

Appendix A Action-Conditioned Temporal Abstraction State Transition

We implement action-conditioned state transition

$$p(z_t|a_t, z_{<t}, m_{<t}) = \begin{cases} \delta(z_t = z_{t-1}) & \text{if } m_{t-1} = 0 \text{ (COPY),} \\ \tilde{p}(z_t|c_t) & \text{otherwise (UPDATE)} \end{cases}$$

where the action input a_t is only affecting the UPDATE operation and we feed it into the deterministic path as the following:

$$c_t = \begin{cases} c_{t-1} & \text{if } m_{t-1} = 0 \text{ (COPY),} \\ f_z(z_{t-1} \| a_t, c_{t-1}) & \text{otherwise (UPDATE).} \end{cases}$$

Appendix B Goal-Oriented Navigation

Algorithm 1 Goal-oriented navigation with imagination (single episode)

Input: environment env , environment model \mathcal{E} , maximum length of imagined trajectories l_{img}
Output: reward r

```

1: Initialize reward  $r \leftarrow 100$ 
2: Initialize environment  $x \leftarrow \text{env.reset}()$ 
3: Initialize context  $X_{\text{ctx}} \leftarrow [x]$ 
4: Sample goal position and extract goal map feature  $g$ 
5: while  $r > 0$  do
6:   Sample action  $a$  from imagination-based planner given  $\mathcal{E}, X_{\text{ctx}}, g, l_{\text{img}}$  (Algorithm 2)
7:   Do action  $x \leftarrow \text{env.step}(a)$ 
8:   Update context  $X_{\text{ctx}} \leftarrow X_{\text{ctx}} + [x]$ 
9:   Update reward  $r \leftarrow r - 1$ 
10:  if current position is at the goal position then
11:    break
12: return reward  $r$ 

```

Algorithm 2 Imagination-based planner

Input: environment model \mathcal{E} , previously observed sequence (context) X_{ctx} , maximum length of imagined trajectories l_{img} , goal map feature g

Output: action a_{\max}

```

1: Initialize model  $\mathcal{E}$  with  $c_0 = f_{\text{ctx}}(X_{\text{ctx}})$ 
2: Initialize  $d_{\max} \leftarrow -\infty$ 
3: Initialize  $a_{\max} \leftarrow \text{None}$ 
4: Set a list  $\mathcal{A}$  of all possible action sequences based on  $l_{\text{img}}$ 
5: for each action sequence  $A$  in  $\mathcal{A}$  do
6:   Get a sequence of states  $\mathcal{S} \in \mathbb{R}^{l_{\text{img}} \times d}$  by doing imagination with model  $\mathcal{E}$  conditioned on  $A$ 
7:   Compute cosine-similarity  $D \in \mathbb{R}^{l_{\text{img}}}$  between all states  $\mathcal{S}$  and goal map feature  $g$ 
8:   if  $\max(D) > d_{\max}$  then
9:     Update  $d_{\max} \leftarrow \max(D)$ 
10:    Update  $a_{\max} \leftarrow A[0]$ 
11: return action  $a_{\max}$ 

```

Appendix C Implementation Details

For bouncing balls, we define the reconstruction loss (data likelihood) by using binary cross-entropy. The images from 3D maze are pre-processed by reducing the bit depth to 5 bits (Kingma & Dhariwal, 2018) and therefore the reconstruction loss is computed by using Gaussian distribution. We used the AMSGrad (Reddi et al., 2018), a variant of Adam, optimizer with learning rate $5e-4$ and all mini-batches are with 64 sequences with length $T = 20$. Both CNN-based encoder and decoder are composed of 4 convolution layers with ELU activations (Clevert et al., 2016). A GRU (Cho et al., 2014) is used for all RNNs with 128 hidden units. The state representations of temporal abstraction and observation abstraction are sampled from 8-dimensional diagonal Gaussian distributions.

Appendix D Evidence Lower Bound (ELBO)

We derive the ELBO without considering recurrent deterministic paths.

Log-likelihood $\log p(X)$

$$\begin{aligned}
\log p(X) &= \log \sum_M \int_{Z,S} p(X, Z, S, M) \\
&\geq \sum_M \int_{Z,S} q(Z, S, M|X) \log \frac{p(X, Z, S, M)}{q(Z, S, M|X)} \\
&= \sum_M \int_{Z,S} q(Z, S, M|X) \log \frac{p(X|Z, S) p(Z, S, M)}{q(Z, S, M|X)} \\
&= \sum_M \int_{Z,S} q(Z, S, M|X) \left[\log p(X|Z, S) + \log \frac{p(Z, S, M)}{q(Z, S, M|X)} \right] \\
&= \underbrace{\mathbb{E}_{q(Z, S, M|X)} [\log p(X|Z, S)]}_{\text{reconstruction}} - \underbrace{\text{KL}[q(Z, S, M|X) || p(Z, S, M)]}_{\text{KL divergence}}
\end{aligned}$$

Decomposing $p(X|Z, S)$ and $p(Z, S, M)$

$$\begin{aligned}
p(X|Z, S) &= \prod_t \underbrace{p(x_t|s_t)}_{\text{decoder}} \\
p(Z, S, M) &= \prod_t \underbrace{p(z_t|z_{t-1}, m_{t-1})}_{\text{temporal abstract transition}} \underbrace{p(s_t|s_{t-1}, z_t, m_{t-1})}_{\text{observation abstract transition}} \underbrace{p(m_t|s_t)}_{\text{boundary prior}}
\end{aligned}$$

Decomposing $q(Z, S, M|X)$

$$\begin{aligned}
q(Z, S, M|X) &= q(M|X) q(Z, S|M, X) \\
&= q(M|X) q(Z|M, X) q(S|Z, M, X) \\
&= q(M|X) \prod_t q(z_t|M, X) q(s_t|z_t, M, X) \\
q(M|X) &= \prod_t q(m_t|X) = \prod_t \text{Bern}(m_t|\sigma(\varphi(X)))
\end{aligned}$$

Reconstruction Term in ELBO

$$\begin{aligned}
&\sum_M \int_{Z,S} q(Z, S, M|X) \log p(X|Z, S) \\
&= \underbrace{\sum_M q(M|X)}_{\text{sample } M} \int_{Z,S} q(Z, S|M, X) \log p(X|Z, S) \\
&\approx \int_{Z,S} q(Z, S|M, X) \sum_t \log p(x_t|s_t) \\
&= \int_{Z,S} \sum_t q(z_t, s_t|M, X) \log p(x_t|s_t) \\
&= \int_{Z,S} \sum_t \underbrace{q(z_t|M, X) q(s_t|z_t, M, X)}_{\text{sampling } z_t \text{ and } s_t} \log p(x_t|z_t, s_t) \\
&\approx \sum_t \log p(x_t|s_t)
\end{aligned}$$

KL Term in ELBO

$$\begin{aligned}
& \log q(Z, S, M|X) - \log p(Z, S, M) \\
&= \log q(M|X) + \log q(Z, S|M, X) - \log p(Z, S, M) \\
&= \log q(M|X) + \log q(Z|M, X) + \log q(S|Z, M, X) - \log p(Z, S, M) \\
&= \log q(M|X) + \sum_t \log \frac{q(z_t|M, X) q(s_t|z_t, M, X)}{p(z_t|z_{t-1}, m_{t-1}) p(s_t|s_{t-1}, z_t, m_{t-1}) p(m_t|s_t)} \\
&= \log q(M|X) + \sum_t \log \frac{q(z_t|M, X)}{p(z_t|z_{t-1}, m_{t-1})} + \log \frac{q(s_t|z_t, M, X)}{p(s_t|s_{t-1}, z_t, m_{t-1})} + \log \frac{1}{p(m_t|s_t)}
\end{aligned}$$

$$\begin{aligned}
& \sum_M \int_{Z,S} q(Z, S, M|X) [\log q(Z, S, M|X) - \log p(Z, S, M)] \\
&= \sum_M q(M|X) \int_{Z,S} q(Z, S|M, X) [\log q(Z, S, M|X) - \log p(Z, S, M)] \\
&= \sum_M q(M|X) \int_{Z,S} q(Z, S|M, X) \left[\log q(M|X) + \sum_t \log \frac{q(z_t|M, X) q(s_t|z_t, M, X)}{p(z_t|z_{t-1}, m_{t-1}) p(s_t|s_{t-1}, z_t, m_{t-1}) p(m_t|s_t)} \right] \\
&= \sum_M q(M|X) \left[\log q(M|X) + \sum_t \int_{Z,S} q(Z, S|M, X) \log \frac{q(z_t|M, X) q(s_t|z_t, M, X)}{p(z_t|z_{t-1}, m_{t-1}) p(s_t|s_{t-1}, z_t, m_{t-1}) p(m_t|s_t)} \right] \\
&= \sum_M q(M|X) \left[\log q(M|X) + \sum_t \int_{z_t, s_t} q(z_t, s_t|M, X) \log \frac{q(z_t|M, X) q(s_t|z_t, M, X)}{p(z_t|z_{t-1}, m_{t-1}) p(s_t|s_{t-1}, z_t, m_{t-1}) p(m_t|s_t)} \right] \\
&= \sum_M q(M|X) \left[\log q(M|X) + \sum_t \text{KL}(q'(z_t)||p'(z_t)) + \int_{z_t, s_t} q(z_t, s_t|M, X) \log \frac{q(s_t|z_t, M, X)}{p(s_t|s_{t-1}, z_t, m_{t-1}) p(m_t|s_t)} \right] \\
&= \sum_M q(M|X) \left[\log q(M|X) + \sum_t \text{KL}(q'(z_t)||p'(z_t)) + \int_{z_t} \underbrace{q'(z_t)}_{\text{sample } z_t} [\text{KL}(q'(s_t)||p'(s_t)) - \log p(m_t|s_t)] \right] \\
&\approx \sum_M q(M|X) \left[\log q(M|X) + \sum_t \text{KL}(q'(z_t)||p'(z_t)) + \text{KL}(q'(s_t)||p'(s_t)) - \log p(m_t|s_t) \right] \\
&= \sum_M q(M|X) \left[\log \frac{\prod_t q(m_t|X)}{\prod_t p(m_t|s_t)} + \sum_t \text{KL}(q'(z_t)||p'(z_t)) + \text{KL}(q'(s_t)||p'(s_t)) \right] \\
&= \sum_{t'} \text{KL}(q'(m_{t'})||p'(m_{t'})) + \sum_M \underbrace{q(M|X)}_{\text{sample } M} \left[\sum_t \text{KL}(q'(z_t)||p'(z_t)) + \text{KL}(q'(s_t)||p'(s_t)) \right] \\
&\approx \sum_t \underbrace{\text{KL}(q'(m_t)||p'(m_t))}_{\text{sequence decomposer}} + \underbrace{\text{KL}(q'(z_t)||p'(z_t))}_{\text{temporal abstraction}} + \underbrace{\text{KL}(q'(s_t)||p'(s_t))}_{\text{observation abstraction}}
\end{aligned}$$

where $p'(m_t) = p(m_t|s_t)$, $p'(z_t) = p(z_t|z_{t-1}, m_{t-1})$, $p'(s_t) = p(s_t|s_{t-1}, z_t, m_{t-1})$
 $q'(m_t) = q(m_t|X)$, $q'(z_t) = q(z_t, s_t|M, X)$, $q'(s_t) = q(s_t|z_t, M, X)$

Appendix E Generated Samples

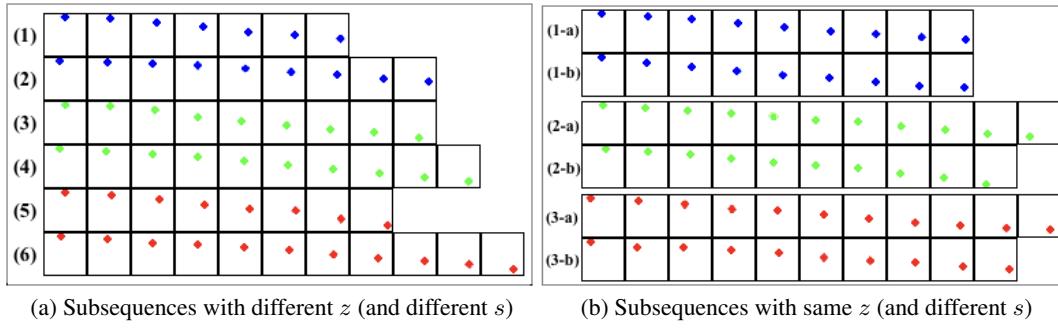


Figure 1: Generated subsequences with different content and temporal structure

References

- Cho, K., van Merriënboer, B., Gülcəhre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D14-1179>.
- Clevert, D., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1511.07289>.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 10215–10224. Curran Associates, Inc., 2018. URL <http://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions.pdf>.
- Reddi, S. J., Kale, S., and Kumar, S. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=ryQu7f-RZ>.