The Impact of Negative Sampling on Contrastive Structured World Models

Ondrej Biza (Northeastern U)
Elise van der Pol (U of Amsterdam)
Thomas Kipf (Google Brain)

BACKGROUND:
The Contrastive Structured World Model (C-SWM, Kipf et al.) can model the dynamics of environments that are

- deterministic,
- fully-observable, and
- decompose into several objects.

In order to minimize C-SWM's contrastive loss, the learned encoding of the current state should be predictive of the encoding of the next state. At the same time, the encoding of the current **(positive)** state should be sufficiently different from some other **(negative)** state.

METHODS
We experiment with three negative sampling strategies:

Baseline negatives: permute positive example in a mini-batch to create negative examples.
Time-aligned negatives: a negative example is sampled from a different episode but the same timestep as the positive example.
Episodic and out-of-episode negatives: mix negative example from the same and different episodes as positive counterparts with a ratio β.

RESULTS
#1: Time-aligned negatives can double C-SWM's 10 step prediction score (main panel) and perform similarly to a recent improvement on C-SWM (right panel, 3rd and 4th row).
#2: Mixing of in- and out-of-episode negatives can both help and hurt depending on how we evaluate the model (right panel, 1st and 2nd row).

Carefully chosen negative examples in the Contrastive Structured World Model greatly improve predictions.