

Model-Based Reinforcement Learning via Latent-Space Collocation

Oleh Rybkin^{*1} Chuning Zhu^{*1} Anusha Nagabandi² Kostas Daniilidis¹ Igor Mordatch³ Sergey Levine⁴

Abstract

The ability to plan into the future while utilizing only raw high-dimensional observations, such as images, can provide autonomous agents with broad capabilities. Visual model-based reinforcement learning (RL) methods that plan future actions directly have shown impressive results on tasks that require only short-horizon reasoning, however, these methods struggle on temporally extended tasks. We argue that it is easier to solve long-horizon tasks by planning sequences of states rather than just actions, as the effects of actions greatly compound over time and are harder to optimize. To achieve this, we draw on the idea of collocation, which has shown good results on long-horizon tasks in optimal control literature, and adapt it to the image-based setting by utilizing learned latent state space models. The resulting latent collocation method (LatCo) optimizes trajectories of latent states, which improves over previously proposed shooting methods for visual model-based RL on tasks with sparse rewards and long-term goals. See the videos on the supplementary website <https://orybkin.github.io/latco/>.

1. Introduction

For autonomous agents to perform complex tasks in open-world settings, they must be able to process high-dimensional sensory inputs, such as images, and reason over long horizons about the potential effects of their actions. Recent work in model-based reinforcement learning (RL) has shown impressive results in autonomous skill acquisition directly from image inputs, demonstrating improvements both in terms of data efficiency and generalization (Ebert et al., 2018; Hafner et al., 2019; Zhang et al., 2019; Schrittwieser et al., 2020). While these advancements have been largely

^{*}Equal contribution ¹University of Pennsylvania ²Covariant ³Google AI ⁴UC Berkeley. Correspondence to: Oleh Rybkin <oleh@seas.upenn.edu>.




Figure 1: **Top:** Latent collocation (LatCo) on a tool use task, where the thermos needs to be pushed with the stick. Each image shows a full plan at that optimization step, visualized via a diagnostic network. LatCo optimizes a latent state sequence and is able to temporarily violate the dynamics during planning, such as the stick flying in the air without an apparent cause. This allows it to rapidly discover the high-reward regions, while the subsequent refinement of the planned trajectory focuses on feasibly achieving it. **Bottom:** in contrast, shooting optimizes an action sequence directly and is unable to discover picking the stick as the actions that lead to that are unlikely to be sampled.

fueled by improvements in modeling (Finn et al., 2016; Hafner et al., 2019), they leave much room for improvement in terms of planning and optimization. Many of the current best-performing deep model-based RL approaches use only gradient-free action sampling as the underlying optimizer (Chua et al., 2018; Nagabandi et al., 2020), and are typically applied to tasks with very short horizons, such as planning a sequence of 5 (Ebert et al., 2018) or 12 (Hafner et al., 2019) actions. On longer-horizon tasks, we observe that these shooting methods struggle with local optima, as credit assignment for individual actions becomes harder. In this work, we study how more powerful planners can be used with these models to achieve better longer-horizon reasoning.

It is appealing instead of optimizing a sequence of actions to optimize a sequence of states. While small deviations in

actions can greatly compound over time and affect the entire trajectory, states can often be easily optimized locally just based on the neighboring states in the trajectory (see Figure 1). An approach that optimizes states directly could then perform credit assignment more easily and be better conditioned. However, it is necessary to ensure that the optimized state sequence is dynamically feasible – that is, each state in the trajectory is reachable from the previous state. Prior deep reinforcement learning methods have addressed this problem by learning reachability functions and performing graph search on the replay buffer (Savinov et al., 2018; Kurutach et al., 2018; Eysenbach et al., 2019). However, it is unclear how to use these methods with partial observability, stochastic dynamics, or reward functions beyond goal reaching. Instead, to arrive at an optimal control solution, we turn to the technique of collocation, which optimizes a sequence of states to maximize the reward, while also eventually ensuring dynamics feasibility by recovering the corresponding actions in a constrained optimization problem:

$$\max_{s_{2:T}, a_{1:T-1}} \sum_t r(s_t) \quad \text{s.t.} \quad s_{t+1} = f(s_t, a_t). \quad (1)$$

Collocation only requires learning a dynamics model and a reward function, and can be used as a plug-and-play optimizer within common model-based reinforcement learning approaches, while providing a theoretically appealing formulation for optimizing sequences of states.

Collocation, as introduced above, can provide many benefits over other optimization techniques, but it has thus far been demonstrated (Liu & Popović, 2002; Ratliff et al., 2009; Schulman et al., 2014; Kalakrishnan et al., 2011; Posa et al., 2014) mostly in conjunction with known dynamics models or when performing optimization over state vectors with a few tens of dimensions (Bansal et al., 2016; Du et al., 2020). In this work, we are interested in autonomous behavior acquisition directly from image inputs, where both the underlying states as well as the underlying dynamics are unknown, and the dimensionality of the observations is in the thousands. Naïvely applying collocation to sequences of images would lead to intractable optimization problems, due to the high-dimensional as well as partially observed nature of images. Instead, we draw on the representation learning literature and leverage latent dynamics models, which learn a latent representation of the observations that is not only Markovian but also compact, and lends itself well to planning. In this learned latent space, we propose to perform collocation over states and actions with the joint objective of maximizing rewards as well as ensuring dynamic feasibility.

The main contribution of this work is an algorithm for latent-space collocation (LatCo), which provides a practical way to utilize collocation methods within a model-based RL algorithm with image observations. LatCo plans sequences of latent states directly from image observations, and is

able to plan under uncertainty. To the best of our knowledge, our paper is the first to scale collocation to visual observations, enabling longer-horizon reasoning by drawing both on techniques from trajectory optimization and deep visual model-based RL. We show experimentally that our approach achieves strong performance on challenging long-horizon visual robotic manipulation tasks where prior shooting-based approaches fail.

2. Related Work

Planning in model-based RL. Many recent papers on deep model-based RL (Chua et al., 2018; Ebert et al., 2018) use the cross-entropy method (Rubinstein & Kroese, 2004) to optimize action trajectories in a shooting formulation. Other work has explored different optimization methods such as the iterative linear-quadratic regulator (Levine & Koltun, 2013; Watter et al., 2015; Zhang et al., 2019), or Monte-Carlo tree search (Schrittwieser et al., 2020). However, these shooting approaches rely on local search in the space of actions, which is prone to local optima. Instead, our collocation approach optimizes over (latent) states as opposed to actions, which we show often enables us to escape local minima and plan better. Another line of work, inspired by classical sampling-based planning (Kavraki et al., 1996; LaValle, 1998), uses graph-based optimization (Kurutach et al., 2018; Savinov et al., 2018; Eysenbach et al., 2019; Liu et al., 2020) or other symbolic planners (Asai & Fukunaga, 2018) to optimize a sequence of states. While this escapes the local minima problem of shooting methods, such graph-based methods require constructing a large graph of possible states and scale poorly to combinatorial state spaces and context-dependent tasks. In contrast, latent collocation provides a principled control method that is able to optimize a sequence of states using continuous optimization, and does not suffer from the drawbacks of graph-based methods.

Recent work has designed hierarchical methods that plan over extended periods of time with intermediate subgoals, and then use separate model-free (Pong et al., 2018; Nasiriany et al., 2019) or model-based (Pertsch et al., 2020b; Nair & Finn, 2021; Pertsch et al., 2020a; Parascandolo et al., 2020) controllers to reach the subgoals. This can be considered a hierarchical form of collocation-based planning. However, in contrast to these approaches, which require a separate controller for reaching subgoals, we focus on the standard model-based RL setup where only a latent dynamics model is learned, and show that latent collocation is able to perform long-horizon tasks without hierarchical planning.

Many of the previously proposed model-based methods (Wahlström et al., 2015; Watter et al., 2015; Kurutach et al., 2018; Buesing et al., 2018; Zhang et al., 2019; Hafner et al., 2019) use latent state-space models for improved prediction quality and runtime. Our proposed method leverages

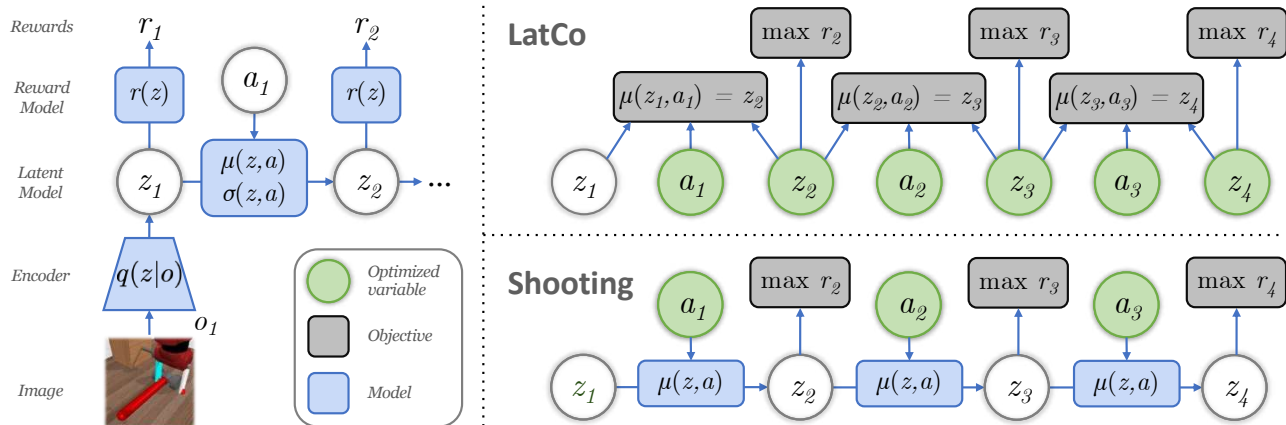


Figure 2: Latent Collocation (LatCo). **Left:** Our latent state-space model, with an encoder $q(z|o)$ and a latent state-space dynamics model $p(z_{t+1}|z_t, a_t) \sim \mathcal{N}(\mu(z_t, a_t), \sigma(z_t, a_t))$. A reward model $r(z_t)$ predicts the reward from the latent state. The model is trained with a variational lower bound to reconstruct the observations (not shown). **Right:** comparison of deterministic LatCo and shooting methods. LatCo optimizes a sequence of latent states and actions $z_{2:T}, a_{1:T}$ to maximize rewards $r(z_t)$ as well as satisfy dynamics $z_{t+1} = \mu(z_t, a_t)$. This joint optimization allows the dynamics constraint to be relaxed at first, which helps escape local minima. In contrast, shooting methods require recursive application of the dynamics and backpropagation through time, which is often difficult to optimize.

this latent state-space design to construct an effective trajectory optimization method with collocation, and we design our method to be model-agnostic, such that it can benefit from future improved latent variable models.

Collocation-based planning.¹ Collocation is a powerful technique for trajectory optimization (Hargraves & Paris, 1987; Witkin & Kass, 1988; Betts, 1998; Tedrake, 2021) that optimizes a sequence of states for the sum of expected reward, while eventually enforcing the constraint that the optimized trajectory conform to a dynamics model for some actions (also see Kelly (2017) for a recent tutorial). Prior work in optimal control and robotics has explored many variants of this approach, with Hamiltonian optimization (Ratliff et al., 2009), explicit handling of contacts (Mordatch et al., 2012; Posa et al., 2014), sequential convex programming (Schulman et al., 2014), as well as stochastic trajectory optimization (Kalakrishnan et al., 2011). Platt Jr et al. (2010); Patil et al. (2015) proposed probabilistic extensions of collocation. These works have demonstrated good results in controlling complex simulated characters, such as humanoid robots, contact-heavy tasks, and tasks with complex constraints. The optimization algorithm we use is most similar to that of Schulman et al. (2014), however, all this prior work assumed availability of a ground truth model of the environment and low-dimensional state descriptors.

Some recent works have attempted using collocation with learned neural network dynamics models (Bansal et al., 2016; Du et al., 2020), but only considered simple or low-

¹In this paper we use the terms “trajectory optimization” and “planning” synonymously as is common in MBRL literature.

dimensional dynamics. In this work, we address how to scale up collocation methods to high-dimensional image observations, where direct optimization over images is intractable, and the dynamics are more challenging to learn. We propose to do this by utilizing a learned latent space.

3. Latent Collocation (LatCo)

We aim to design a collocation method that plans trajectories from raw image observations. A naive approach would learn an image dynamics model, and directly optimize an image sequence using Equation 1. However, this is impractical for several reasons. First, optimizing over images directly is difficult due to the high dimensionality of the images and the fact that valid images lie on a thin manifold. Second, images typically do not constitute a Markovian state space. We propose to instead learn a Markovian and compact space by means of a latent variable model, and then use this learned latent space for collocation.

3.1. Learning Latent Dynamics

The design of dynamics models for visual observations is challenging. Recent work has proposed learning latent state-space models that represent the observations in a compact latent space z_t . Specifically, this work learns a latent dynamics model $p_\phi(z_{t+1}|z_t, a_t)$, as well as observation $p_\phi(o_t|z_t)$ and reward $r(z_t)$ decoders (see Figure 2, left). This approach is powerful due to high-capacity neural network latent dynamics models, and is computationally efficient as the latent space is compact. Importantly, the Markov property of the latent state is enforced, allowing a convenient

interpretation of the latent dynamics model as a Markov decision process. The model is shown in Figure 2 (left).

Many environments of interest are stochastic or partially observable, which necessitates accounting for uncertainty. The latent dynamics distribution $p_\phi(z_{t+1}|z_t, a_t)$ should then reflect the stochasticity in the observed data. To achieve this, we train the model by maximizing the likelihood of the observed data $r_{1:T}, o_{1:T}$. While maximizing exact likelihood is often intractable, we optimize a variational lower bound on it using a variational distribution $q_\phi(z_{t+1}|o_{t+1}, a_t, z_t)$ (Chung et al., 2015; Fraccaro et al., 2016):

$$\begin{aligned} \ln p_\phi(o_{2:T}, r_{1:T} | o_1, a_{1:T}, r_{1:T}) &\geq \mathcal{L}_{\text{ELBO}}(o_{1:T}, a_{1:T}, r_{1:T}) = \\ &\mathbb{E}_{q_\phi(z_{1:T}|o_{1:T}, a_{1:T}, z_0)} \sum_t [\ln p_\phi(o_{t+1}, r_{t+1} | z_{t+1}) - \\ &\text{KL}(q_\phi(z_{t+1}|o_{t+1}, a_t, z_t) || p_\phi(z_{t+1}|z_t, a_t))]. \end{aligned} \quad (2)$$

3.2. Latent Collocation with Probabilistic Dynamics

Given probabilistic latent-space dynamics $p_\phi(z_{t+1}|z_t, a_t)$, a reward function $r(z_t)$, and the current state z_1 , our method aims to select the actions with maximum expected return:

$$\max_{a_{1:T-1}} \mathbb{E}_{p_\phi(z_{t+1}|z_t, a_t)} \left[\sum_t r(z_t) \right]. \quad (3)$$

Prior work (Ebert et al., 2018; Hafner et al., 2019) used shooting methods (Betts, 1998; Tedrake, 2021) to directly optimize actions with this objective. However, this is known to be poorly conditioned due to the recursive application of the dynamics, which results in hard credit assignment. Instead, LatCo leverages the structure of the problem to construct an objective with only pairwise dependencies between temporally adjacent latents, and no recursive application.

We would like to formulate the trajectory optimization problem in Equation 3 as a constrained optimization problem, optimizing over sequences of both latent states and actions. However, this requires reformulating the problem in terms of probability distributions, since the model has stochastic dynamics and each observation corresponds to a distribution over latent states. To handle this, we formulate collocation with distributional constraints, analogously to belief-space planning (Platt Jr et al., 2010; Patil et al., 2015). This problem can be defined abstractly as optimizing a sequence of distributions $q(z_t)$, each representing an uncertain estimate of what the latent z_t will be at time t :

$$\begin{aligned} \max_{q(z_{2:T}), a_{1:T-1}} \sum_t \mathbb{E}_{q(z_t)} [r(z_t)] \\ \text{s.t. } q(z_{t+1}) &= \mathbb{E}_{q(z_t)} p_\phi(z_{t+1}|z_t, a_t). \\ q(z_1) &= p(z_1). \end{aligned} \quad (4)$$

When the constraint is satisfied, this is equivalent to the original problem in Equation 3. A simplified version of this

approach is illustrated in Figure 2 (right). We can express the distributional constraint in the form of a Bregman divergence or moment matching. For computational simplicity, we follow the latter approach

$$\begin{aligned} \max_{q(z_{2:T}), a_{1:T-1}} \sum_t \mathbb{E}_{q(z_t)} [r(z_t)] \\ \text{s.t. } \text{mean}[p(z_t)] &= \text{mean}[q(z_t)]. \\ \text{var}[p(z_t)] &= \text{var}[q(z_t)]. \end{aligned} \quad (5)$$

where we denoted with some abuse of notation $p(z_{t+1}) = \mathbb{E}_{q(z_t)} p_\phi(z_{t+1}|z_t, a_t)$. Further moments can be used for distributions with more degrees of freedom.

3.3. LatCo with Deterministic Plans

In many cases, finding a good plan without accounting for stochasticity is an effective and computationally efficient strategy. To do this, we will represent the plan distribution $q(z_{2:T})$ with particles $z_{2:T}^i$. In practice, we will simply use a single particle. The moment matching constraints from Equation 7 therefore reduce to a single constraint: $\text{mean}[p_\phi(z_t|z_{t-1}, a_{t-1})] = z_t$, while the variance constraint disappears since the variance of a set of one particle is a constant. This constraint, for $\mu_\phi(z_t, a_t) = \mathbb{E}_{p_\phi(z_{t+1}|z_t, a_t)} [z_{t+1}]$, yields a simplified planning problem:

$$\max_{z_{2:T}, a_{1:T-1}} \sum_t r(z_t) \quad \text{s.t. } z_{t+1} = \mu_\phi(z_t, a_t). \quad (6)$$

This approximation is equivalent to assuming that the underlying dynamics model is deterministic using its expected value, and performing standard deterministic collocation. This is a common assumption known as the certainty equivalence principle, and has appealing properties for certain kinds of distributions (Aoki, 1967; Bar-Shalom & Tse, 1974; Kwakernaak & Sivan, 1972). Other approximations with a single particle are possible, such as maximizing the likelihood of the particle, but these require introduction of additional balance hyperparameters, while our constrained optimization approach is principled and automatically tunes this balance. We visualize this approach in Figure 2.

3.4. LatCo with Gaussian Plans

The simple deterministic approximation from the previous subsection can perform well in environments that are close to deterministic, however, since it approximates the distribution $q(z_t)$ with just a single particle it is not expressive enough to represent the uncertainty in the plan. Instead, we can use more expressive distributions in our LatCo framework to perform planning under uncertainty. Specifically, since our latent dynamics $p_\phi(z_{t+1}|z_t, a_t)$ follow a Gaussian distribution, we parametrize q as a diagonal-covariance Gaussian $q(z_{1:T}) = \mathcal{N}(\mu_{1:T}, \sigma_{1:T}^2)$. More expressive distributions can be used, but we observed that

diagonal-covariance Gaussian already yields good performance on stochastic environments. To evaluate the moment matching terms in Equation 5 we use μ, σ^2 directly for the mean and variance of $q(z_t)$, while the mean and variance of $p(z_t)$ are estimated with samples. Gradients are estimated with reparametrization and 50 samples.

3.5. LatCo for Visual MBRL

We discussed two objectives that either represent the plan $q(z_t)$ deterministically with a single particle or with a Gaussian distribution. Either objective can be used in our framework for latent collocation. Under some regularity conditions, we can reformulate the equivalent dual version of Equation 5 as the saddle point problem on the Lagrangian

$$\begin{aligned} \min_{\lambda} \max_{q(z_{2:T}), a_{1:T-1}} \sum_t \mathbb{E}_{q(z_t)} [r(z_t)] \\ - \lambda (\|\text{mean}[p(z_t)] - \text{mean}[q(z_t)]\|^2 - \epsilon) \\ - \lambda (\|\text{var}[p(z_t)] - \text{var}[q(z_t)]\|^2 - \epsilon), \end{aligned} \quad (7)$$

by introducing Lagrange multipliers λ . In our implementation, we used squared constraints $\|\text{mean}[p(z_t)] - \text{mean}[q(z_t)]\|^2 = \epsilon$, however a non-squared constraint can also be used. Similar to dual descent (Nocedal & Wright, 2006), we address this min-max problem by taking alternating maximization and minimization steps, as also used with deep neural networks by Goodfellow et al. (2014); Haarnoja et al. (2018). While this strategy is not guaranteed to converge to a saddle point, we found it to work well in practice. Following prior work, we do not learn a terminal value function and simply use the reward of a truncated state sequence; however, a value function would be straightforward to learn with TD-learning on optimized plans. We summarize the main design choices below, deterministic LatCo in Algorithm 1, and provide further details in Appendix C.

Latent state models. We implement our latent dynamics model with convolutional neural networks for the encoder and decoder, and a recurrent neural network for the transition dynamics, following Hafner et al. (2019); Denton & Fergus (2018). The latent state includes a stochastic component with a conditional Gaussian transition function, and the hidden state of the recurrent neural network with deterministic transitions. The latent state space model predicts the reward directly from the latent state z_t using an approximation $r_\phi(z_t)$, so we never need to decode images during the optimization procedure, which makes it memory-efficient.

Dynamics relaxation. The dynamics constraint can be relaxed in the beginning of the optimization, leading LatCo to rapidly discover high-reward state space regions (potentially violating the dynamics), and then gradually modifying the trajectory to be more dynamically consistent, as illustrated in Figure 1. This is in contrast to shooting methods, which suffer from local optima in long horizon tasks, since the

Algorithm 1 Latent Collocation (LatCo)

- 1: Start with any available data \mathcal{D}
 - 2: **while** not converged **do**
 - 3: **for** each time step $t = 1 \dots T_{\text{tot}}$ **with step** T_{cache} **do**
 - 4: Infer latent state: $z_t \sim q(z_t|o_t)$
 - 5: Define the Lagrangian:

$$\begin{aligned} \mathcal{L}(z_{t+1:t+H}, a_{t:t+H}, \lambda) = \sum_t [r(z_t) \\ - \lambda_t^{\text{dyn}} (\|z_{t+1} - \mu(z_t, a_t)\|^2 - \epsilon^{\text{dyn}}) \\ - \lambda_t^{\text{act}} (\max(0, |a_t| - a_m)^2 - \epsilon^{\text{act}})]. \end{aligned} \quad (8)$$
 - 6: **for** each optimization step $k = 1 \dots K$ **do**
 - 7: Update plan:

$$z_{t+1:t+H}, a_{t:t+H} += \tilde{\nabla} \mathcal{L} \quad \triangleright \text{Eq (9)}$$
 - 8: Update dual variables:

$$\lambda_{t:t+H} := \text{UPDATE}(\mathcal{L}, \lambda_{t:t+H}) \quad \triangleright \text{Eq (10)}$$
 - 9: **end for**
 - 10: Execute $a_{t:t+T_{\text{cache}}}$ in environment:

$$o_{t:t+T_{\text{cache}}}, r_{t:t+T_{\text{cache}}} \sim p_{\text{env}}$$
 - 11: **end for**
 - 12: Add episode to replay buffer:

$$\mathcal{D} := \mathcal{D} \cup (o_{1:T_{\text{tot}}}, a_{1:T_{\text{tot}}}, r_{1:T_{\text{tot}}})$$
 - 13: **for** training iteration $i = 1 \dots I$ **do**
 - 14: Sample minibatch from replay buffer:

$$(o_{1:T}, a_{1:T}, r_{1:T})_{1:b} \sim \mathcal{D}$$
 - 15: Train dynamics model:

$$\phi += \alpha \nabla \mathcal{L}_{\text{ELBO}}(o_{1:T}, a_{1:T}, r_{1:T})_{1:b} \quad \triangleright \text{Eq (2)}$$
 - 16: **end for**
 - 17: **end while**
-

algorithm must simultaneously drive the states toward the high-reward regions *and* discover the actions that will get them there. The ability to disentangle these two stages by first finding the high-reward region and only then optimizing for the actions that achieve that reward allows LatCo to solve more complex and temporally extended tasks while suffering less from local optima. In our approach, this dynamics relaxation is not explicitly enforced, but is simply a consequence of using primal-dual constrained optimization.

MPC and Online training. LatCo is a planning method that can be used within an online model-based reinforcement learning training loop. In this setup, we use model predictive control (MPC) within a single episode, i.e. we carry out the plan only up to T_{cache} actions, and then re-plan to provide closed-loop behavior. Further, for online training, we perform several gradient updates to the dynamics model after every collected episode. This provides us with a model-based RL agent that can be trained either from scratch by only collecting its own data, or by seeding the replay buffer with any available offline data.




Figure 3: Sparse MetaWorld tasks, featuring temporally extended planning and sparse rewards. The agent observes the environment only through the visual input shown here. LatCo creates effective visual plans and performs well on all tasks.

4. Optimization for Latent Collocation

The latent collocation framework described in the previous section can be used with any optimization procedure, such as gradient descent or Adam (Kingma & Ba, 2015). However, we found that the choice of the optimizer for both the latent states and the Lagrange multipliers has a large influence on runtime. We detail our specific implementation below.

Levenberg-Marquardt optimization. We use the Levenberg-Marquardt optimizer for the states and actions, which pre-conditions the gradient direction with the matrix $(J^T J)^{-1}$, where J is the Jacobian of the objective with respect to states and actions. This preconditioner approximates the Hessian inverse, significantly improving convergence speed:

$$\tilde{\nabla} = (J^T J + \lambda I)^{-1} J^T \rho. \quad (9)$$

The Levenberg-Marquardt optimizer has cubic complexity in the number of optimized dimensions. However, by noting that the Jacobian of the problem is block-tridiagonal, we can implement a more efficient optimizer that scales linearly with the planning horizon (Mordatch et al., 2012). This efficient optimizer converges 10-100 times faster than gradient descent in wall clock in our experiments.

The Levenberg-Marquardt algorithm optimizes the sum of squares $\sum_i \rho_i^2$, defined in terms of residuals ρ . Any bounded objective can be expressed in terms of residuals by a combination of shifting and square root operations. For the dynamics constraint, we use the $z_{t+1} - \mu(z_t, a_t)$ differences as residuals directly, with one residual per state dimension. We similarly constrain the planned actions to be within the environment range a_m using the residual $\max(0, |a_t| - a_m)$. This corresponds to using a squared constraint instead of a linear one. For the reward objective, we form residuals as the softplus of the negative reward: $\rho = \ln(1 + e^{-r})$, which we found to be an effective way of forming a nonnegative cost without the need to estimate the maximum reward.

Constrained optimization. The naive gradient ascent update rule $\lambda += \|z_{t+1} - \mu(z_t, a_t)\|^2 - \epsilon$ for the multiplier λ works poorly when the current value of the multiplier is either much smaller or larger than the cost value. While it yields good plans, it suffers from slow convergence and suboptimal runtime. We can design a better behaved update rule by applying a monotonic function that rescales the

magnitude of the update while preserving the sign, which can be seen as a time-dependent learning rate. Specifically, we observed that scaling the update with the value of the multiplier itself λ as well as using log of the constrained value $\log \|z_{t+1} - \mu(z_t, a_t)\|^2 - \log \epsilon$ provided better scaled updates and led to faster convergence:

$$\lambda += \alpha \log \left(\frac{\|z_{t+1} - \mu(z_t, a_t)\|^2}{\epsilon} + \eta \right) \lambda, \quad (10)$$

where $\eta = 0.01$ ensures numerical stability and the learning rate $\alpha = 0.1$. Using a small non-zero ϵ is beneficial for the optimization and ensures fast convergence, as the exact constraint might be hard to reach.

5. Experiments

We evaluate long-horizon planning capabilities of LatCo for model-based reinforcement learning on several challenging manipulation and locomotion tasks. Each subsection below corresponds to a distinct scientific question that we study.

5.1. Experimental Setup

Environments. To evaluate on challenging visual planning tasks, we adapt the MetaWorld benchmark (Yu et al., 2020) to visual observations and sparse rewards, with a reward of 1 given only when the task is solved and 0 otherwise. Figure 3 shows example visual observations provided to the agent on these seven tasks. The robot and the object position are randomly initialized. The Thermos and Hammer tasks require a complex two-stage solution using tools to manipulate another object. We use $H = 30, T_{\text{cache}} = 30, T_{\text{tot}} = 150$ for all tasks except Pushing, Thermos, and Hammer, where we use $H = 50, T_{\text{cache}} = 25, T_{\text{tot}} = 150$. In addition, we evaluate on the standard continuous control tasks with shaped rewards from the DeepMind Control Suite (Tassa et al., 2020). According to the protocol from (Hafner et al., 2020), we use an action repeat of 2 and set $H = 12, T_{\text{cache}} = 6, T_{\text{tot}} = 1000$ for all DMC tasks.

Comparisons. To specifically ascertain the benefits of collocation, we must determine whether they stem from gradient-based optimization, optimizing over states, or both. Therefore, we include a prior method based on zeroth-order CEM optimization, *PlaNet* (Hafner et al., 2019); another

Table 1: MBRL results on the visual Sparse Metaworld tasks and DM Control. On sparse reward tasks, Shooting only solves the simpler tasks, while the powerful trajectory optimization with LatCo finds good trajectories more consistently.

	Reaching	Button	Window	Drawer	Pushing	Reacher Easy	Cheetah Run	Quadruped Walk
Shaped reward	×	×	×	×	×	✓	✓	✓
LatCo (Ours)	91 ± 3%	55 ± 4%	49 ± 8%	46 ± 3%	38 ± 3%	559 ± 15	245 ± 12	121 ± 9
PlaNet	20 ± 1%	13 ± 2%	31 ± 2%	22 ± 2%	22 ± 3%	561 ± 14	326 ± 3	72 ± 4
MPPI	16 ± 1%	10 ± 2%	30 ± 2%	21 ± 3%	21 ± 4%	657 ± 17	298 ± 1	57 ± 4
Shoot. GD	8 ± 1%	7 ± 0%	28 ± 1%	18 ± 2%	19 ± 3%	756 ± 7	246 ± 2	101 ± 2

Table 2: MBRL with offline and online data. Shooting fails to construct adequate plans on these challenging long-horizon tasks, while LatCo performs significantly better.

	Thermos	Hammer
Shaped reward	×	×
LatCo (Ours)	37 ± 21%	13 ± 2%
PlaNet	1 ± 1%	1 ± 0%
MPPI	0 ± 0%	0 ± 0%
Shoot. GD	0 ± 0%	3 ± 1%

sampling-based shooting method, *MPPI* (Williams et al., 2016; Nagabandi et al., 2020); as well as a gradient-based method that optimizes actions directly using the objective in Equation 3, which we denote *Shooting GD*. To isolate the effects of different planning methods, we use the same dynamics model architecture for all agents. We train all methods online according to Algorithm 1. The hyperparameters are detailed in Appendix A. We use an action repeat of 2 for Thermos and Hammer environments for all methods, and no action repeat on other tasks.

5.2. Is LatCo better able to solve sparse reward tasks in visual model-based reinforcement learning?

First, we evaluate LatCo’s performance in the standard model-based reinforcement learning setting, where the agent learns the model from scratch and collects the data online using the LatCo planner, according to Algorithm 1. We evaluate the deterministic LatCo and other model-based agents on the Sparse MetaWorld tasks from visual input without any state information. The performance is shown in Table 1 and learning curves in the Appendix. We observe that LatCo is able to learn a good model and construct effective plans on these sparse reward tasks. Shooting-based trajectory optimization methods, however, struggle, not being able to optimize a sparse reward signal that requires longer-horizon planning.

We further examine how the performance changes with the required planning horizon. We plot performance with respect to different distances to the goal in Figure 4. Shooting baselines degrade significantly on harder tasks that require longer horizons, while LatCo is able to solve even these harder tasks well. This confirms our hypothesis that collocation scales better to long-horizon tasks.

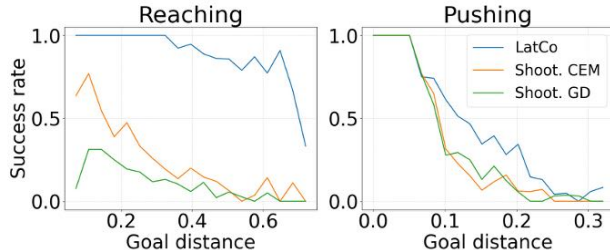


Figure 4: Performance with respect to task difficulty. We observe that LatCo maintains good performance even for harder tasks with long-horizon goals, whereas shooting is only able to solve easier tasks.

In addition to the sparse reward tasks, we further evaluate our method on the more standard tasks from DM control suite, also in Table 1 and dense MetaWorld tasks in App. Figure 10. We use the experimental protocol from (Hafner et al., 2020) and our results are consistent with their Fig. 10. Since these tasks generally have dense rewards that are easy to optimize, we do not expect a significant improvement from LatCo. However, LatCo performs competitively with other methods and outperforms them on some environments, showcasing its generality.

5.3. Is LatCo able to plan for and solve long-horizon robotic tasks from images?

Our aim is to evaluate LatCo on tasks with longer horizons than used in prior work (e.g., Ebert et al. (2018); Hafner et al. (2019)), focusing specifically on the performance of the planner. Solving such tasks from scratch also presents a challenging exploration problem, with all methods we tried failing to get non-zero reward, and is outside of the scope of this work. Therefore, to focus on comparing planning methods, we use a pre-collected dataset with some successful and some failed trajectories to jump-start all agents, similarly to the protocol of Rajeswaran et al. (2018); Nair et al. (2020). We initialize the replay buffer with this dataset, and train according to Algorithm 1 for the Thermos and Hammer tasks shown in Figure 3. Further, during evaluation, we reinitialize the optimization several times for LatCo and Shooting GD and select the best solution. These optimization runs are performed in parallel and do not significantly impact runtime. We do not reinitialize CEM as it already incorporates sampling in the optimization, and it requires a large batch size, making parallel initializations infeasible.




Figure 5: Planned and executed trajectories on the hammer and thermos tasks. Plans are visualized by passing the latent states through the convolutional decoder. Both tasks require picking up a tool and using it to manipulate another object. The sparse reward is only given after completing the full task, and the planner needs to infer all stages required to solve the task. LatCo produces feasible and effective plans, and is able to execute them on these challenging long-horizon tasks.

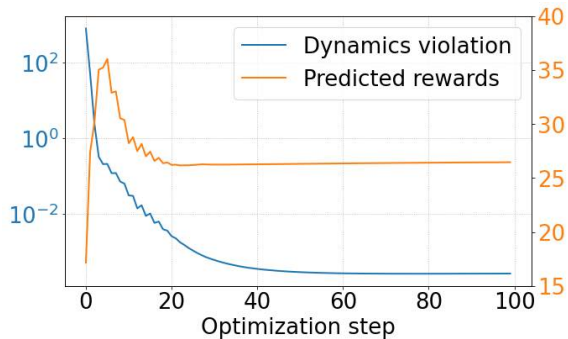


Figure 6: Dynamics violation and reward of the plan over 100 optimization steps. First, LatCo explores high-reward regions, converging to the goal state. As dynamics are enforced, the plan becomes both feasible and effective.

From Table 2, we observe that LatCo exhibits superior performance to shooting-based methods on both long-horizon tasks. These tasks are challenging for all of the methods, as the dynamics are more complex and they require accurate long-term predictions to correctly optimize the sparse signal. Specifically, since there is no partial reward for grasping the object, the planner has to look ahead into the high reward final state and reason backwards that it needs to approach the object. We observed that both of these tasks require a planning horizon of at least 50 steps for good performance. Shooting-based methods struggle to find a good plan on these tasks, often getting stuck in local minimum of not reaching for the object. LatCo outperforms shooting methods by a considerable margin, and is often able to construct and execute effective plans, as shown in Figure 5.

5.4. What is important for LatCo performance?

In this section, we analyze different design choices in LatCo. First, we analyze the ability to temporarily violate dynamics in order to effectively plan for long-term reward. We show the dynamics violation costs and the predicted rewards over the course of optimization quantitatively in Figure 6 and qualitatively in Figure 1. Since the dynamics constraint is only gradually enforced with the increase of the Lagrange

Table 3: Ablating LatCo components on the Button task.

	Success	Runtime (FPS)
LatCo	55%	1.6
LatCo no relaxation	38%	1.6
LatCo no constrained opt.	40%	1.6
LatCo no deterministic latent	43%	1.6
LatCo no second-order opt.	48%	0.1
Image Collocation	18%	0.2

multipliers, the first few optimization steps allow dynamics violations in order to focus on discovering high rewards (steps 0 to 10 in Figure 6). Later in the optimization, the constraints are optimized until the trajectory becomes dynamically feasible. If this ability to violate dynamics is removed by initializing the Lagrange multipliers λ to large values, the optimization performs similarly to shooting and struggles to discover rewards, as shown in App. Figure 11.

We further evaluate the importance of various design choices quantitatively in Table 3 through ablations. We see that LatCo without dynamics relaxation performs poorly, confirming our qualitative analysis above. Using a constant balance weight instead of Lagrange multipliers (LatCo no constrained opt.) requires extensive tuning of the weight, but can perform well with the optimal value of the weight. This highlights the importance of our constrained optimization framework that removes the need for this additional tuning. Using a latent dynamics model with only probabilistic states (LatCo no deterministic latent) degrades the performance slightly as this architecture produces inferior predictions, consistent with the findings of (Hafner et al., 2019). Using gradient descent instead of the Levenberg-Marquardt as well as for the Lagrange multiplier update (LatCo no second order) produces reasonable performance, but has much higher runtime complexity as it requires many more optimization steps. Optimizing images directly (Image Collocation) as opposed to optimizing latents performs better than shooting, but substantially worse than LatCo as the optimization problem is more complex.

5.5. Does probabilistic LatCo handle uncertainty well?

Deterministic LatCo evaluated in the previous sections performs well in standard environments as they do not require reasoning about uncertainty. However, uncertainty is important for many practical tasks (Kaelbling et al., 1998; Thrun, 1999; Van Den Berg et al., 2012). In this section, we show that deterministic LatCo and common shooting methods fail a simple lottery task, unlike Gaussian LatCo (Section 3.4).

The Lottery task has two goals (*Top* and *Bottom*). Reaching the Top goal provides 1 reward. Reaching the Bottom goal provides either 20 (65% chance) or -40 reward (35% chance). The agent needs to navigate to the goal using continuous $x - y$ actions, and only one goal can be chosen. An agent that treats uncertainty incorrectly might plan to achieve 20 reward with the Bottom goal, which leads to an expected reward of -1. The correct solution is the Top goal.

Figure 7 shows the performance of all methods, pretrained on a dataset of reaching either goal in equal proportion. As CEM chooses best sampled trajectories, it plans optimistically and prefers the incorrect Bottom goal. Deterministic LatCo and Shooting GD plan for the mean latent future, which is either optimistic or pessimistic depending on the latent space geometry. These methods visit both goals in similar proportions. Only Gaussian LatCo plans correctly by estimating the future state distribution preferring the Top goal. We further observed that while Gaussian LatCo only approximates the future using a Gaussian, this yields accurate prediction results in practice.

In Appendix E.2, we show that Gaussian LatCo performs comparably to the deterministic version on MetaWorld tasks. We observed similar performance on sparse reward tasks, but are unable to include the full experiment due to compute requirements. In contrast to commonly used shooting methods, LatCo performs well in both stochastic and deterministic environments, further showcasing its generality.

6. Discussion

Conclusion. We presented LatCo, a method for latent collocation that improves performance of visual-model based reinforcement learning agents. In contrast to the common shooting methods, LatCo uses powerful trajectory optimization to plan for long-horizon tasks where prior work fails. By improving the planning capabilities and removing the need for reward shaping, LatCo can scale to complex tasks more easily and with less manual instrumentation.

Limitations. As collocation usually requires many more optimization variables than shooting, it can be slower or take more optimization steps. While we address this with a problem-specific optimizer, future work might learn smooth latent spaces for easier planning. Further, we observed

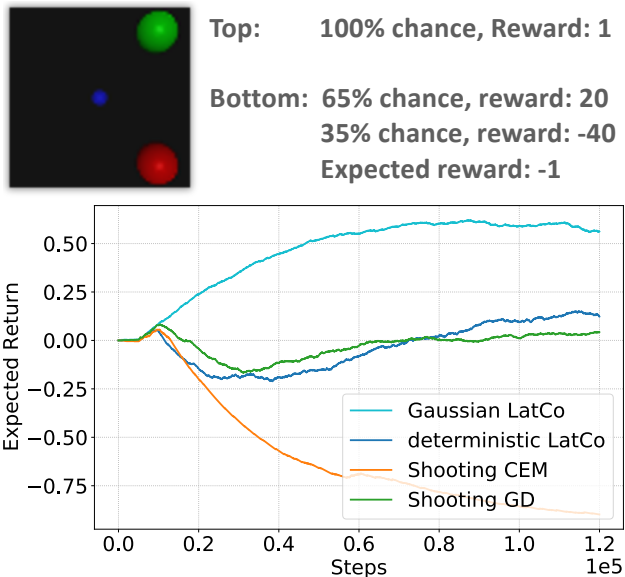


Figure 7: Planning under uncertainty on the Lottery task. Top: the environment description, bottom: training curves. Gaussian LatCo is able to correctly plan in this environment, while deterministic latco and other shooting baselines fail as they are not designed to plan under uncertainty. LatCo correctly estimates the expected cumulative reward, accounting for both optimistic and pessimistic outcomes.

collocation can still suffer from local optima, although less so than shooting. This issue may be addressed with even larger overparametrized latent states or better optimization.

Future work. By formulating a principled approach to optimizing sequences of latent states for RL, LatCo provides a useful building block for future work. Ideas presented in this paper may be applied to settings such as state-only imitation learning or hierarchical subgoal-based RL. With state optimization, state constraints in addition to rewards could be used for more flexible task specification. As LatCo does not require sampling from the dynamics model, it allows the use of more powerful models such as autoregressive, energy-based or score-based models for which sampling is expensive. Finally, the benefits of state optimization may translate into the policy learning setting by learning an amortized policy to output future states with the LatCo objective.

Acknowledgements

We thank Clark Zhang and Matt Halm for generous advice on trajectory optimization, and Aviral Kumar, Danijar Hafner, Dinesh Jayaraman, Edward Hu, Karl Schmeckpeper, Russell Mendonca, Deepak Pathak, and members of GRASP for discussions, as well as anonymous reviewers for helpful comments and suggestions. Support was provided by the ARL DCIST CRA W911NF-17-2-0181, ONR N00014-17-1-2093, and by Honda Research Institute.

References

- Aoki, M. *Optimization of stochastic systems: topics in discrete-time systems*. Academic Press, 1967.
- Asai, M. and Fukunaga, A. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. *Proceedings of AAAI Conference on Artificial Intelligence*, 2018.
- Bansal, S., Akametalu, A. K., Jiang, F. J., Laine, F., and Tomlin, C. J. Learning quadrotor dynamics using neural network for flight control. In *IEEE Conference on Decision and Control (CDC)*. IEEE, 2016.
- Bar-Shalom, Y. and Tse, E. Dual effect, certainty equivalence, and separation in stochastic control. *IEEE Transactions on Automatic Control*, 1974.
- Betts, J. T. Survey of numerical methods for trajectory optimization. *Journal of guidance, control, and dynamics*, 1998.
- Buesing, L., Weber, T., Racaniere, S., Eslami, S., Rezende, D., Reichert, D. P., Viola, F., Besse, F., Gregor, K., Hassabis, D., et al. Learning and querying fast generative models for reinforcement learning. *arXiv preprint arXiv:1802.03006*, 2018.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, 2018.
- Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A. C., and Bengio, Y. A recurrent latent variable model for sequential data. In *Advances in Neural Information Processing Systems*, 2015.
- Denton, E. and Fergus, R. Stochastic video generation with a learned prior. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Du, Y., Lin, T., and Mordatch, I. Model-based planning with energy-based models. In *Proceedings of the Conference on Robot Learning*, 2020.
- Ebert, F., Finn, C., Dasari, S., Xie, A., Lee, A., and Levine, S. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv:1812.00568*, 2018.
- Eysenbach, B., Salakhutdinov, R. R., and Levine, S. Search on the replay buffer: Bridging planning and reinforcement learning. In *Advances in Neural Information Processing Systems*, 2019.
- Finn, C., Goodfellow, I., and Levine, S. Unsupervised learning for physical interaction through video prediction. In *Advances in Neural Information Processing Systems*, 2016.
- Fraccaro, M., Sønderby, S. K., Paquet, U., and Winther, O. Sequential neural models with stochastic layers. In *Advances in Neural Information Processing Systems*, pp. 2199–2207, 2016.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. In *Advances in Neural Information Processing Systems*, 2014.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Hafner, D., Lillicrap, T., Ba, J., and Norouzi, M. Dream to control: learning behaviors by latent imagination. In *International Conference on Learning Representations*, 2020.
- Hargraves, C. R. and Paris, S. W. Direct trajectory optimization using nonlinear programming and collocation. *Journal of guidance, control, and dynamics*, 1987.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 1998.
- Kalakrishnan, M., Chitta, S., Theodorou, E., Pastor, P., and Schaal, S. Stomp: Stochastic trajectory optimization for motion planning. In *2011 IEEE international conference on robotics and automation*. IEEE, 2011.
- Kavraki, L. E., Svestka, P., Latombe, J.-C., and Overmars, M. H. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 1996.
- Kelly, M. An introduction to trajectory optimization: How to do your own direct collocation. *SIAM Review*, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- Kurutach, T., Tamar, A., Yang, G., Russell, S. J., and Abbeel, P. Learning plannable representations with causal infogan. In *Advances in Neural Information Processing Systems*, 2018.

- Kwakernaak, H. and Sivan, R. *Linear optimal control systems*, volume 1. Wiley-interscience New York, 1972.
- LaValle, S. M. Rapidly-exploring random trees: A new tool for path planning. 1998.
- Levine, S. and Koltun, V. Guided policy search. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.
- Liu, C. K. and Popović, Z. Synthesis of complex dynamic character motion from simple animations. *ACM Transactions on Graphics (TOG)*, 2002.
- Liu, K., Kurutach, T., Tung, C., Abbeel, P., and Tamar, A. Hallucinative topological memory for zero-shot visual planning. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Mordatch, I., Todorov, E., and Popović, Z. Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (TOG)*, 2012.
- Nagabandi, A., Konolige, K., Levine, S., and Kumar, V. Deep dynamics models for learning dexterous manipulation. In *Conference on Robot Learning*, 2020.
- Nair, A., Dalal, M., Gupta, A., and Levine, S. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.
- Nair, S. and Finn, C. Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation. *Proceedings of International Conference on Learning Representations (ICLR)*, 2021.
- Nasiriany, S., Pong, V., Lin, S., and Levine, S. Planning with goal-conditioned policies. In *Advances in Neural Information Processing Systems*, 2019.
- Nocedal, J. and Wright, S. *Numerical optimization*. Springer Science & Business Media, 2006.
- Parascandolo, G., Buesing, L., Merel, J., Hasenclever, L., Aslanides, J., Hamrick, J. B., Heess, N., Neitz, A., and Weber, T. Divide-and-conquer monte carlo tree search for goal-directed planning. *arXiv preprint arXiv:2004.11410*, 2020.
- Patil, S., Kahn, G., Laskey, M., Schulman, J., Goldberg, K., and Abbeel, P. Scaling up gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation. In *Algorithmic foundations of robotics XI*. Springer, 2015.
- Pertsch, K., Rybkin, O., Ebert, F., Finn, C., Jayaraman, D., and Levine, S. Long-horizon visual planning with goal-conditioned hierarchical predictors. In *Advances in Neural Information Processing Systems*, 2020a.
- Pertsch, K., Rybkin, O., Yang, J., Derpanis, K., Lim, J., Daniilidis, K., and Jaegle, A. Keyin: Discovering subgoal structure with keyframe-based video prediction. In *Annual Conference on Learning for Dynamics and Control*, 2020b.
- Platt Jr, R., Tedrake, R., Kaelbling, L., and Lozano-Perez, T. Belief space planning assuming maximum likelihood observations. In *Proceedings of Robotics: Science and Systems*, 2010.
- Pong, V., Gu, S., Dalal, M., and Levine, S. Temporal difference models: Model-free deep RL for model-based control. In *International Conference on Learning Representations*, 2018.
- Posa, M., Cantu, C., and Tedrake, R. A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research*, 2014.
- Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., and Levine, S. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. In *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- Ratliff, N., Zucker, M., Bagnell, J. A., and Srinivasa, S. Chomp: Gradient optimization techniques for efficient motion planning. In *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009.
- Rubinstein, R. Y. and Kroese, D. P. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*. Springer-Verlag New York, 2004.
- Savinov, N., Dosovitskiy, A., and Koltun, V. Semi-parametric topological memory for navigation. In *International Conference on Learning Representations*, 2018.
- Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 2020.
- Schulman, J., Duan, Y., Ho, J., Lee, A., Awwal, I., Bradlow, H., Pan, J., Patil, S., Goldberg, K., and Abbeel, P. Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 2014.
- Singh, A., Yang, L., Hartikainen, K., Finn, C., and Levine, S. End-to-end robotic reinforcement learning without reward engineering. In *Proceedings of Robotics: Science and Systems (RSS)*, 2019.
- Tassa, Y., Erez, T., and Todorov, E. Synthesis and stabilization of complex behaviors through online trajectory

- optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012.
- Tassa, Y., Tunyasuvunakool, S., Muldal, A., Doron, Y., Liu, S., Bohez, S., Merel, J., Erez, T., Lillicrap, T., and Heess, N. *dm_control: Software and tasks for continuous control*, 2020.
- Tedrake, R. Underactuated robotics: Algorithms for walking, running, swimming, flying, and manipulation (course notes for mit 6.832). 2021. <http://underactuated.mit.edu/>.
- Thrun, S. Monte carlo pomdps. In *Advances in Neural Information Processing Systems*, 1999.
- Van Den Berg, J., Patil, S., and Alterovitz, R. Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research*, 2012.
- Wahlström, N., Schön, T. B., and Deisenroth, M. P. From pixels to torques: Policy learning with deep dynamical models. *arXiv preprint arXiv:1502.02251*, 2015.
- Watter, M., Springenberg, J., Boedecker, J., and Riedmiller, M. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, 2015.
- Williams, G., Drews, P., Goldfain, B., Rehg, J. M., and Theodorou, E. A. Aggressive driving with model predictive path integral control. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016.
- Witkin, A. and Kass, M. Spacetime constraints. *ACM Siggraph Computer Graphics*, 1988.
- Yu, T., Quillen, D., He, Z., Julian, R., Hausman, K., Finn, C., and Levine, S. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, 2020.
- Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M., and Levine, S. Solar: Deep structured representations for model-based reinforcement learning. In *International Conference on Machine Learning*, 2019.