Structured deep generative models for sampling on constraint manifolds in sequential manipulation

Learning and Intelligent Systems

Joaquim Ortiz-Haro   Jung-Su Ha   Danny Driess   Marc Toussaint   *TU Berlin*

imprs-is

Deep Generative Constraint Sampling (DGCS)

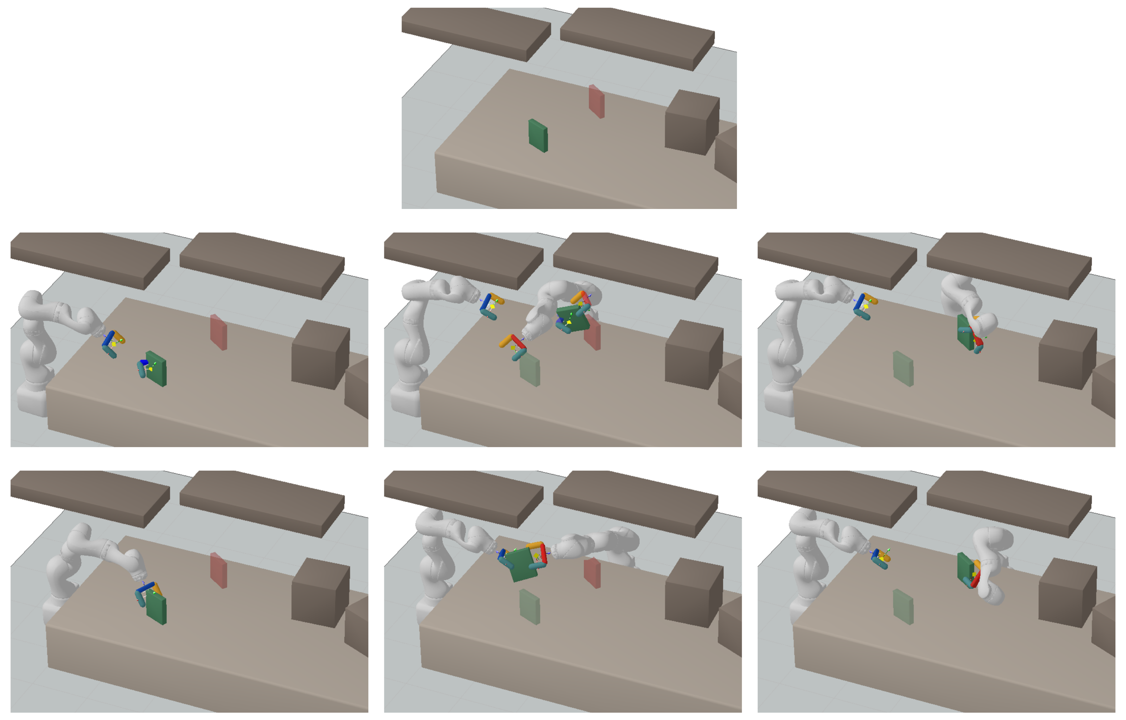

0 - New instance
Image representation

1 - Approximate sample with deep generative model

2 - Solution with nonlinear optimizer

Sampling Framework

Goal: Sample on a parametric constraint manifold,

$$\mathcal{M}_\tau = \{x \in \mathbb{R}^n \quad \text{s.t} \quad \phi_{\text{eq}}(x;\tau) = 0,\ \phi_{\text{ineq}}(x;\tau) \leq 0\}\ ,$$

$\phi_{\text{eq}}(x;\tau)$, $\phi_{\text{ineq}}(x;\tau)$ nonlinear piecewise differentiable vector functions.
$\tau \in \mathbb{R}^m$ parametrization of current problem instance.

Sample on a constraint manifold

1. Sample $x_0 \sim \mathbb{P}_\theta(\tau)$
2. $x \leftarrow \Pi(x_0)$, Project to $\mathcal{M}_\tau$ with a nonlinear optimizer

$$\min_x ||x - x_0||^2 \text{ s.t } \phi_{\text{eq}}(x;\tau) = 0;\ \phi_{\text{ineq}}(x;\tau) \leq 0$$

Training Deep Generative Models

Deep generative model $\tilde{x} \sim \mathbb{P}_\theta(\tau)$ with $\tilde{x} = G_\theta(z, \tau)$, $z \sim \mathbb{P}_z$

$\mathbb{P}_z$ is a multidimensional Gaussian and G_θ is a neural network.

Analytical features of the support of the distribution,

$$\phi(x;\tau) = [\phi_{\text{eq}}(x;\tau), \max(0, \phi_{\text{ineq}}(x;\tau))]$$

Sample diversity: regularize with respect to a reference distribution $\mathbb{P}_r$ (Dataset of solutions $\{x_i, \tau_i\}$) with Wasserstein distance W,

$$\min \mathop{\mathbb{E}}_{\tau}\ W(\mathbb{P}_\theta(\tau), \mathbb{P}_r(\tau)) + \beta \mathop{\mathbb{E}}_{\tilde{x}\sim\mathbb{P}_\theta} ||\phi(\tilde{x};\tau)||^2\ ,$$

Wasserstein GAN [1] formulation. Minimax game (stochastic gradient descent) between the critic network D and the generator G,

$$\min_G \max_D \mathop{\mathbb{E}}_{\tau} \mathop{\mathbb{E}}_{x\sim\mathbb{P}_r} D(x;\tau) - \mathop{\mathbb{E}}_{\tilde{x}\sim\mathbb{P}_\theta} D(\tilde{x};\tau) - \lambda \mathop{\mathbb{E}}_{\hat{x}\sim\mathbb{P}_{\hat{x}}} (\|\nabla D(\hat{x};\tau)\| - 1)^2 + \beta \mathop{\mathbb{E}}_{\tilde{x}\sim\mathbb{P}_\theta} \|\phi(\tilde{x};\tau)\|^2$$

Structure: Constraint Graphs

Exploit factorization of the problem [3]. Constraint graph representation.

$$x = \{x_1, \ldots, x_N\}\ , \quad \phi(x;\tau) = \{\phi_1, \ldots, \phi_L\}\ .$$

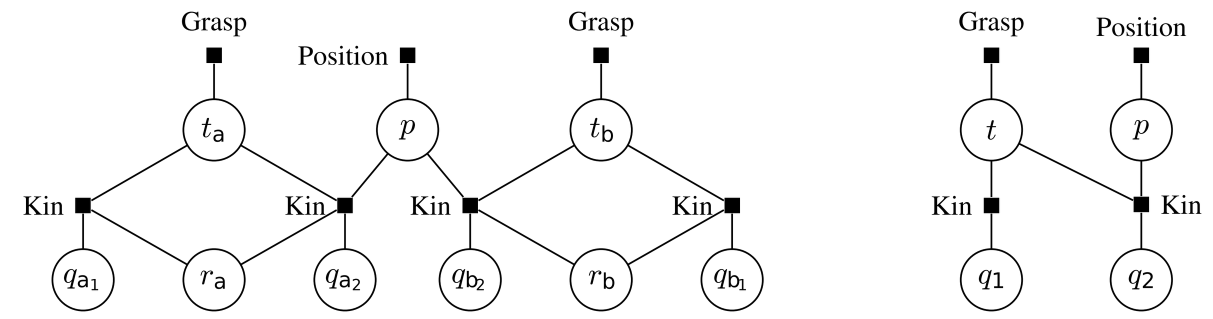

q robot joint configuration
r mobile base pose
t transformation object-gripper
p position of the object

Constraints: Kinematic, Grasp, Position and Collision avoidance.

Directed Graphical Model and Sequential Sampling

- Factorization of the joint probability based on the graph structure.
- Train conditional models with the marginals of the original data.

Example: Pick and Place

$$P(p, t, q_1, q_2) = P(p)\, P(t|p)\, P(q_1|t)\, P(q_2|t, p)\ ,$$

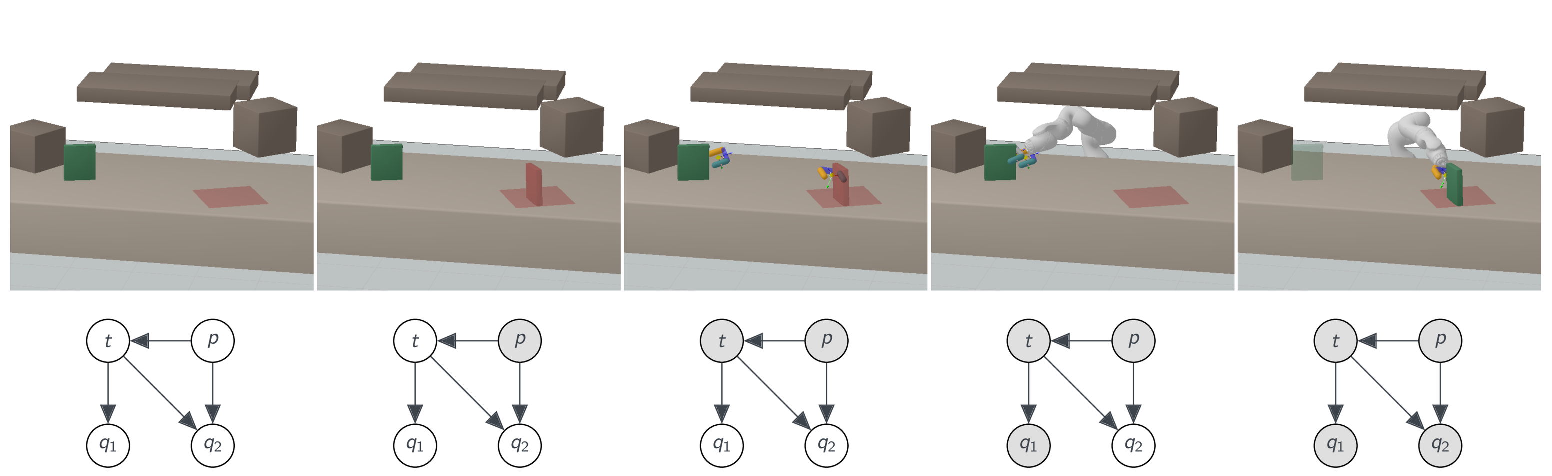

Advantages:
1. Reduce sample complexity and partition of conditioning
2. Improve Multimodality and training stability

Image-Based Problem Parametrization

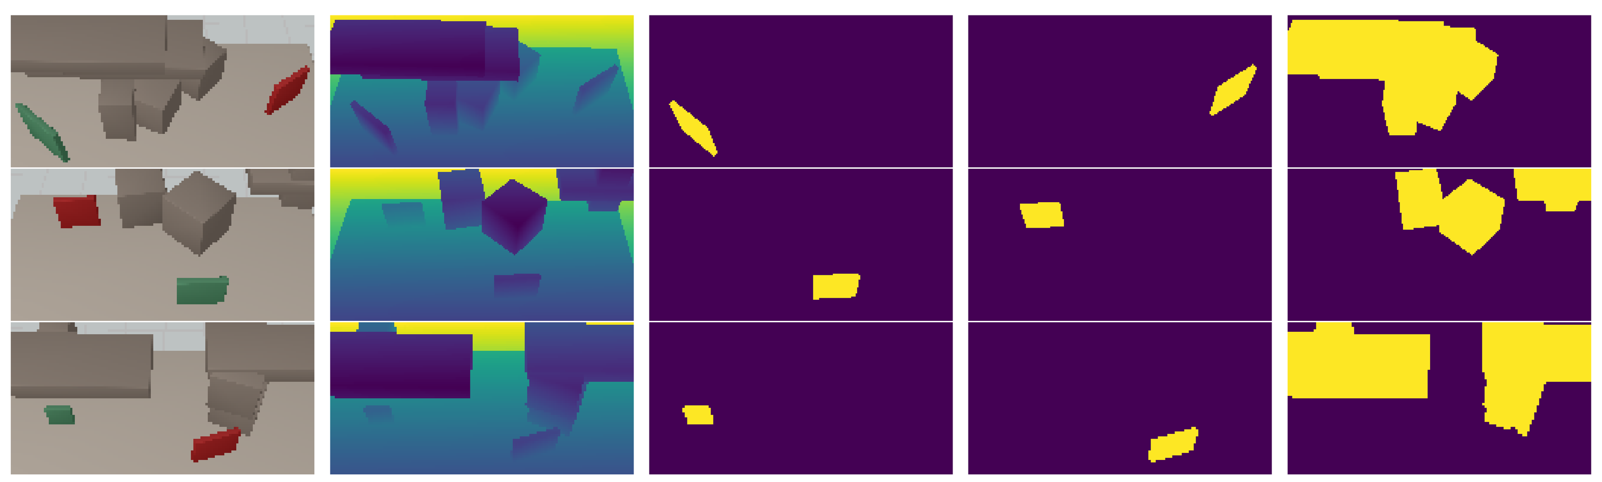

Experiments

1. Pick and Place
2. Handover
3. Assembly

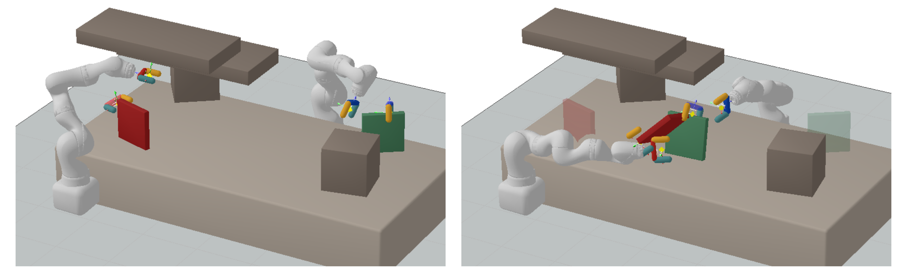

Ablation Study (Pick and Place)

	Seeds			Solutions		
	Coverage	Precision	Error	Coverage	Precision	Success
Big NN	0.81	0.7	8.38	0.58	0.39	0.46
Big NN + analytical	0.79	0.53	1.21	0.75	0.41	0.43
Structure NN	0.6	0.62	8.09	0.41	0.44	0.56
Structure NN + analytical	0.57	0.47	1.46	0.44	0.28	0.78

Benchmark: Warmstart Nonlinear Optimization

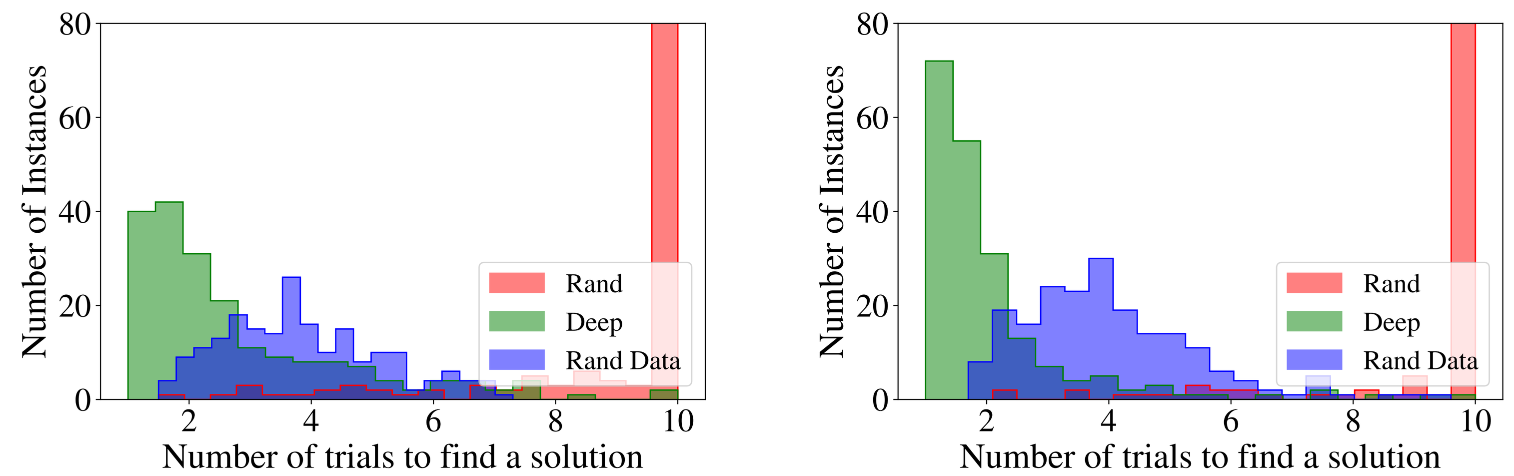

(a) Handover (b) Assembly

Estimated number of trials necessary to solve an instance.

Deep (deep generative model with structure + analytical); Baselines: *Rand* (random value) and *Rand Data* (random point of dataset of solutions).

Conclusion

- DGCS combines a deep generative model with nonlinear optimization
- Contributions: Structure and analytical features
- Outperform heuristic warmstart

Future Work 1 - Generalization to different problem classes 2 - Optimization as the last layer of the neural network.

References

[1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

[2] Danny Driess, Jung-Su Ha, and Marc Toussaint. Deep visual reasoning: Learning to predict action sequences for task and motion planning from an initial scene image. In *Robotics: Science and Systems 2020*, 2020.

[3] Caelan Reed Garrett, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Sampling-based methods for factored task and motion planning. *CoRR*, abs/1801.00680, 2018.