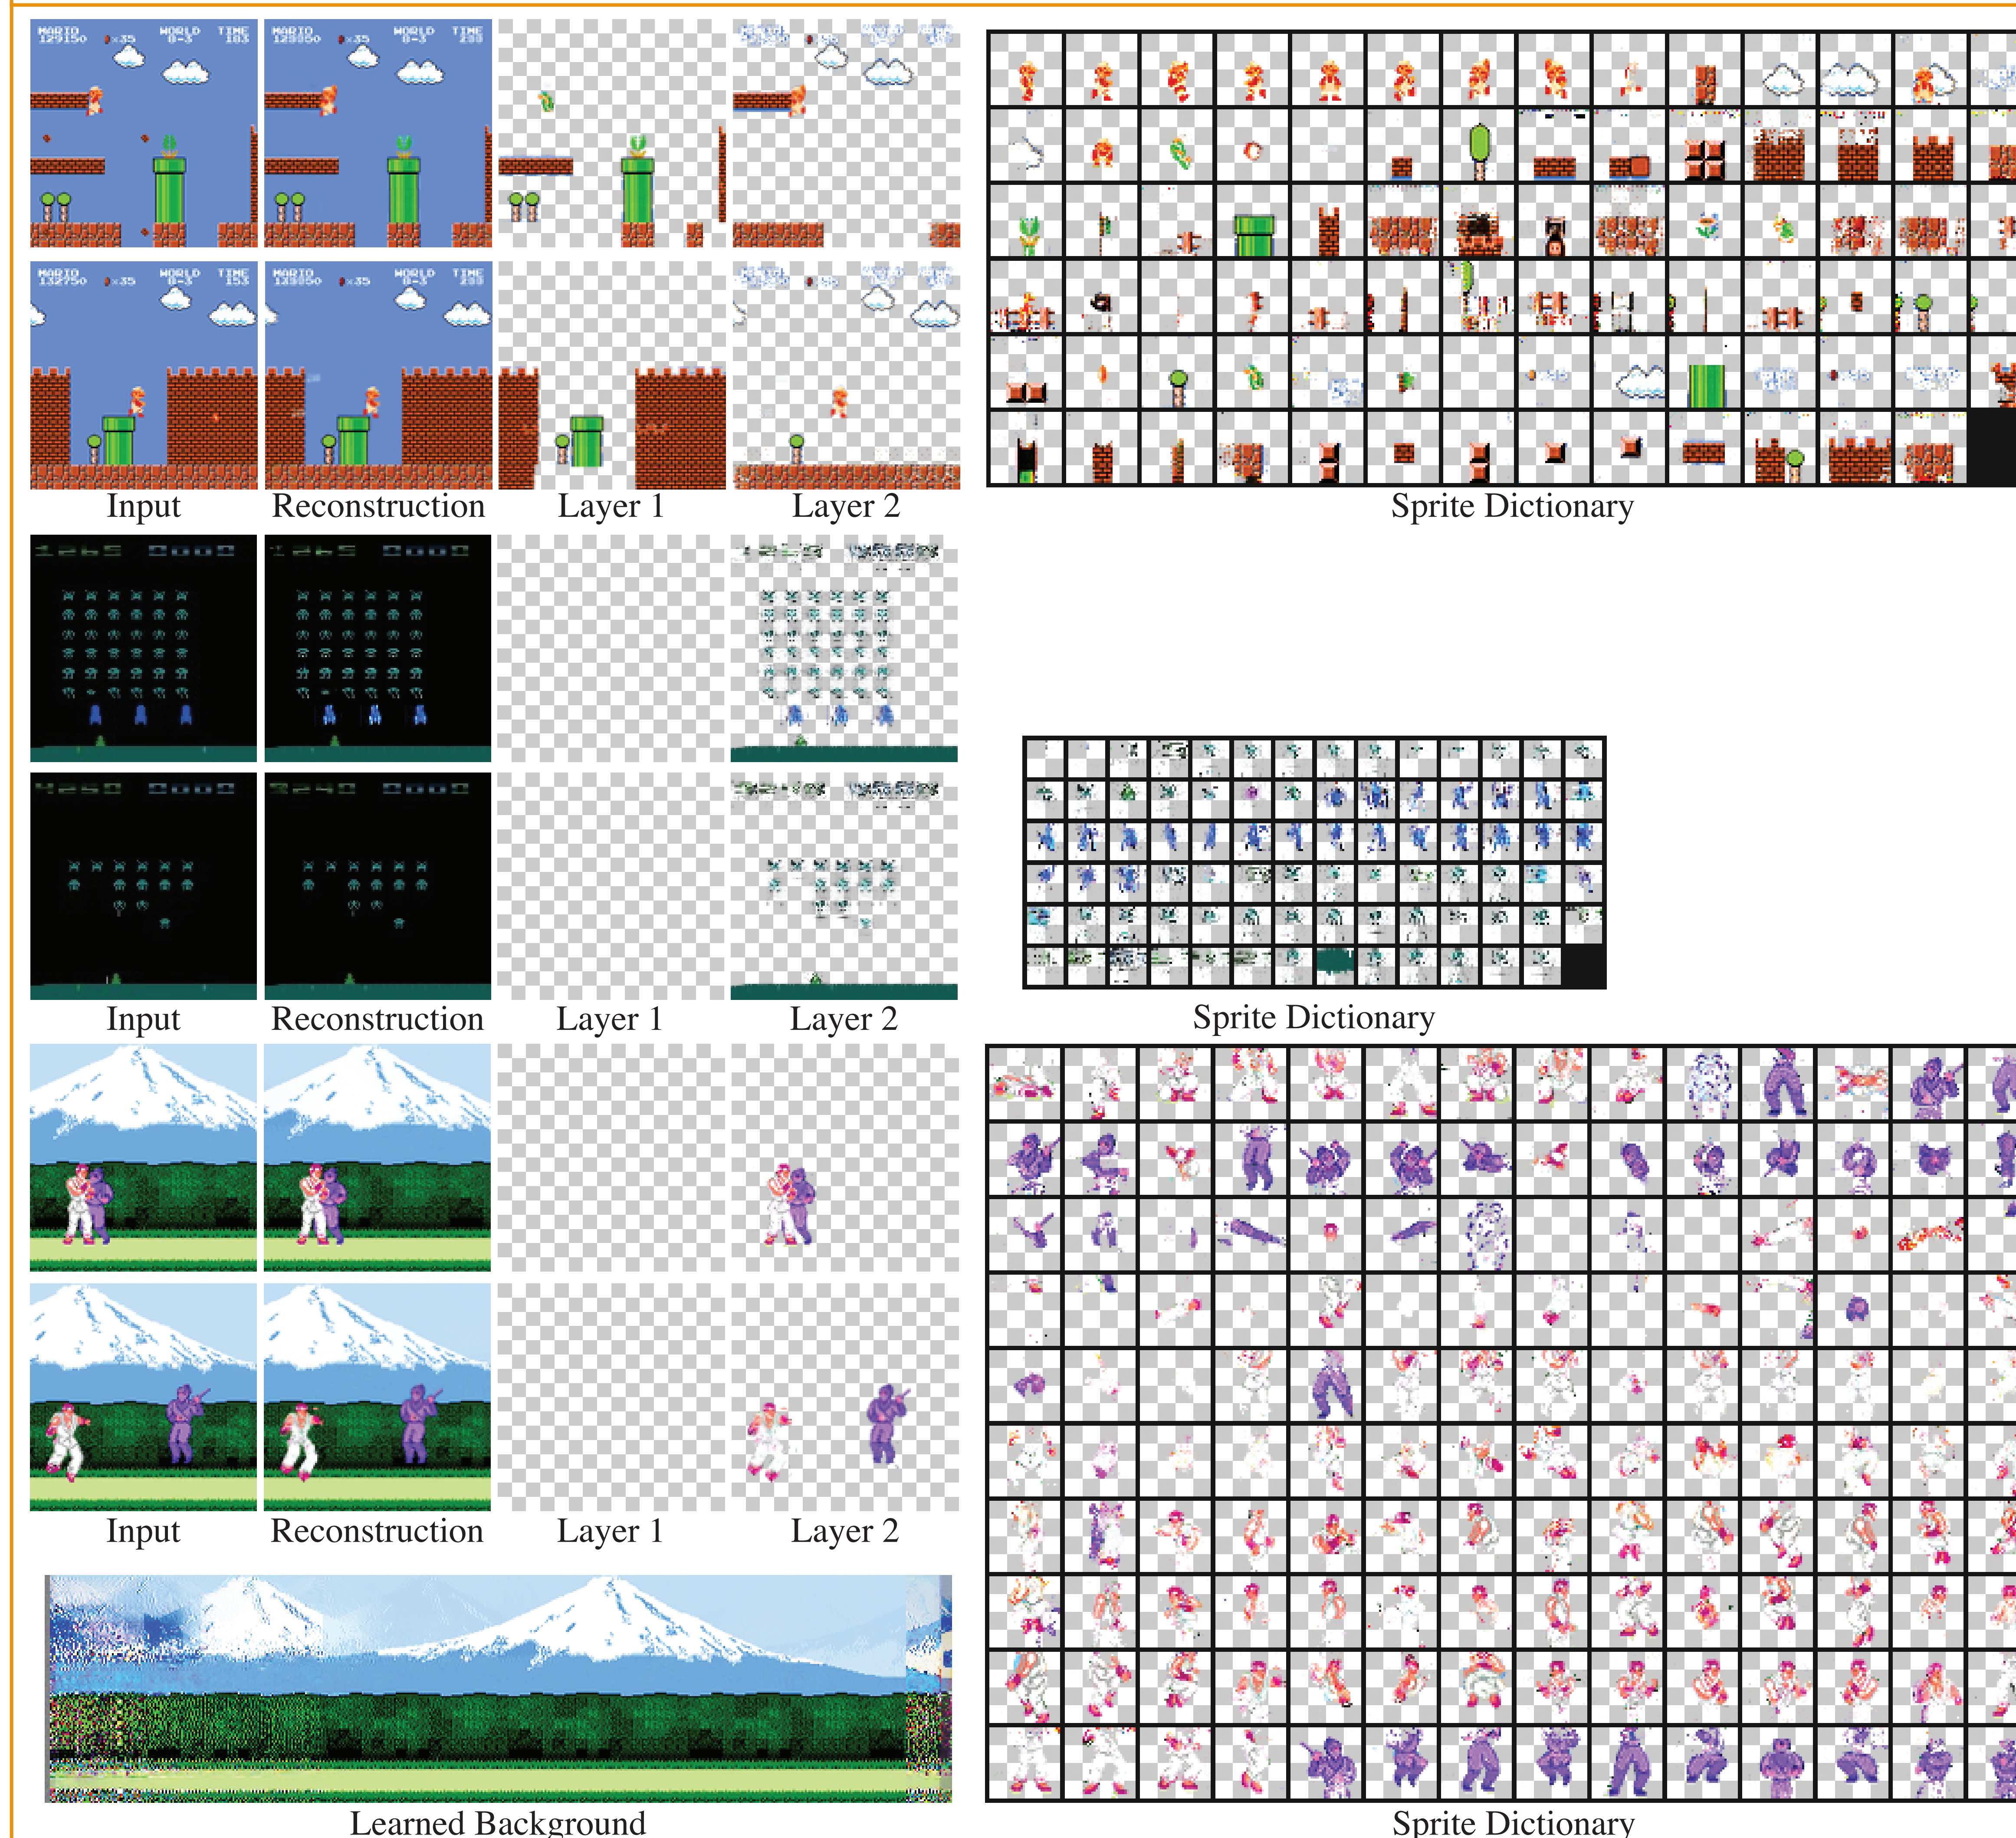
SUMMARY

We propose a self-supervised deep learning approach that decomposes sprite-based video animations into a disentangled representation of recurring graphic elements. By jointly learning a dictionary of patches and a model that places them onto a canvas, we deconstruct sprite-based content into a consistent and explicit representation that can be easily used in downstream tasks, like editing or analysis.

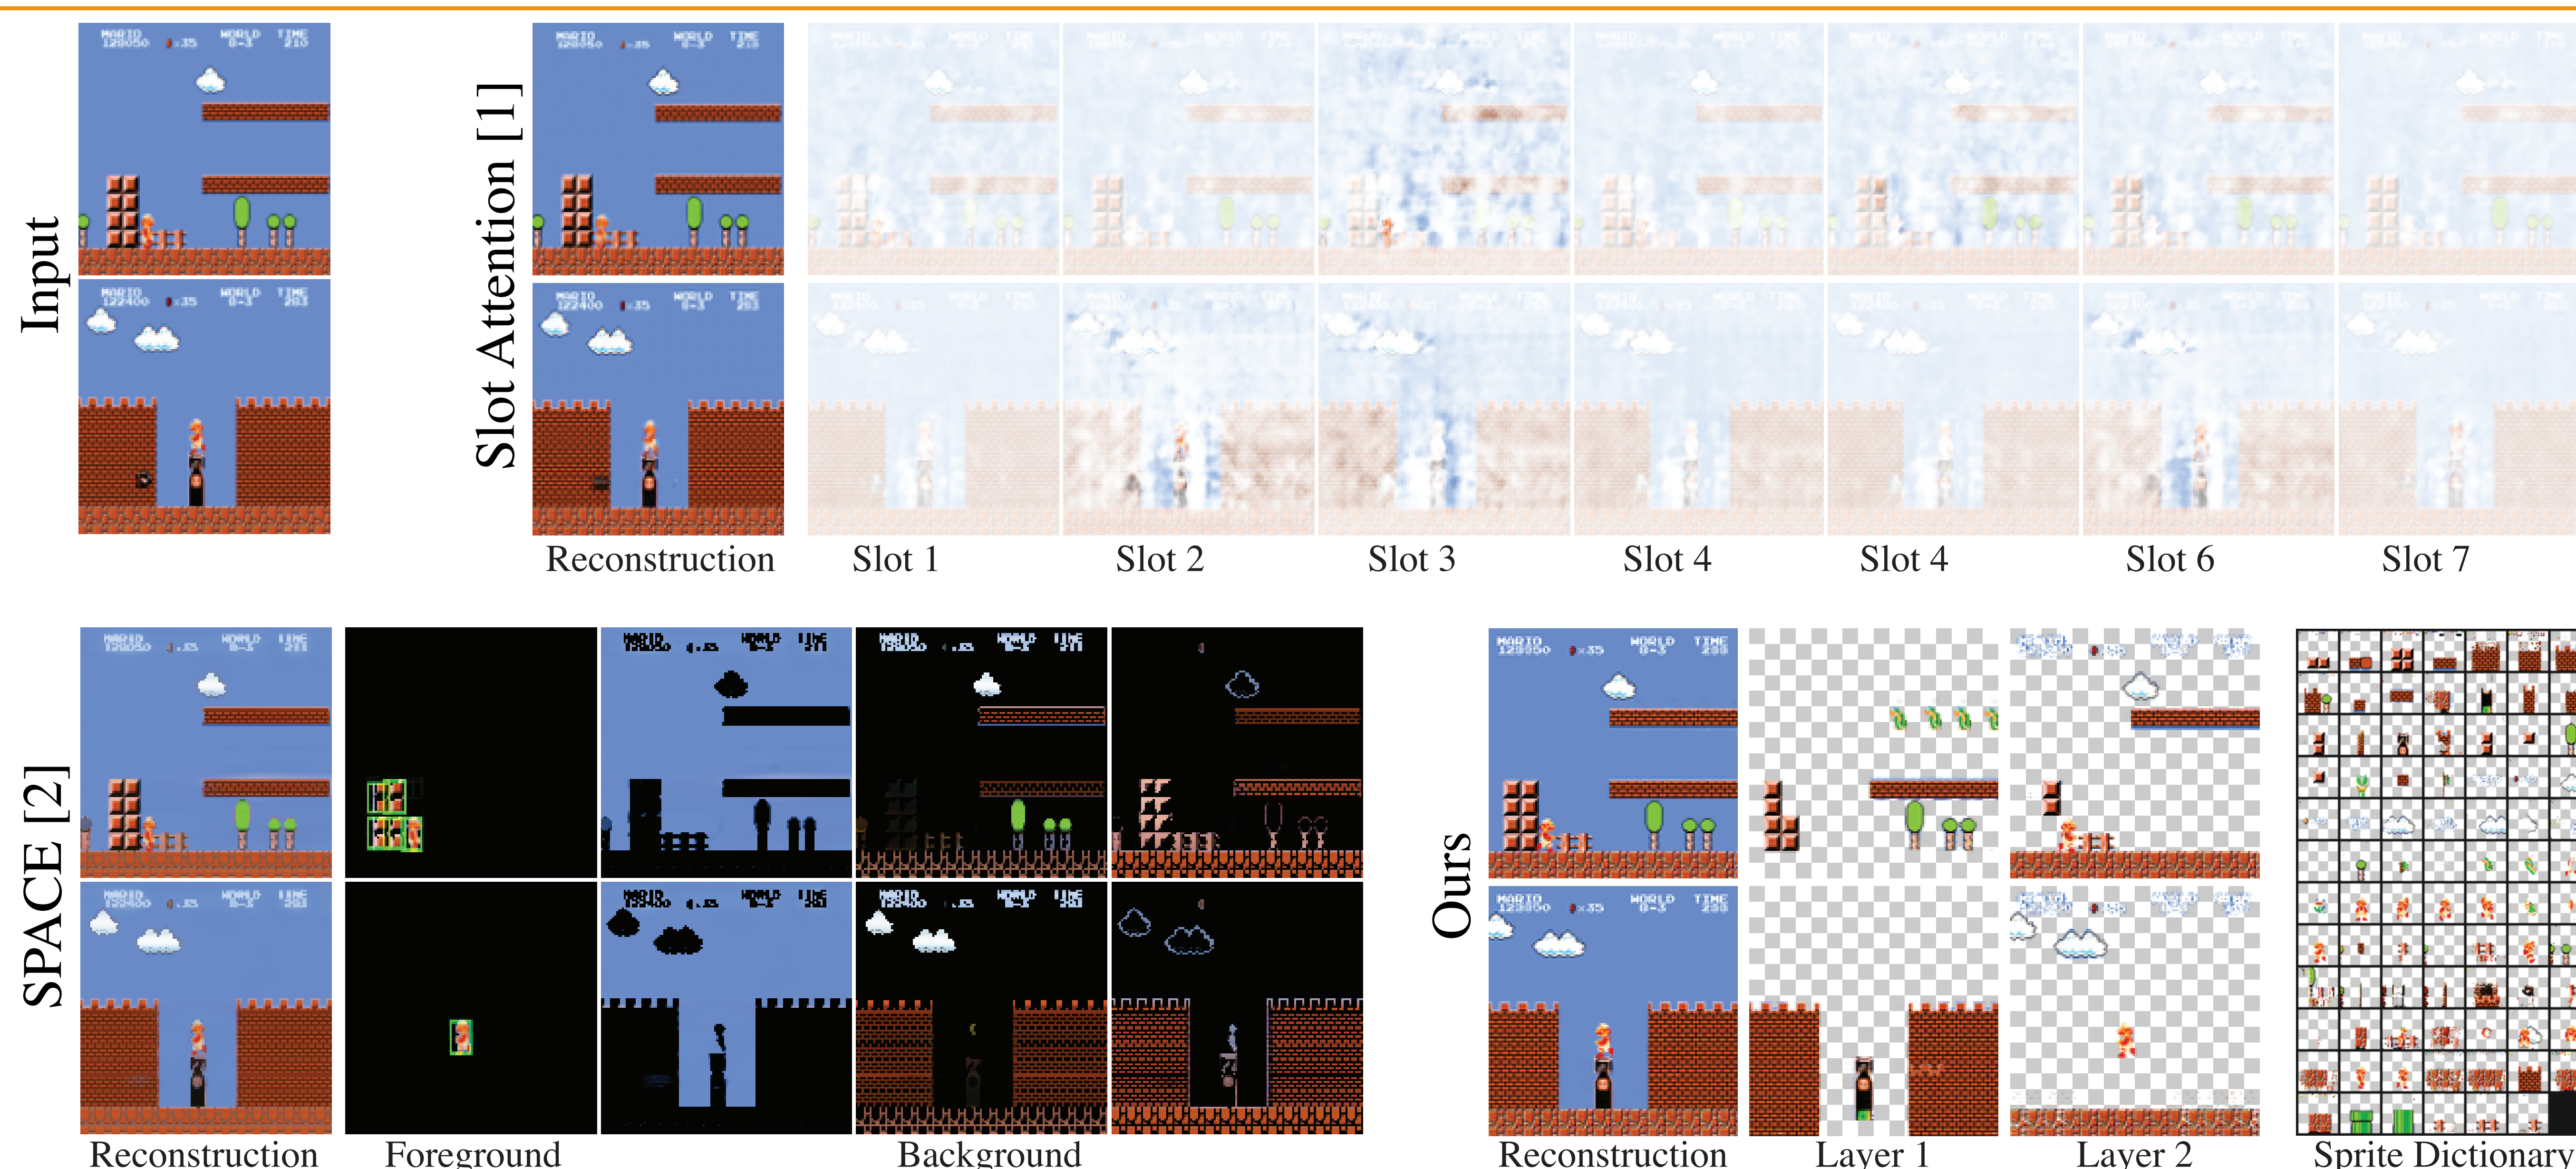
PIPELINE OVERVIEW



RESULTS: SPRITE-BASED GAME DECONSTRUCTION



COMPARISONS



RESULTS: UNSUPERVISED NATURAL VIDEO SEGMENTATION



[1] F. Locatello, D. Dirk Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, and A. Dosovitskiy. "Object-Centric Learning with Slot Attention." arXiv 2020.
[2] Z. Lin, Y. Wu, V. Peri, W. Sun, G. Singh, F. Deng, J. Jiang, and S. Ahn. "SPACE: Unsupervised Object-Oriented Scene Representation via Spatial Attention and Decomposition." ICLR 2020.

PROJECT WEBPAGE

<https://people.csail.mit.edu/smirnov/marionette/>

