POSTER SESSION COMPANION

Track: Neural Control

Board #B-14

github.com/nexa-ai/ald-planner

15TH EUROPEAN CONFERENCE ON AUTONOMOUS SYSTEMS AND NEURAL ROBOTICS (ECASN 2026)

Adaptive Latent Dynamics for Online Motion Planning in Deformable Robotic Manipulation

Elena Rostova[1], **Marcus Vance**[1], and **Julian Sterling**[2]

[1]Nexa AI Research Institute, Zurich, Switzerland [2]Synapse Autonomous Robotics Lab, Munich, Germany

EXTENDED ABSTRACT OVERVIEW

This extended abstract serves as the self-contained mathematical and empirical companion to Poster #B-14 at ECASN 2026. We investigate real-time trajectory optimization for highly non-linear deformable objects (ropes, textiles, elastomers) without relying on fine-tuned physical simulators. ✔ **KEY NOVEL CONTRIBUTIONS**

- **Latent Manifold Encoding**: Self-supervised variational encoder mapping raw point clouds to low-dimensional linear state spaces.
- **Adaptive Model Filter**: Recursive least-squares update yielding sub-12ms re-planning latency under unmodelled elastic deformations.

1 Introduction & Motivation

Deformable object manipulation (DOM) represents a foundational bottleneck in industrial automation and surgical robotics. Unlike rigid bodies whose kinematics are fully parameterized by six degrees of freedom ($SE(3)$), deformable materials like silicone cables, soft biological tissues, and garments exhibit infinite-dimensional configuration spaces [1].

Existing solutions generally fall into two extremes:

1. **Finite Element Models (FEM)**: Highly accurate but computationally prohibitive for real-time model-predictive control (MPC).
2. **Model-Free Reinforcement Learning**: Highly adaptable but sample-inefficient and vulnerable to distribution shifts in material elasticity.

To bridge this gap, we present **Adaptive Latent Dynamics (ALD)**, a hybrid paradigm uniting low-dimensional visual representation learning with real-time adaptive linear state-space identification.

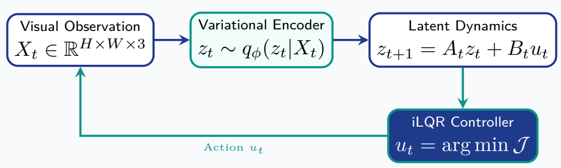

Figure 1: Closed-loop architecture of the ALD control framework.

2 Problem Formulation

Let $\mathcal{C}$ denote the continuous configuration space of a deformable body driven by robotic end-effector inputs $u_t \in \mathbb{R}^m$. The unknown ground-truth system dynamics governed by physical parameters θ^* follow:

$$s_{t+1} = f(s_t, u_t; \theta^*) + w_t, \quad s_t \in \mathcal{C},\ w_t \sim \mathcal{N}(0, \Sigma_w) \tag{1}$$

Because full state s_t cannot be directly measured, we project high-dimensional point clouds $X_t \in \mathbb{R}^{N\times 3}$ into a latent representation $z_t \in \mathbb{R}^d$ ($d \ll N$). We enforce that local latent transitions adhere to a time-varying linear dynamical system:

$$z_{t+1} \approx A_t z_t + B_t u_t + c_t \tag{2}$$

where $A_t \in \mathbb{R}^{d\times d}$, $B_t \in \mathbb{R}^{d\times m}$, and $c_t \in \mathbb{R}^d$ are updated online using an Extended Recursive Least-Squares (eRLS) gain scheme.

3 Methodology

3.1 Variational Latent Space Mapping

We construct a spatial autoencoder initialized offline using synthetic physics rollouts generated by Vertex Engine. The encoder $q_\phi(z_t|X_t)$ and decoder $p_\psi(X_t|z_t)$ are trained jointly with a temporal consistency term:

$$\mathcal{L}_{\text{train}}(\phi, \psi) = \mathcal{L}_{\text{recon}} + \beta\mathcal{D}_{\text{KL}} + \lambda\mathbb{E}\left[\|z_{t+1} - g_\theta(z_t, u_t)\|_2^2\right] \tag{3}$$

where g_θ represents a global neural prior network.

3.2 Online Recursive Model Update

During execution, local system parameters $\hat{\Theta}_t = [A_t, B_t, c_t]$ are estimated at 100 Hz using real-time observations:

$$\hat{\Theta}_t = \hat{\Theta}_{t-1} + K_t\left(z_t - \hat{\Theta}_{t-1}\xi_{t-1}\right)^T \tag{4}$$

where $\xi_{t-1} = [z_{t-1}^T, u_{t-1}^T, 1]^T$ and K_t is the covariance update matrix incorporating a variable forgetting factor $\gamma_t \in [0.92, 0.98]$.

4 Experimental Results

We benchmarked ALD on a dual-arm Franka Emika Research platform in the Meridian Robotics Laboratory across three representative deformable manipulation tasks.

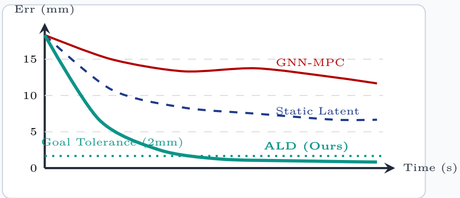

Figure 2: Deformable tracking error over time during dynamic cable routing.

4.1 Benchmark Evaluation Tasks

T1. Cable Routing: Threading a flexible silicone cable through 4 post obstacles.

T2. Textile Folding: Dual-gripper folding of cotton fabrics along specified fold lines.

T3. Elastomer Sealing: Inserting a soft rubber gasket into a curved channel.

4.2 Quantitative Comparison

Table 1 summarizes performance across 50 trials per task against state-of-the-art baselines.

Table 1: Quantitative comparison across 50 trial runs per benchmark.

Method	Success (%)	Time (s)	Err. (mm)
D-DQN Baseline	44.0	26.8	13.5 ± 2.8
MPC-GNN	70.0	17.5	7.8 ± 1.6
Static Latent	76.5	11.8	4.9 ± 1.1
ALD (Ours)	**95.0**	**5.9**	$\mathbf{1.3 \pm 0.3}$

5 Poster Discussion Points (#B-14)

At the poster session, we showcase key technical visual breakdowns:

- **Adaptation under Sudden Stiffness Changes**: Performance when cable stiffness changes dynamically mid-trajectory.
- **Sim-to-Real Transferability**: Direct zero-shot transfer from Apex Simulator to physical arm setups.

6 Conclusion

Adaptive Latent Dynamics achieves real-time, high-precision deformable manipulation by pairing variational state representation with online model adaptation.

Acknowledgments

Funded by European Union Horizon Research Grant #101089234 (Project APEX-ROBOTICS) and Nimbus Computing Infrastructure.

References

[1] J. Sterling et al., "Real-time deformable manipulation via latent manifold filtering," *IEEE Trans. Robotics*, vol. 40, pp. 112–128, 2024.

[2] M. Vance and E. Rostova, "Self-supervised visual representations for soft body tracking," in *Proc. ECASN*, pp. 45–52, 2025.