Supplementary Material

SketchINR: A First Look into Sketches as Implicit Neural Representations

Hmrishav Bandyopadhyay[1] Ayan Kumar Bhunia[1] Pinaki Nath Chowdhury[1] Aneeshan Sain[1]
Tao Xiang[1,2] Timothy Hospedales [3] Yi-Zhe Song[1,2]
[1]SketchX, CVSSP, University of Surrey, United Kingdom.
[2]iFlyTek-Surrey Joint Research Centre on Artificial Intelligence.
[3]University of Edinburgh, United Kingdom
{h.bandyopadhyay, a.bhunia, p.chowdhury, a.sain, t.xiang, y.song}@surrey.ac.uk
t.hospedales@ed.ac.uk

More on inference time calculations: We present two primary models: one trained with latent vectors and the other as a VAE. While the latter works well with sketches from *Vector-MNIST* and *Quick-Draw!*, we use learned latent vectors (in Tab. 1) for more complex datasets like *FS-COCO* and *Sketchy* (having $\sim$3000 and $\sim$600 number of stroke points, respectively). For a fair comparison in Tab. 1, we evaluate against (i) the "Decoder only" model of Sketch-RNN and (ii) only the decoder from CoSE. Methods like BézierSketch [1] and SketchODE [2] are trained in a per-sketch optimisation setting as ours. We compute the inference time for all methods, as time taken to decode a sketch from its latent representation. We particularly emphasise on SketchINR's scalability to very complex sketches in *FS-COCO*, where even per-sample optimisation versions of BézierSketch and SketchODE fail considerably.

References

[1] Ayan Das, Yongxin Yang, Timothy Hospedales, Tao Xiang, and Yi-Zhe Song. Béziersketch: A generative model for scalable vector sketches. In *ECCV*, 2020. 1

[2] Ayan Das, Yongxin Yang, Timothy Hospedales, Tao Xiang, and Yi-Zhe Song. Sketchode: Learning neural sketch representation in continuous time. In *ICLR*, 2021. 1