Expansive Latent Space Trees for Planning from Visual Inputs

Robert Gieselmann (robgie@kth.se), Florian T. Pokorny (fpokorny@kth.se)

Division of Robotics, Perception and Learning, KTH Royal Institute of Technology, Stockholm

ROYAL INSTITUTE OF TECHNOLOGY

Planning in Latent Spaces

Long-horizon decision-making from **visual observations** is challenging due to the high dimensionality and complexity of the space of images. Planning in learned latent spaces provides an intriguing alternative due to the reduced dimensionality of the state space.

Expansive Latent Space Trees - Overview

We present **Expansive Latent Space Trees (ELAST)**, a latent planning method which explores solutions by growing a search tree within the estimated support region of the latent space. Our method does not require costly training of image generative models and can be trained in a **self-supervised** fashion given offline datasets of random trajectories.

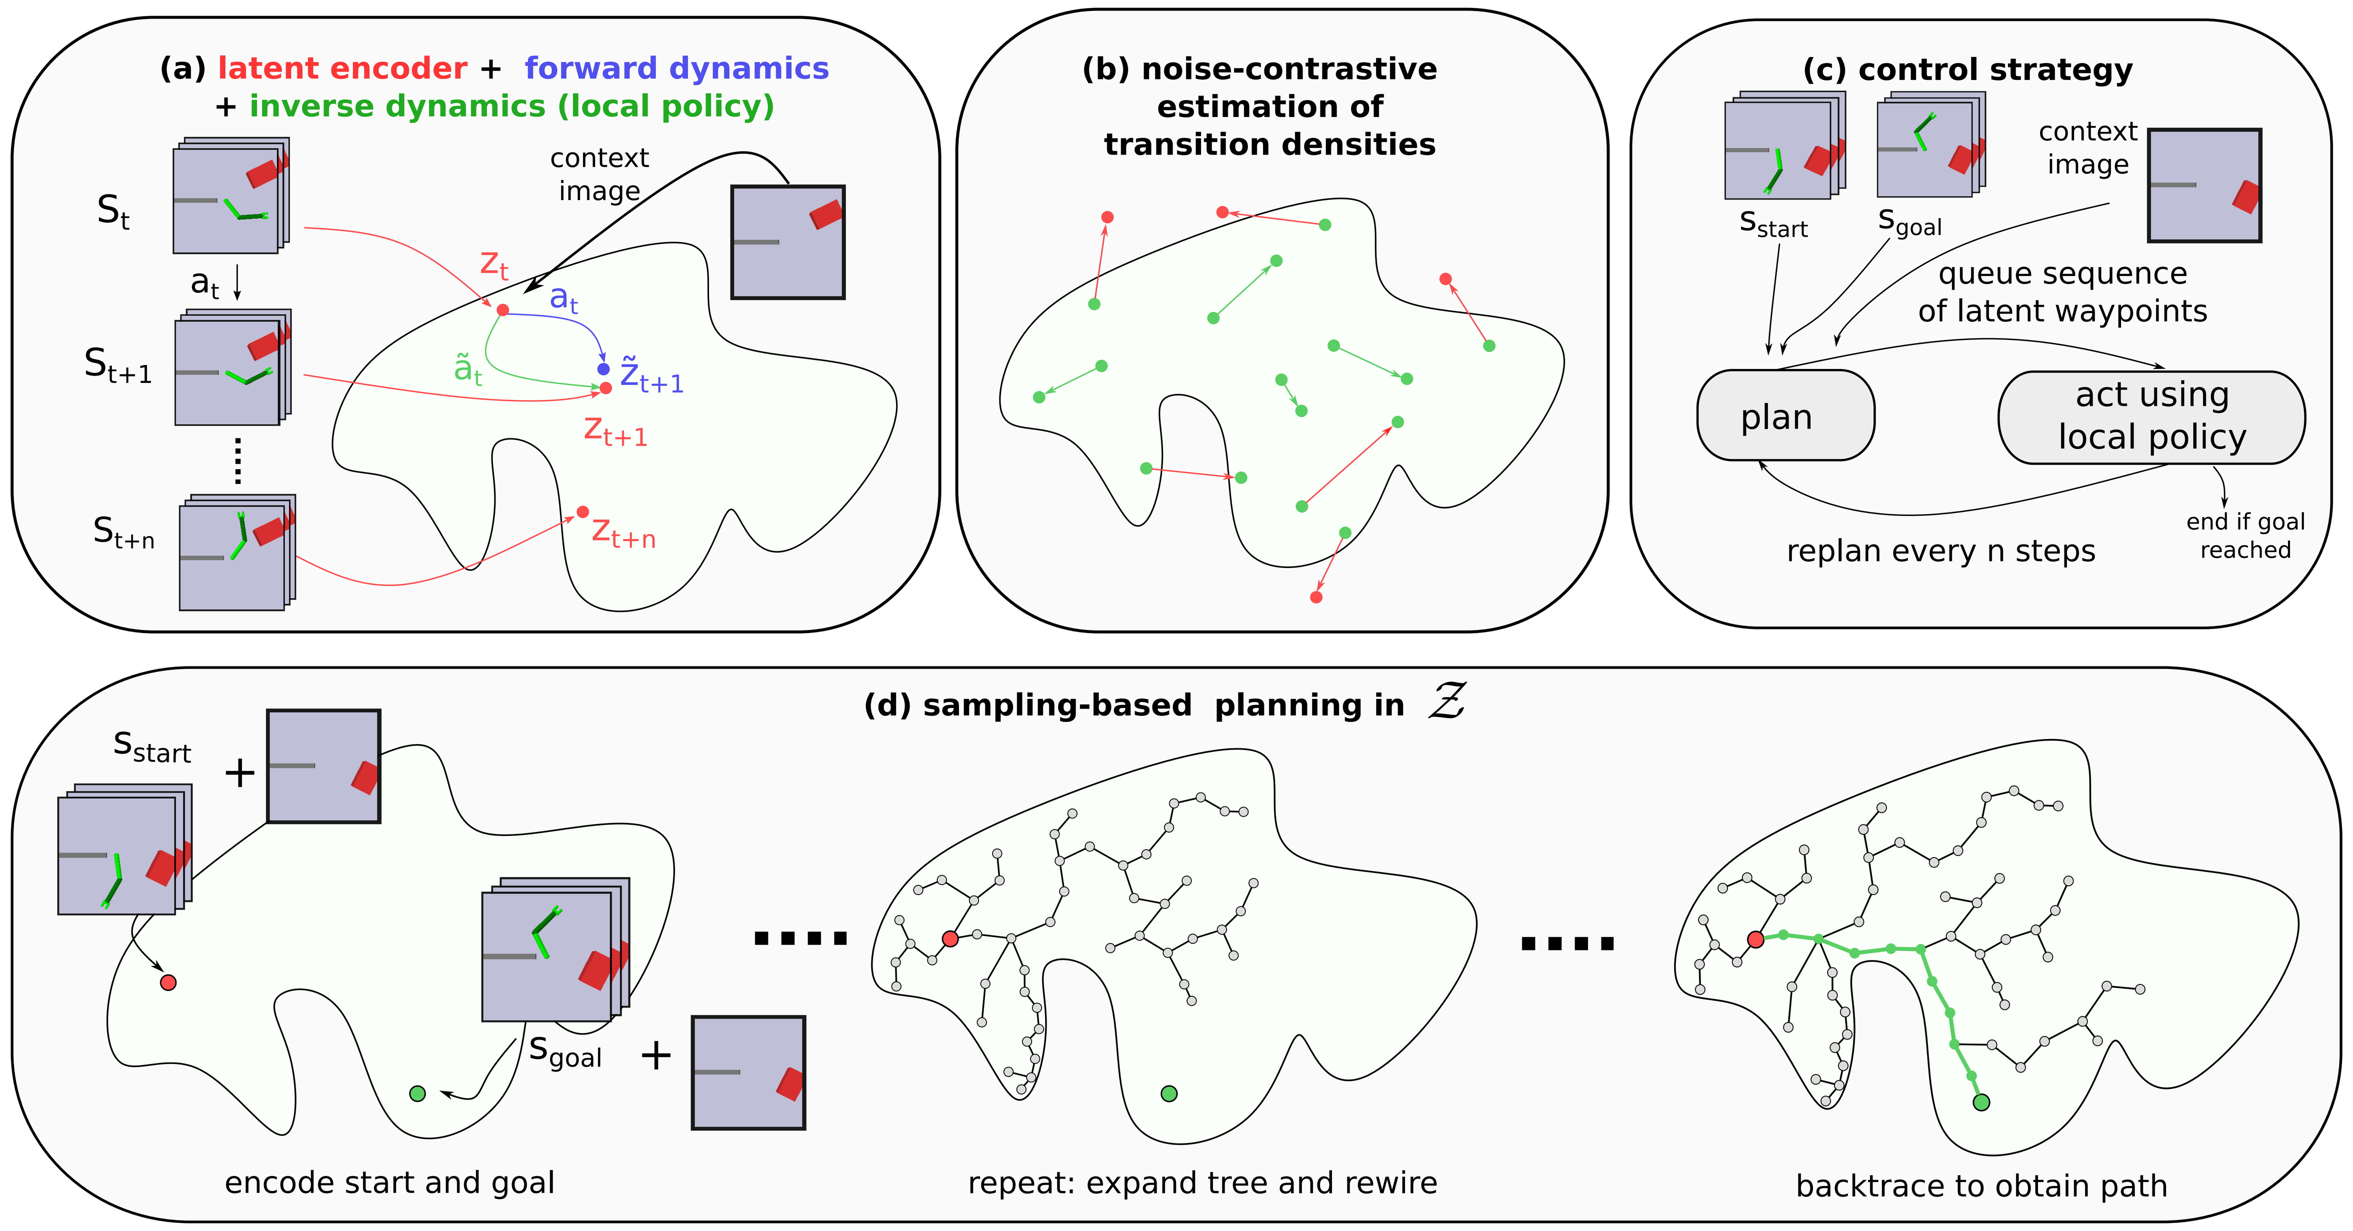

Contrastive State Representations for Control

We employ Representation Learning via **Contrastive Predictive Coding (CPC)** [1] to embed image observations into a lower-dimensional vector space $\mathcal{Z}$ in which temporal vicinity of states is enforced.

Acknowledgement

This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Unsupervised Transition Density Estimation

Inspired by the EST algorithm [2], ELAST utilizes a learned dynamics model to randomly grow a tree within $\mathcal{Z}$. During exploration, the nodes are rewired in order to reduce the traveling distance in the tree. The rewiring mechanism requires to determine if a transition from a latent state z_a to z_b is possible.

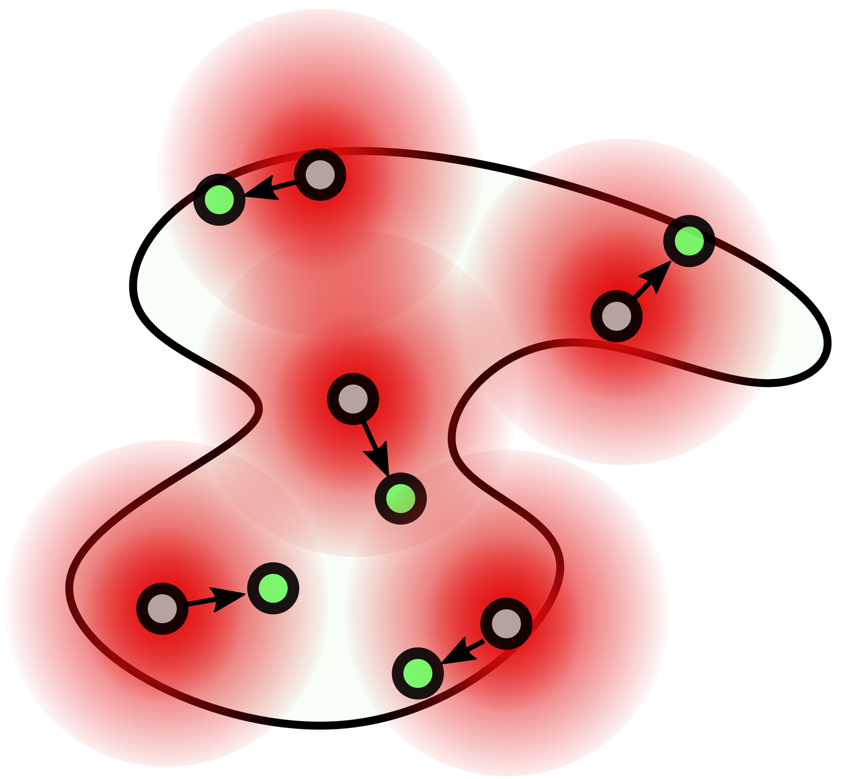

We formulate this task as density estimation of latent transitions $p(z_b|z_a)$ and train a parametric density estimator unsupervised using **Noise-Contrastive Estimation (NCE)** [3]. Our key insight is that due to the locality of temporally correlated states in $\mathcal{Z}$, we can define noise distributions p_n by centering a multivariate Gaussian on each conditioning state z_a.

Initial Experimental Results

We evaluated ELAST on a set of simulated planning domains. The figure below illustrates several latent paths that were planned with ELAST and projected into the 2D Isomap embedding.

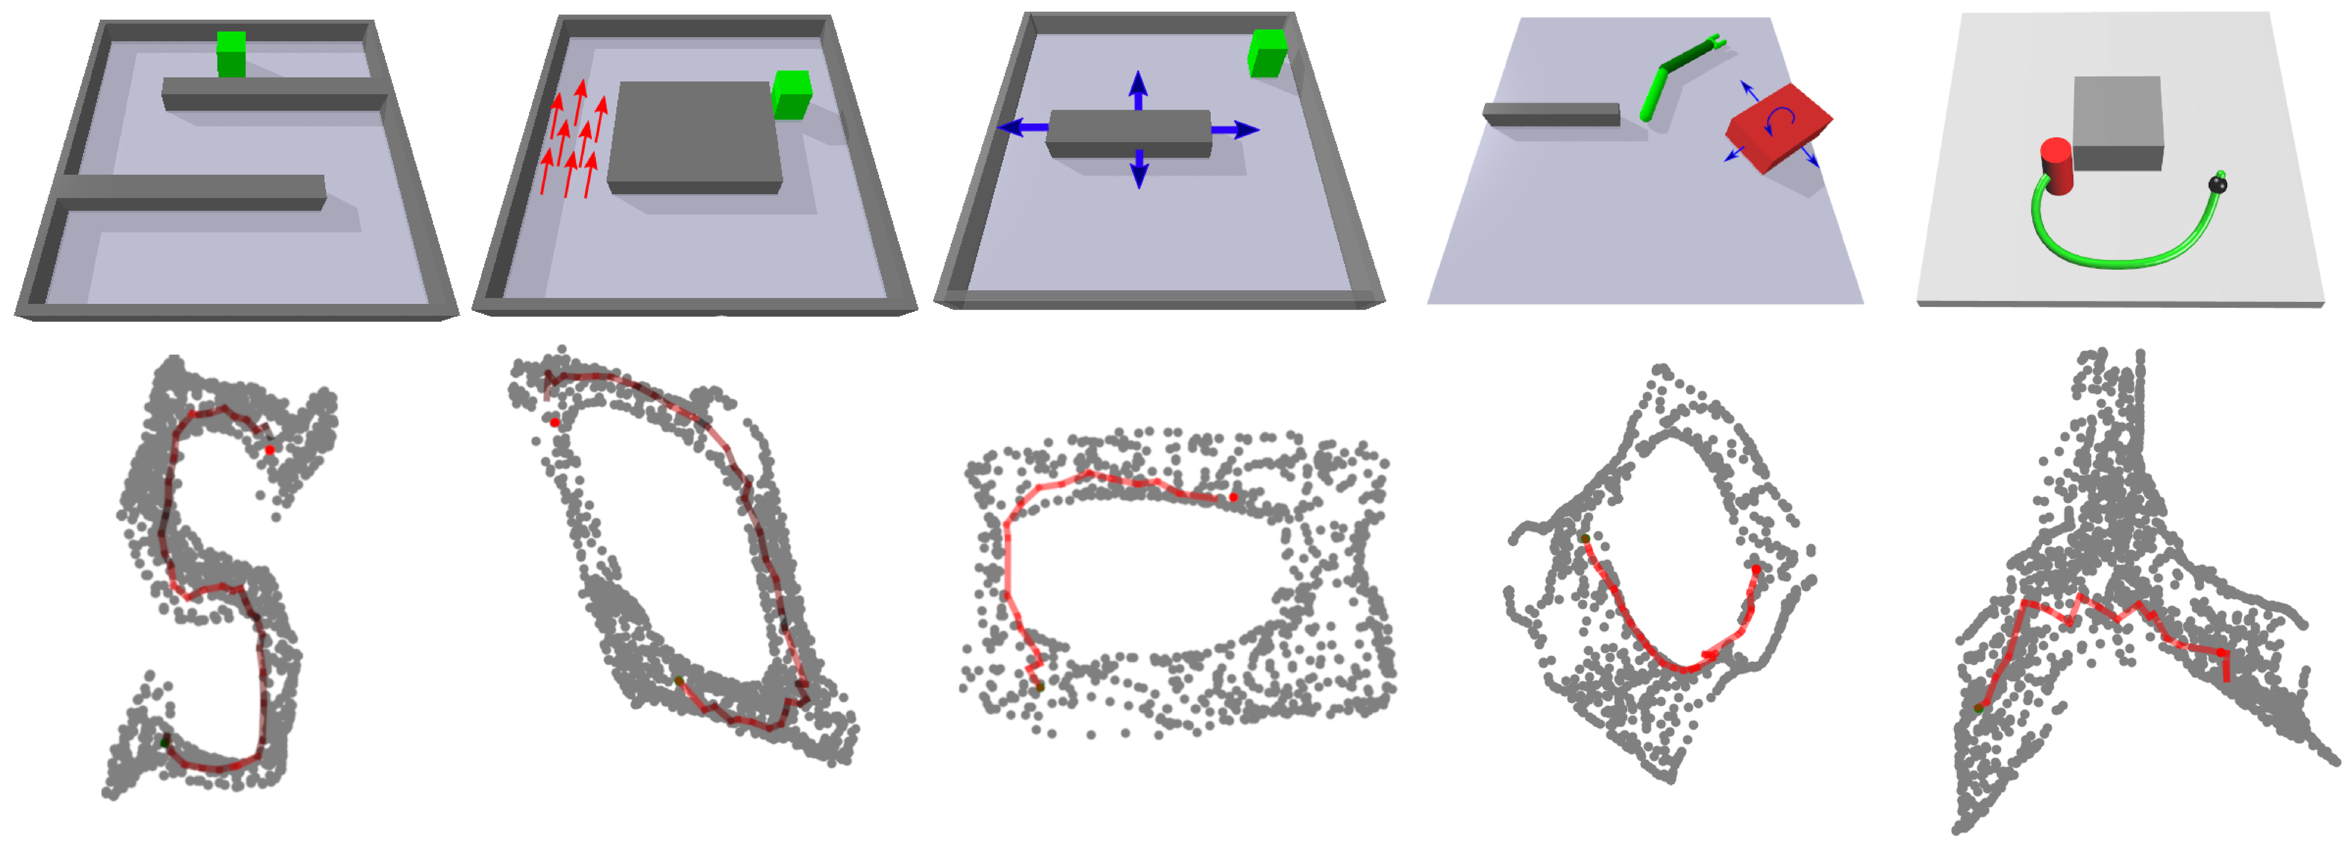

For a quantitative evaluation and comparison with existing baselines, we refer to our workshop paper.

References

[1] Aaeron van den Oord, Yazhe Li, and Oriol Vinyals "Representation learning with contrastive predictive coding." CoRR, abs/1807.03748 2018

[2] D. Hsu, J.-C. Latombe, and R. Motwani "Path planning in expansive configuration spaces." Proceedings of International Conference on Robotics and Automation, volume 3, pages 2719-2726 vol.3 1997

[3] Michael U. Gutmann and Aapo Hyvaerinen "Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics" Journal of Machine Learning Research, 13(11):307-361 2012