
PID-Inspired Inductive Biases for Deep Reinforcement Learning in Partially Observable Control Tasks

Ian Char

Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213
i`char`@cs.cmu.edu

Jeff Schneider

Machine Learning Department, Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
s`chneide`@cs.cmu.edu

Abstract

Deep reinforcement learning (RL) has shown immense potential for learning to control systems through data alone. However, one challenge deep RL faces is that the full state of the system is often not observable. When this is the case, the policy needs to leverage the history of observations to infer the current state. At the same time, differences between the training and testing environments makes it critical for the policy not to overfit to the sequence of observations it sees at training time. As such, there is an important balancing act between having the history encoder be flexible enough to extract relevant information, yet be robust to changes in the environment. To strike this balance, we look to the PID controller for inspiration. We assert the PID controller’s success shows that only summing and differencing are needed to accumulate information over time for many control tasks. Following this principle, we propose two architectures for encoding history: one that directly uses PID features and another that extends these core ideas and can be used in arbitrary control tasks. When compared with prior approaches, our encoders produce policies that are often more robust and achieve better performance on a variety of tracking tasks. Going beyond tracking tasks, our policies achieve 1.7x better performance on average over previous state-of-the-art methods on a suite of locomotion control tasks.¹

1 Introduction

Deep reinforcement learning (RL) holds great potential for solving complex tasks through data alone, and there have already been exciting applications of RL in video game playing [70], language model tuning [53], and robotic control [2]. Despite these successes, there still remain significant challenges in controlling real-world systems that stand in the way of realizing RL’s full potential [20]. One major hurdle is the issue of partial observability, resulting in a Partially Observable Markov Decision Process (POMDP). In this case, the true state of the system is unknown and the policy must leverage its history of observations. Another hurdle stems from the fact that policies are often trained in an imperfect simulator, which is likely different from the true environment. Combining these two challenges necessitates striking a balance between extracting useful information from the history and avoiding overfitting to modelling error. Therefore, introducing the right inductive biases to the training procedure is crucial.

The use of recurrent network architectures in deep RL for POMDPs was one of the initial proposed solutions [29] and remains a prominent approach for control tasks [47, 75, 50]. These architectures are certainly flexible; however, it is unclear whether they are the best choice for control tasks,

¹Code available at <https://github.com/IanChar/GPIDE>

especially since they were originally designed with other applications in mind such as natural language processing.

In contrast with deep RL methods, the Proportional-Integral-Derivative (PID) controller remains a cornerstone of modern control systems despite its simplicity and the fact it is over 100 years old [5, 48]. PID controllers are single-input single-output (SISO) feedback controllers designed for tracking problems, where the goal is to maintain a signal at a given reference value. The controller adjusts a single actuator based on the weighted sum of three terms: the current error between the signal and its reference, the integral of this error over time, and the temporal derivative of this error. PID controllers are far simpler than recurrent architectures and yet are still able to perform well in SISO tracking problems despite having no model for the system’s dynamics. We assert that PID’s success teaches us that in many cases only two operations are needed for successful control: summing and differencing.

To investigate this assertion, we conduct experiments on a variety of SISO and multi-input multi-output (MIMO) tracking problems using the same featurizations as a PID controller to encode history. We find that this encoding often achieves superior performance and is significantly more resilient to changes in the dynamics during test time. The biggest shortcoming with this method, however, is that it can only be used for tracking problems. As such, we propose an architecture that is built on the same principles as the PID controller, but is general enough to be applied to arbitrary control problems. Not only does this architecture exhibit similar robustness benefits, but policies trained with it achieve an average of *1.7x better performance* than previous state-of-the-art methods on a suite of locomotion control tasks.

2 Preliminaries

The MDP and POMDP We define the discrete time, infinite horizon Markov Decision Process (MDP) to be the tuple $(\mathcal{S}, \mathcal{A}, r, T, T_0, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, $T : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition function, $T_0 \subset \Delta(\mathcal{S})$ is the initial state distribution, and γ is the discount factor. We use $\Delta(\mathcal{S})$ to denote the space of distributions over \mathcal{S} . Importantly, the Markov property holds for the transition function, i.e. the distribution over a next state s' depends only on the current state, s , and current action, a . Knowing previous states and actions does not provide any more information. The objective is to learn a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that maximizes the objective $J(\pi) = \mathbb{E} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, s_{t+1})]$, where $s_0 \sim T_0$, $a_t \sim \pi(s_t)$, and $s_{t+1} \sim T(s_t, a_t)$. When learning a policy, it is often key to learn a corresponding value function, $Q^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, which outputs the expected discounted returns after playing action a at state s and then following π afterwards.

In a Partially Observable Markov Decision Process (POMDP), the observations that the policy receives are not the true states of the process. In control this may happen for a variety of reasons such as noisy observations made by sensors, but in this work we specifically focus on the case where aspects of the state space remain unmeasured. In any case, the POMDP is defined as the tuple $(\mathcal{S}, \mathcal{A}, r, T, T_0, \Omega, \mathcal{O}, \gamma)$, where Ω is the space of possible observations, $\mathcal{O} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\Omega)$ is the conditional distribution of seeing an observation, and the rest of the elements of the tuple remain the same as before. The objective remains the same as the MDP, but now the policy and value functions are not allowed access to the state.

Crucially, the Markov property does not hold for observations in the POMDP. That is, where $o_{1:t+1} := o_1, o_2, \dots, o_{t+1}$ are observations seen at times 1 through $t + 1$, $o_{1:t-1} \not\perp o_{t+1} | o_t, a_t$. A naive solution to this problem is to instead have the policy take in the history of the episode so far. Of course, it is usually infeasible to learn a policy that takes in the entire history for long episodes since the space of possible histories grows exponentially with the length of the episode. Instead, one can encode the information into a more compact representation. In particular, one can use an encoder ϕ which outputs an encoding $z_t = \phi(o_{1:t}, a_{1:t-1}, r_{1:t-1})$ (note that encoders need not always take in the actions and rewards). Then, the policy and Q-value functions are augmented to take in (o_t, z_t) and (o_t, a_t, z_t) , respectively.

Tracking Problems and PID Controllers. We first focus on the tracking problem, in which there are a set of signals that we wish to maintain at given reference values. For example, in espresso machines the temperature of the boiler (i.e. the signal) must be maintained at a constant reference temperature, and a controller is used to vary the boiler’s on-off time so the temperature

is maintained at that value [43]. Casting tracking problems as discrete time POMDPs, we let $o_t = (x_t^{(1)}, \dots, x_t^{(M)}, \sigma_t^{(1)}, \dots, \sigma_t^{(M)})$ be the observation at time t , where $x_t^{(i)}$ and $\sigma_t^{(i)}$ are the i^{th} signal and corresponding reference value, respectively. The reward at time t is simply the negative error summed across dimensions, i.e. $-\sum_{m=1}^M |x_t^{(m)} - \sigma_t^{(m)}|$.

When dealing with a single-input single-output (SISO) system (with one signal and one actuator that influences the signal), one often uses a Proportional-Integral-Derivative (PID) controller: a feedback controller that is often paired with feedforward control. This controller requires no knowledge of the dynamics, and simply sets the action via a linear combination of three terms: the error (P), **the integral of the error (I)**, and **the derivative of the error (D)**. When comparing other architectures to the PID controller, we will use **orange colored text** and **blue colored text** to highlight similarities between the I and D terms, respectively. Concretely, the policy corresponding to a discrete-time PID controller is defined as

$$\pi^{\text{PID}}(o_t) = K_P(x_t^{(1)} - \sigma_t^{(1)}) + K_I \sum_{i=1}^t (x_i^{(1)} - \sigma_i^{(1)}) dt + K_D \frac{(x_t^{(1)} - \sigma_t^{(1)}) - (x_{t-1}^{(1)} - \sigma_{t-1}^{(1)})}{dt} \quad (1)$$

where K_P , K_I , and K_D are scalar values known as gains that must be tuned. PID controllers are designed for SISO control problems, but many real-world systems are multi-input multi-output (MIMO). In the case of MIMO tracking problems, where there are M signals with M corresponding actuators, one can control the system with M separate PID controllers. However, this assumes there is a clear breakdown of which actuator influences which signal. Additionally, there are often interactions between the different signals, which the PID controllers do not account for. Beyond tracking problems, it is less clear how to use PID controllers without substantial engineering efforts.

3 Methodology

To motivate the following, consider the task of controlling a tokamak: a toroidal device that magnetically confines plasma and is used for nuclear fusion. Nuclear fusion holds the promise of providing an energy source with few drawbacks and an abundant fuel source. As such, there has recently been a surge of interest in applying machine learning [1, 13], and especially RL [14, 45, 17, 71, 65, 64, 44], for tokamak control. However, applying deep RL has the same problems as mentioned earlier; the state is partially observable since there are aspects of the plasma’s state that cannot be measured in real time, and the policy must be trained before-hand on an imperfect simulator since operation of the actual device is extremely expensive.

How should one choose a historical encoder with these challenges in mind? Previous works [50, 46] suggest using Long Short Term Memory (LSTM) [31], Gated Recurrent Units [15], or transformers [69]. These architectures have been shown to be powerful tools in natural language processing, where there exist complicated relationships between words and how they are positioned with respect to each other. However, do the same complex temporal relationships exist in something like tokamak control? The fact that PID controllers have been successfully applied for feedback control on tokamaks suggests this may not be the case [72, 27]. In reality, the extra flexibility of these architectures may become a hindrance when deployed on the physical device if they overfit to quirks in the simulator.

In this section, we present two historical encoders that we believe have good inductive biases for control. They are inspired by the PID controller in that they only sum and difference in order to combine information throughout time. Following this, in Section 5, we empirically show the benefits of these encoders on a number of control tasks including tokamak control.

The PID Encoder. Under the framework of a policy that uses a history encoder, the standard PID controller (1) is simply a linear policy with an encoder that outputs the tracking error, the integral of the tracking error, and the derivative of the tracking error. This notion can be extended to MIMO problems and arbitrary policy classes, resulting in the *PID-Encoder (PIDE)*. Given input $o_{1:t}$, this encoder outputs a $3M$ dimensional vector consisting of $(x_t^{(m)} - \sigma_t^{(m)})$, $\sum_{i=1}^t (x_i^{(m)} - \sigma_i^{(m)}) dt$, and $\frac{(x_t^{(m)} - \sigma_t^{(m)}) - (x_{t-1}^{(m)} - \sigma_{t-1}^{(m)})}{dt} \quad \forall m = 1, \dots, M$. For SISO problems, policies with this encoder have access to the same information as a PID controller. However, for MIMO problems the policy has

access to all the information that each PID controller, acting in isolation, would have. Ideally a sophisticated policy would coordinate each actuator setting well.

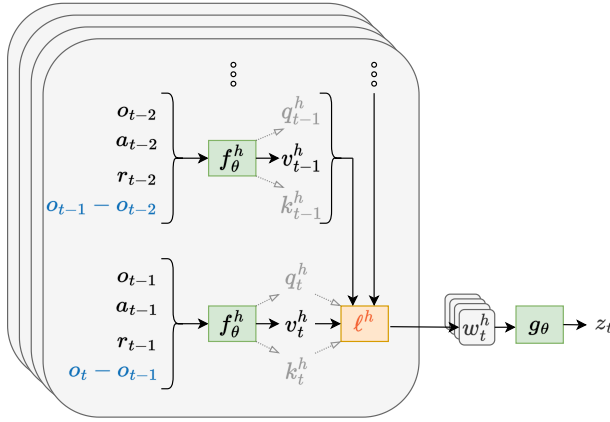


Figure 1: **Architecture for GPIDE.** The diagram shows how one encoding, z_t , is formed. Each of the gray, rounded boxes corresponds to one of the heads that makes up GPIDE. Each green box shows a function to be learned from data, and the orange box shows the weighted summation of all previous vectors, $v_{1:t}^h$. We write the difference in observations in blue text to highlight the part of GPIDE that relates to a PID controller’s D term. Note that $q_{1:t}^h$ and $k_{1:t}^h$ only play a role in this process if head h uses attention; as such, we write these terms in gray text.

The Generalized PID Encoder. A shortcoming of PIDE is that it is only applicable to tracking problems since it operates over tracking error explicitly. A more general encoder should instead accumulate information over arbitrary features of each observation. With this in mind, we introduce the *Generalized-PID-Encoder* (GPIDE).

GPIDE consists of a number of “heads”, each accumulating information about the history in a different manner. When there are H heads, GPIDE forms history encoding, z_t , through the following:

$$\begin{aligned} v_i^h &= f_\theta^h(\text{concatenate}(o_{i-1}, a_{i-1}, r_{i-1}, o_i - o_{i-1})) & \forall i \in \{1, \dots, t\}, h \in \{1 \dots, H\} \\ w_t^h &= \ell^h(v_{1:t}^h) & \forall h \in \{1 \dots, H\} \\ z_t &= g_\theta(\text{concatenate}(w_t^1, w_t^2, \dots, w_t^H)) \end{aligned}$$

Here, GPIDE is parameterized by θ . For head h , f_θ^h is a linear projection of the previous observation, action, reward, and difference between the current and previous observation to \mathbb{R}^D , and ℓ^h is a weighted summation of these projections. g_θ is a decoder which combines all of the information from the heads. A diagram of this process is shown in Figure 1. Note that θ is trained along with the policy and Q networks end-to-end with a gradient based optimizer.

Notice that the key aspects of the PID controller are present here. The **difference in observations** is explicitly taken before the linear projection f_θ^h . We found that this simple method works best for representing differences when the observations are scalar descriptions of the state (e.g. joint positions). Although we do not consider image observations in this work, we imagine a similar technique could be done by taking the differences in image encodings. Like the integral term of the PID, ℓ^h also accumulates information over time. In the following, we consider several possibilities for ℓ^h , and we will refer to these different choices as “head types” throughout this work. We omit the superscript h below for notational convenience.

Summation. Most in line with PID, the projections can be summed, i.e. $\ell(v_{1:t}) = \sum_{i=1}^t v_i$.

Exponential Smoothing. In order to weight recent observations more heavily, exponential smoothing can be used. That is, $\ell(v_{1:t}) = (1 - \alpha)^{t-1}v_1 + \sum_{i=2}^t \alpha(1 - \alpha)^{t-i}v_i$, where $0 \leq \alpha \leq 1$ is the smoothing parameter. Unlike summation, this head type cannot accumulate information in the same way because it is a convex combination.

Attention. Instead of hard-coding a weighted summation of the projections, this weighting can be learned through attention [69]. Attention is one of the key components of transformers because of its ability to learn relationships between tokens. To implement this, two additional linear functions should be learned that project $\text{concatenate}(o_{i-1}, a_{i-1}, r_{i-1}, o_i - o_{i-1})$ to \mathbb{R}^D . These new projections are referred to as they key and query vectors, denoted as k_i and q_i respectively. The softmax between their inner products is then used to form the weighting scheme for $v_{1:t}$. We can rewrite the first two

steps of GPIDE for a head that uses attention as

$$v_i, k_i, q_i = f_\theta(\text{concatenate}(o_{i-1}, a_{i-1}, r_{i-1}, o_i - o_{i-1})) \quad \forall i \in \{1, \dots, t\}$$

$$w_{1:t} = \ell(q_{1:t}, k_{1:t}, v_{1:t}) = \text{softmax} \left(\frac{q_{1:t} k_{1:t}^T}{\sqrt{D}} \right) v_{1:t}$$

Here, $q_{1:t}$, $k_{1:t}$, and $v_{1:t}$ are treated as $t \times D$ dimensional matrices. Since it results in a convex combination, attention has the capacity to reproduce exponential smoothing but not summation.

To anchor the GPIDE architecture back to the PID controller, we note that the P, I, and D terms can be formed exactly. At a high level, this is achieved when f_θ^h simply subtracts the target from the state measurement and when using summation and exponential smoothing heads with $\alpha = 1$. We write down the specific instance of GPIDE that results in these terms in Appendix A.1. While it is trivial for GPIDE to reconstruct the P, I, and D terms, it is less clear how an LSTM or GRU would achieve this, especially because of the I term. At the same time, GPIDE is much more flexible than the PID-representation since altering f_θ^h results in different representations at each time step and altering the type of head results in different temporal relationships.

4 Related Work

A control task may be partially observable for a myriad of reasons including unmeasured state variables [28, 75, 29], sensor noise [47], and unmeasured system parameters [76, 54]. When there are unmeasured system parameters, this is usually framed as a meta-reinforcement learning (MetaRL) [73] problem. This is a specific subclass of POMDPs where there is a collection of MDPs, and each episode, an MDP is sampled from this collection. Although these works do consider system parameters varying between episodes, the primary focus of the experiments usually tends to be on the multi-task setting (i.e. different reward functions instead of transition functions) [77, 18, 60]. We consider not only differing system parameters but also the presence of unmeasured state variables; therefore, the class of POMDPs considered in this paper is broader than the one studied in MetaRL.

Using recurrent networks has long been an approach for tackling POMDPs [29], and is still a common way to do so in a wide variety of settings [19, 73, 70, 50, 47, 75, 66, 11, 2]. Moreover implementations are publicly available both for on-policy [41, 30] and off-policy [50, 75, 11] algorithms, making it an easy pick for those wanting a quick solution. Some works [32, 77, 28, 18, 3] use recurrent networks to estimate the belief state [37], which is a distribution over the agent’s true state. However, Ni et al. [50] recently showed that well-implemented, recurrent versions of SAC [26] and TD3 [23] perform competitively with many of these specialized algorithms. In either case, we believe works that estimate the belief state are not in conflict with our own since their architectures can be modified to use GPIDE instead of a recurrent unit.

Beyond recurrent networks, there has been a surge of interest in applying transformers to reinforcement learning [40]. However, we were unable to find many instances of transformers being used as history encoders in the online setting, perhaps because of their difficulty to train. Parisotto et al. [55] introduced a new architecture to remedy these difficulties; however, Melo [46] applied transformers to MetaRL and asserted that careful weight initialization is the only thing needed for stability in training. We note that GPIDE with only attention heads is similar to a single multi-headed self-attention block that appears in many transformer architectures; however, we show that attention is the least important type of head in GPIDE and often hurts performance (see Section 5.3).

Perhaps closest to our proposed architecture is PEARL [60], which does a multiplicative combination of Gaussian distributions corresponding to each state-action-reward tuple. However, their algorithm is designed for the MetaRL setting specifically. Additionally, we note that the idea of summations and averaging has been shown to be powerful in prior works. Specifically, Oliva et al. [52] introduced the Statistical Recurrent Unit, an alternative architecture to LSTMs and GRUs that leverages moving averages and performs competitively across several supervised learning tasks.

There are many facets of RL where improvements can be made to robustness, and many works focus on altering the training procedure. They use techniques such as optimizing the policy’s worst-case performance [59, 36] or using variational information bottlenecks (VIB) [4] to limit the information used by the policy [42, 33, 21]. In contrast, our work specifically focuses on how architecture choices of history encoders affect robustness, but we note our developments can be used in conjunctions with

these other directions, possibly resulting in improved robustness. We perform additional experiments that consider VIB in Appendix F.1.

Lastly, we note that there is a plethora of work interested in the intersection of reinforcement learning and PID control [35, 25, 39, 22, 74, 12]. These works focus on using reinforcement learning to tune the coefficients of PID controllers (often in MIMO settings). We view these as important works on how to improve PID control using reinforcement learning; however, we view our own work as how to improve deep reinforcement learning by leveraging ideas from PID control.

5 Experiments

In this section, we experimentally compare PIDE and GPIDE against recurrent and transformer encoders. In particular, we explore the following questions:

- How does the performance of a policy using PIDE or GPIDE do on tracking problems? In addition, how well can policies adapt to different system parameters and how robust to modelling error are they on these problems? (Section 5.1)
- Going beyond tracking problems, how well does GPIDE perform on higher dimensional locomotion control tasks (Section 5.2)
- How important is each type of head in GPIDE? (Section 5.3)

For the following tracking problems we use the Soft Actor Critic (SAC) [26] algorithm with each of the different methods for encoding observation history. Following Ni et al. [50], we make two separate instantiations of the encoders for the policy and value networks, respectively. Since the tracking problems are relatively simple, we use a small policy network consisting of 1 hidden layer with 24 units; however, we found that we still needed to use a relatively large Q network consisting of 2 hidden layers with 256 units each to solve the problems. All hyperparameters remain fixed across baselines and tracking tasks; only the history encoders change.

For the recurrent encoder, we use a GRU and follow the implementation of Ni et al. [50] closely. Our transformer encoder closely resembles the GPT2 architecture [58], and it also includes positional encodings for the observation history. For GPIDE, we use $H = 6$ heads: one summation head, two attention heads, and three exponential smoothing heads (with $\alpha = 0.25, 0.5, 1.0$). This choice was not optimized, but rather was picked so that all types of heads were included and so that GPIDE has roughly the same amount of parameters as our GRU baseline. As a reference point for these RL methods, we also evaluate the performance of a tuned PID controller. Not only do PID controllers have an incredibly small number of parameters compared to the other RL-based controllers, but the training procedure is also much more straightforward since it can be posed as a black-box optimization over the returns. While there exists many sophisticated extensions of the PID controller (especially in MIMO systems [8]), we only consider the vanilla PID controller since we believe it serves as a good reference point. All methods are built on top of the rllkit library [57]. More details about implementations, hyperparameters, and computation can be found in Appendices B, C, and D, respectively. We also include additional experiments regarding variational information bottlenecks (VIB) and lookback size ablations in Appendices F.1 and F.2. Implementations can be found at <https://github.com/IanChar/GPIDE>.

5.1 Tracking Problems

In this subsection we consider a number of tracking problems. For each environment, the observation consists of the current signals, the reference values, and additional information about the last action made. Unless stated otherwise, the reward is as described in Section 2. More information about environments can be found in Appendix E. To make a fair comparison against PID controls, we choose to only encode the history of observations. For evaluation, we use 100 fixed settings of the environment (each setting consists of targets and system parameters). To avoid overfitting to these 100 settings, we used a separate set of 100 settings and averaged over 3 seeds when developing our methods. We evaluate policies throughout training, but report the average over the last 10% of evaluations as the final returns. We allow each policy to collect one million environment transitions, and all scores are averaged over 5 seeds. Lastly, each table shows scores formed by scaling the returns by the best and worst average returns across all methods in a particular variant of the environment, where scores of 0 and 100 correspond to the worst and best returns respectively.

Environment (Train/Test)	PID Controller	GRU	Transformer	PIDE	GPIDE
MSD Fixed/Fixed	0.00 ± 3.96	83.73 ± 3.48	85.79 ± 1.98	100.00 ± 0.66	83.72 ± 2.86
MSD Small/Small	0.00 ± 5.58	100.00 ± 1.59	73.27 ± 4.98	75.51 ± 1.31	80.21 ± 8.59
MSD Fixed/Large	36.58 ± 2.86	0.00 ± 3.42	53.70 ± 1.71	34.92 ± 0.93	29.55 ± 2.32
MSD Small/Large	43.52 ± 2.82	87.63 ± 2.28	81.44 ± 0.82	53.21 ± 1.31	68.03 ± 4.43
MSD Large/Large	45.60 ± 1.71	100.00 ± 0.61	92.60 ± 1.49	69.88 ± 0.69	93.03 ± 1.27
Average	25.14	74.27	77.36	66.70	70.91
DMSD Fixed/Fixed	24.33 ± 3.97	0.00 ± 8.69	22.05 ± 3.58	100.00 ± 1.08	76.23 ± 6.26
DMSD Small/Small	16.17 ± 3.09	0.00 ± 7.79	43.74 ± 3.70	100.00 ± 0.94	86.74 ± 3.94
DMSD Fixed/Large	63.59 ± 2.91	0.00 ± 2.28	59.84 ± 1.13	78.77 ± 1.16	63.89 ± 2.16
DMSD Small/Large	70.35 ± 1.44	39.26 ± 2.37	73.81 ± 1.60	88.52 ± 0.83	89.66 ± 1.33
DMSD Large/Large	78.77 ± 1.97	52.01 ± 2.01	84.45 ± 1.41	86.90 ± 0.18	100.00 ± 0.91
Average	50.64	18.25	56.78	90.84	83.30
Total Average	37.89	46.26	67.07	78.77	77.11

Table 1: **Mass Spring Damper Task Results.** The scores presented are averaged over five seeds and we show the standard error for each score.

Mass Spring Damper Tracking The first tracking task is the control of a classic 1D toy physics system in which there is a mass attached to a wall by a spring and damper. The goal is then to apply a force to the mass in order to move it to a given reference location. There are three system parameters to consider here: the mass, spring constant, and damping factor. We also consider the substantially more difficult problem in which there are two masses sandwiched between two walls, and the masses are connected to the walls and each other by springs and dampers (see Appendix E.1 for a diagram of this). Overall there are eight system parameters (three spring constants, three damping factors, and two masses) and two actuators (a force applied to each mass). We refer to the first problem as Mass-Spring-Damper (MSD) and the second problem as Double-Mass-Spring-Damper (DMSD).

Additionally, we test how adaptive these policies are by changing system parameters in a MetaRL-type fashion (i.e. for each episode we randomly select system parameters and then fix them for the rest of the episode). Similar to Packer et al. [54], we train the policies on three versions of the environment: one with no variation in system parameters, one with a small amount of variation, and one with a large amount of variation. We evaluate all policies on the version of the environment with large system parameter variation to test generalization capabilities.

Table 1 shows the scores achieved for each of the settings. While GRU and transformers seem to do a good job at encoding history for the MSD environment, both are significantly worse on the more complex DMSD task when compared to our proposed encoders. This is true especially for GRU, which performs worse than two independent PID controllers for every configuration. Additionally, while it seems that GRU can generalize to large amounts of variation in system parameters when a small amount is present, it fails horribly when trained on fixed system parameters. On the other hand, transformers are able to generalize surprisingly well when trained on both fixed system parameters and with small variation. We hypothesize the autoregressive nature of GRU may make it particularly susceptible to overfitting. Comparing PIDE and GPIDE, we see that PIDE tends to shine in the straightforward cases where there is little change in system parameters, whereas GPIDE is able to adapt when there is a large variation in parameters since it has additional capacity.

Navigation Environment To emulate the setting where the policy is trained on an imperfect simulator, we consider an environment in which the agent is tasked with moving itself across a surface to a specified 2D target as quickly and efficiently as possible. At every point in time, the agent can apply some force to move itself, but a penalty term proportional to the magnitude of the force is subtracted from the reward. Suppose that we have access to a simulator of the environment that is perfect except for the fact that it does not model friction between the agent and the surface. We refer to this simulator and the real environment as the “No Friction“ and “Friction” environment, respectively. In both environments, the mass of the agent is treated as a system parameter that is sampled for each episode; however, the Friction environment has a larger range of masses and also randomly samples the coefficient of friction each episode.

Figure 2 shows the average returns recorded during training for both navigation environments and when the policies trained in No Friction are evaluated in Friction. A table of final scores can be found in Appendix G.3. One can see that GPIDE not only achieves the best returns in the environments it was trained in, but is also robust when going from the frictionless environment to the one with friction. On the other hand, PIDE has less capacity and therefore cannot achieve the same results; however, it




Figure 2: **Average Returns for Navigation Environments.** The curves show the average over five seeds and the shaded region shows the standard error. For this plot, we allowed for 5x the normal amount budget to allow all methods to converge. We omit the PID controllers from this plot since it gets substantially worse returns.

is immediately more robust than the other methods, although it begins to overfit over time. It is also clear that using GRU is less sample efficient and less robust to changes in the test environment.

Tokamak Control For our last tracking experiment we return to tokamak control. In particular, we focus on the DIII-D tokamak, a device operated by General Atomics in San Diego, California. We aim to control two quantities: β_N , the normalized ratio between plasma and magnetic pressure, and rotation, i.e. how fast the plasma is spinning around the toroid. These are important quantities to track because β_N serves as an approximate economic indicator and rotation control of the plasma has been suggested to be key for stability [7, 67, 10, 61, 56]. The policy has control over the eight neutral beams [24], which are able to inject power and torque by blasting neutrally charged particles into the plasma. Importantly, two of the eight beams can be oriented in the opposite direction from the others, which decouples the total combined power and torque to some extent (see Figure 3).




Figure 3: **Illustration of DIII-D from Above.** [14] Each beamline in the figure contains two independent beams (yellow boxes). The plasma is rotating counter-clockwise and the two beams in the bottom left of the figure are oriented in the counter-current direction, allowing power and torque to be decoupled. This figure gives a rough idea of beam positioning but is not physically accurate.

To emulate the sim-to-real training experience, we create a simulator based on the equations described in Boyer et al. [9] and Scoville et al. [63]. This simulator has two major shortcomings: it assumes that certain states of the plasma (e.g. its shape) are fixed for entire episodes, and it assumes that there are no events that cause loss of confinement of the plasma. We make up for part of the former by randomly sampling plasma states each episode. The approximate

Environment (Train/Test)	PID Controller	GRU	Transformer	PIDE	GPIDE
β_N -Track Sim/Sim	40.69 ± 0.32	100.00 ± 0.20	97.56 ± 0.19	0.00 ± 1.05	98.33 ± 0.41
β_N -Track Sim/Real	89.15 ± 0.99	40.96 ± 5.45	40.05 ± 11.91	0.00 ± 21.04	55.21 ± 4.44
β_N -Track Real/Real	98.45 ± 0.77	98.24 ± 0.38	98.74 ± 0.29	100.00 ± 0.23	99.30 ± 0.64
Average	76.10	79.73	78.79	33.33	84.28
β_N -Rot-Track Sim/Sim	0.00 ± 0.83	99.06 ± 0.22	96.22 ± 0.94	67.98 ± 0.50	100.00 ± 0.29
β_N -Rot-Track Sim/Real	83.71 ± 2.64	39.76 ± 5.84	33.31 ± 0.69	0.00 ± 8.89	51.00 ± 1.92
β_N -Rot-Track Real/Real	92.02 ± 0.84	98.34 ± 0.52	96.32 ± 0.31	98.21 ± 0.23	100.00 ± 0.46
Average	58.58	79.05	75.28	55.40	83.67
Total Average	67.34	79.39	77.03	44.36	83.97

Table 2: **Tokamak Control Task Results.** The scores presented are averaged over five seeds and we show the standard error for each score.

Environment	PPO-GRU	TD3-GRU	VRM	SAC-Transformer	SAC-GPIDE
HalfCheetah-P	27.09 ± 7.85	85.80 ± 5.15	-107.00 ± 1.39	37.00 ± 9.97	82.63 ± 3.46
Hopper-P	49.00 ± 5.22	84.63 ± 8.33	3.53 ± 1.63	59.54 ± 19.64	93.27 ± 13.56
Walker-P	1.67 ± 4.39	29.08 ± 9.67	-3.89 ± 1.25	24.89 ± 14.80	96.61 ± 1.60
Ant-P	39.48 ± 3.74	-36.36 ± 3.35	-36.39 ± 0.17	-10.57 ± 2.34	66.66 ± 2.94
HalfCheetah-V	19.68 ± 11.71	59.03 ± 2.88	-80.49 ± 2.97	-41.31 ± 26.15	20.39 ± 29.60
Hopper-V	13.86 ± 4.80	57.43 ± 8.63	10.08 ± 3.51	0.28 ± 8.49	90.98 ± 4.28
Walker-V	8.12 ± 5.43	-4.63 ± 1.30	-1.80 ± 0.70	-8.21 ± 1.31	36.90 ± 16.59
Ant-V	1.43 ± 3.26	17.03 ± 6.55	-13.41 ± 0.12	0.81 ± 1.31	18.03 ± 5.10
Average	20.04	36.50	-28.67	7.80	63.18

Table 3: **PyBullet Task Results.** Each score is averaged over four seeds and we report the standard errors. Unlike before, we scale the returns by the returns of an oracle policy (i.e. one which sees position and velocity) and a policy which does not encode any history. For the environment names, “P” and “V” denote only position or only velocity in the observation, respectively.

“real” environment addresses these shortcomings by using a data-driven simulator. This approach to simulating has been shown to be relatively accurate [14, 65, 64, 1], and we use an adapted version of the simulator appearing in Char et al. [14] for our work. This simulator accounts for a greater set of the plasma’s state, and the additional information is rich enough that loss of confinement events play a role in the dynamics.

We consider two versions of this task: the first is a SISO task where total power is controlled to achieve a β_N target, and the second is a MIMO task where total power and torque is controlled to achieve β_N and rotation targets. The results for both of these tasks are shown in Table 2. Most of the RL techniques are able to do well if tested in the same environment they were trained in; the exception of this is PIDE, which curiously is unable to perform well in the simulator environment. While no reinforcement learning method matches the robustness of a PID controller, policies trained with GPIDE fare significantly better.

5.2 PyBullet Locomotion Tasks

Moving past tracking problems, we evaluate GPIDE on the PyBullet [16] benchmark proposed by Han et al. [28] and adapted in Ni et al. [50]. The benchmark has four robots: halfcheetah, hopper, walker, and ant. For each of these, either the current position information or velocity information is hidden from the agent. Except for GPIDE and transformer encoder, we use all of the performance traces given by Ni et al. [50]. In addition to SAC, they also train using PPO [62], A2C [49], TD3 [23], and VRM [28], a variational method that uses recurrent units to estimate the belief state. We reproduce as much of the training and evaluation procedure as possible, including using the same hyperparameters in the SAC algorithm and giving the history encoders access to actions and rewards. For more information see Appendix C.2. Table 3 shows the performance of GPIDE along with a subset of best performing methods (more results can be found in Appendix G.5). These results make it clear that GPIDE is powerful in arbitrary control tasks besides tracking since the average score achieved across all tasks is a 73% improvement over TD3-GRU, which we believe is the previous state-of-the-art for this benchmark at the time of this work.

Moreover, GPIDE dominates performance for every robot except HalfCheetah. The only setting where GPIDE achieves significantly worse performance is HalfCheetah-V. This setting seemed to cause stability issues in some seeds, lowering the average score. We believe that these stability issues stem from the attention heads, and we found that removing these heads fixed stability issues and resulted in a competitive average score (see Appendix G.5).

5.3 GPIDE Ablations

To investigate the role of each type of head, we reran all experiments with three variants of GPIDE: one with six exponential smoothing heads (ES), one with five exponential smoothing heads and

	MSD	DMSD	Navigation	β_N Track	β_N -Rot Track	PyBullet
ES	+2.69%	-11.14%	-0.11%	+2.57%	+0.29%	+5.81%
ES + Sum	-8.33%	+5.49%	-1.65%	+4.22%	+0.76%	+11.00%
Attention	-0.36%	-54.95%	-3.91%	-8.85%	-7.55%	-39.44%

Table 4: **GPIDE Ablation Percent Difference for Average Scores.** All final scores can be found in Appendix G.

one summation head (ES + Sum), and one with six attention heads (see Appendix C.3 for details). We choose these three configurations specifically to better understand the roles that attention and summation play.

Table 4 shows the differences in the average scores for each environment. The first notable takeaway is that having summation is often important in some of the more complex environments. The other takeaway is that much of the heavy lifting is being done by the exponential smoothing. GPIDE fares far worse when only having attention heads, especially in DMSD and the PyBullet environments.

We visualize some of the attention schemes learned by GPIDE for MSD with small variation and HalfCheetah (Figure 4). While the attention scheme learned for MSD could potentially be useful since it recalls information from near the beginning of the episode when the most movement is happening, it appears that the attention scheme for HalfCheetah is simply a poor reproduction of exponential smoothing, making it redundant and suboptimal. In fact, we found this phenomenon to be true across all attention heads and PyBullet tasks. We believe that the periodicity that appears here is due to the oscillatory nature of the problem and lack of positional encoding (although we found including positional encoding degrades performance).




Figure 4: **Averaged Attention Schemes for MSD-Small and HalfCheetah-P.** Each y-position on the grid corresponds to an amount of history being recorded, and each x-position corresponds to a time point in that history. As such, each of the left-most points are the oldest observation in the history, and the diagonals correspond to the most recent observation. The darker the blue, the greater the weight that is assigned to that time point.

6 Discussion

In this work, we introduced the PIDE and GPIDE history encoders to be used for reinforcement learning in partially observable control tasks. Although both are far simpler than prior methods of encoding, they often result in powerful yet robust controllers. We hope that this work inspires the research community to think about how pre-existing control methods can inform architecture choices.

Limitations There are many different ways a control task may be partially observable, and we do not believe that our proposed methods are solutions to all of them. For example, we do not think GPIDE is necessarily suited for tasks where the agent needs to remember events (e.g. picking up a key to unlock a door).

As with any bias, the PID-inspired biases that we propose in this work come at the cost of flexibility. For the experiments considered in this work, this trade off is beneficial and results in better policies. However, it is unclear whether this trade off is always worth making. It is possible that in higher dimensional environments or environments with more complex dynamics that having more flexibility is preferable to our proposed architecture.

Lastly, some tasks may require the policy to act on images as observations. We are optimistic that PIDE and GPIDE are still useful architectures in this setting, but we speculate that this is contingent on training an image encoder that is well-suited for these architectures, and we leave this research direction for future work.

7 Acknowledgements

We would like to thank Conor Igoe for his helpful discussions and advice on this work. We would also like to thank the NeurIPS 2023 reviewers assigned to our paper for their discussion and feedback.

This work was supported in part by US Department of Energy grants under contract numbers DE-SC0021414 and DE-AC02-09CH1146. This work is also supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1745016 and DGE2140739. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] Joseph Abbate, R Conlin, and E Kolemen. Data-driven profile prediction for diii-d. *Nuclear Fusion*, 61(4):046027, 2021.
- [2] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on Robot Learning*, pages 403–415. PMLR, 2023.
- [3] Kei Akuzawa, Yusuke Iwasawa, and Yutaka Matsuo. Estimating disentangled belief about hidden state and hidden task for meta-reinforcement learning. In *Learning for Dynamics and Control*, pages 73–86. PMLR, 2021.
- [4] Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.
- [5] Lucas Alpi. Pid control: Breaking the time barrier, November 2019. URL https://www.novusautomation.com/en/article_PID_control#:~:text=In%201911%20the%20first%20PID,still%20widely%20used%20in%20automation.
- [6] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [7] Laszlo Bardoczi, NC Logan, and EJ Strait. Neoclassical tearing mode seeding by nonlinear three-wave interactions in tokamaks. *Physical Review Letters*, 127(5):055002, 2021.
- [8] Stephen Boyd, Martin Hast, and Karl Johan Åström. Mimo pid tuning via iterated lmi restriction. *International Journal of Robust and Nonlinear Control*, 26(8):1718–1731, 2016.
- [9] MD Boyer, KG Erickson, BA Grierson, DC Pace, JT Scoville, J Rauch, BJ Crowley, JR Ferron, SR Haskey, DA Humphreys, et al. Feedback control of stored energy and rotation with variable beam energy and perveance on diii-d. *Nuclear Fusion*, 59(7):076004, 2019.
- [10] RJ Buttery, RJ La Haye, P Gohil, GL Jackson, H Reimerdes, EJ Strait, and DIII-D Team. The influence of rotation on the β_n threshold for the 2/1 neoclassical tearing mode in diii-d. *Physics of Plasmas*, 15(5):056115, 2008.
- [11] Massimo Caccia, Jonas Mueller, Taesup Kim, Laurent Charlin, and Rasool Fakoor. Task-agnostic continual reinforcement learning: In praise of a simple baseline. *arXiv preprint arXiv:2205.14495*, 2022.
- [12] Ignacio Carlucho, Mariano De Paula, and Gerardo G Acosta. An adaptive deep reinforcement learning approach for mimo pid control of mobile robots. *ISA transactions*, 102:280–294, 2020.
- [13] Ian Char, Youngseog Chung, Willie Neiswanger, Kirthevasan Kandasamy, Andrew O Nelson, Mark Boyer, Egemen Kolemen, and Jeff Schneider. Offline contextual bayesian optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- [14] Ian Char, Joseph Abbate, László Bardóczi, Mark Boyer, Youngseog Chung, Rory Conlin, Keith Erickson, Viraj Mehta, Nathan Richner, Egemen Kolemen, et al. Offline model-based reinforcement learning for tokamak control. In *Learning for Dynamics and Control Conference*, pages 1357–1372. PMLR, 2023.
- [15] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [16] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021.
- [17] Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897): 414–419, 2022.

- [18] Ron Dorfman, Idan Shenfeld, and Aviv Tamar. Offline meta learning of exploration. *arXiv preprint arXiv:2008.02598*, 2020.
- [19] Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- [20] Gabriel Dulac-Arnold, Nir Levine, Daniel J Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Machine Learning*, 110(9):2419–2468, 2021.
- [21] Ben Eysenbach, Russ R Salakhutdinov, and Sergey Levine. Robust predictable control. *Advances in Neural Information Processing Systems*, 34:27813–27825, 2021.
- [22] Fumitake Fujii, Akinori Kaneishi, Takafumi Nii, Ryu’ichiro Maenishi, and Soma Tanaka. Self-tuning two degree-of-freedom proportional–integral control system based on reinforcement learning for a multiple-input multiple-output industrial process that suffers from spatial input coupling. *Processes*, 9(3):487, 2021.
- [23] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pages 1587–1596. PMLR, 2018.
- [24] BA Grierson, MA Van Zeeland, JT Scoville, B Crowley, I Bykov, JM Park, WW Heidbrink, A Nagy, SR Haskey, and D Liu. Testing the diii-d co/counter off-axis neutral beam injected power and ability to balance injected torque. *Nuclear Fusion*, 61(11):116049, 2021.
- [25] Zhe Guan and Toru Yamamoto. Design of a reinforcement learning pid controller. *IEEJ Transactions on Electrical and Electronic Engineering*, 16(10):1354–1360, 2021.
- [26] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- [27] Sang-hee Hahn, YJ Kim, BG Penafior, JG Bak, H Han, JS Hong, YM Jeon, JH Jeong, M Joung, JW Juhn, et al. Progress and plan of kstar plasma control system upgrade. *Fusion Engineering and Design*, 112:687–691, 2016.
- [28] Dongqi Han, Kenji Doya, and Jun Tani. Variational recurrent models for solving partially observable control tasks. *arXiv preprint arXiv:1912.10703*, 2019.
- [29] Nicolas Heess, Jonathan J Hunt, Timothy P Lillicrap, and David Silver. Memory-based control with recurrent neural networks. *arXiv preprint arXiv:1512.04455*, 2015.
- [30] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. Stable baselines. <https://github.com/hill-a/stable-baselines>, 2018.
- [31] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [32] Maximilian Igl, Luisa Zintgraf, Tuan Anh Le, Frank Wood, and Shimon Whiteson. Deep variational reinforcement learning for pomdps. In *International Conference on Machine Learning*, pages 2117–2126. PMLR, 2018.
- [33] Maximilian Igl, Kamil Ciosek, Yingzhen Li, Sebastian Tschiatschek, Cheng Zhang, Sam Devlin, and Katja Hofmann. Generalization in reinforcement learning with selective noise injection and information bottleneck. *Advances in neural information processing systems*, 32, 2019.
- [34] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015.

- [35] Hozefa Jesawada, Amol Yerudkar, Carmen Del Vecchio, and Navdeep Singh. A model-based reinforcement learning approach for pid design. *arXiv preprint arXiv:2206.03567*, 2022.
- [36] Yuankun Jiang, Chenglin Li, Wenrui Dai, Junni Zou, and Hongkai Xiong. Monotonic robust policy optimization with model discrepancy. In *International Conference on Machine Learning*, pages 4951–4960. PMLR, 2021.
- [37] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- [38] Andrej Karpathy. nanogpt, 2022–. URL <https://github.com/karpathy/nanoGPT>.
- [39] Nathan P Lawrence, Gregory E Stewart, Philip D Loewen, Michael G Forbes, Johan U Backstrom, and R Bhushan Gopaluni. Reinforcement learning based design of linear fixed structure controllers. *IFAC-PapersOnLine*, 53(2):230–235, 2020.
- [40] Wenzhe Li, Hao Luo, Zichuan Lin, Chongjie Zhang, Zongqing Lu, and Deheng Ye. A survey on transformers in reinforcement learning. *arXiv preprint arXiv:2301.03044*, 2023.
- [41] Eric Liang, Richard Liaw, Philipp Moritz, Robert Nishihara, Roy Fox, Ken Goldberg, Joseph E Gonzalez, Michael I Jordan, and Ion Stoica. Rllib: Abstractions for distributed reinforcement learning. arxiv e-prints, page. *arXiv preprint arXiv:1712.09381*, 2017.
- [42] Xingyu Lu, Kimin Lee, Pieter Abbeel, and Stas Tiomkin. Dynamics generalization via information bottleneck in deep reinforcement learning. *arXiv preprint arXiv:2008.00614*, 2020.
- [43] La Marzocco. A brief history of the pid, October 2015. URL <https://home.lamarzoccousa.com/history-of-the-pid/>.
- [44] Viraj Mehta, Biswajit Paria, Jeff Schneider, Stefano Ermon, and Willie Neiswanger. An experimental design perspective on model-based reinforcement learning. *arXiv preprint arXiv:2112.05244*, 2021.
- [45] Viraj Mehta, Ian Char, Joseph Abbate, Rory Conlin, Mark D Boyer, Stefano Ermon, Jeff Schneider, and Willie Neiswanger. Exploration via planning for information about the optimal trajectory. *arXiv preprint arXiv:2210.04642*, 2022.
- [46] Luckeciano C Melo. Transformers are meta-reinforcement learners. In *International Conference on Machine Learning*, pages 15340–15359. PMLR, 2022.
- [47] Lingheng Meng, Rob Gorbet, and Dana Kulić. Memory-based deep reinforcement learning for pomdps. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5619–5626. IEEE, 2021.
- [48] Nicolas Minorsky. Directional stability of automatically steered bodies. *Journal of the American Society for Naval Engineers*, 34(2):280–309, 1922.
- [49] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.
- [50] Tianwei Ni, Benjamin Eysenbach, and Ruslan Salakhutdinov. Recurrent model-free RL can be a strong baseline for many POMDPs. In *International Conference on Machine Learning*, pages 16691–16723. PMLR, 2022.
- [51] Fernando Nogueira. Bayesian Optimization: Open source constrained global optimization tool for Python, 2014–. URL <https://github.com/fmfn/BayesianOptimization>.
- [52] Junier B Oliva, Barnabás Póczos, and Jeff Schneider. The statistical recurrent unit. In *International Conference on Machine Learning*, pages 2671–2680. PMLR, 2017.
- [53] OpenAI. Gpt-4 technical report. *arXiv*, 2023.

- [54] Charles Packer, Katelyn Gao, Jernej Kos, Philipp Krähenbühl, Vladlen Koltun, and Dawn Song. Assessing generalization in deep reinforcement learning. *arXiv preprint arXiv:1810.12282*, 2018.
- [55] Emilio Parisotto, Francis Song, Jack Rae, Razvan Pascanu, Caglar Gulcehre, Siddhant Jayakumar, Max Jaderberg, Raphael Lopez Kaufman, Aidan Clark, Seb Noury, et al. Stabilizing transformers for reinforcement learning. In *International conference on machine learning*, pages 7487–7498. PMLR, 2020.
- [56] PA Politzer, CC Petty, RJ Jayakumar, TC Luce, MR Wade, JC DeBoo, JR Ferron, P Gohil, CT Holcomb, AW Hyatt, et al. Influence of toroidal rotation on transport and stability in hybrid scenario plasmas in diii-d. *Nuclear Fusion*, 48(7):075001, 2008.
- [57] Vitchyr Pong and Ashvin Nair. rlkit. <https://github.com/rail-berkeley/rlkit>, 2018–.
- [58] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [59] Aravind Rajeswaran, Sarvjeet Ghotra, Balaraman Ravindran, and Sergey Levine. Epopot: Learning robust neural network policies using model ensembles. *arXiv preprint arXiv:1610.01283*, 2016.
- [60] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, pages 5331–5340. PMLR, 2019.
- [61] H Reimerdes, AM Garofalo, GL Jackson, M Okabayashi, EJ Strait, Ming-Sheng Chu, Y In, RJ La Haye, MJ Lanctot, YQ Liu, et al. Reduced critical rotation for resistive-wall mode stabilization in a near-axisymmetric configuration. *Physical review letters*, 98(5):055001, 2007.
- [62] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [63] JT Scoville, DA Humphreys, JR Ferron, and P Gohil. Simultaneous feedback control of plasma rotation and stored energy on the diii-d tokamak. *Fusion engineering and design*, 82(5-14): 1045–1050, 2007.
- [64] J Seo, Y-S Na, B Kim, CY Lee, MS Park, SJ Park, and YH Lee. Development of an operation trajectory design algorithm for control of multiple 0d parameters using deep reinforcement learning in kstar. *Nuclear Fusion*, 62(8):086049, 2022.
- [65] Jaemin Seo, Y-S Na, B Kim, CY Lee, MS Park, SJ Park, and YH Lee. Feedforward beta control in the kstar tokamak by deep reinforcement learning. *Nuclear Fusion*, 61(10):106010, 2021.
- [66] Open Ended Learning Team, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob Bauer, Jakub Sygnowski, Maja Trebacz, Max Jaderberg, Michael Mathieu, et al. Open-ended learning leads to generally capable agents. *arXiv preprint arXiv:2107.12808*, 2021.
- [67] B Tobias, M Chen, IGJ Classen, CW Domier, R Fitzpatrick, BA Grierson, NC Luhmann Jr, CM Muscatello, M Okabayashi, KEJ Olofsson, et al. Rotation profile flattening and toroidal flow shear reversal due to the coupling of magnetic islands in tokamaks. *Physics of Plasmas*, 23(5):056107, 2016.
- [68] ITER Physics Expert Group on Confinement Transport, , ITER Physics Expert Group on Confinement Modelling Database, , and ITER Physics Basis Editors. Chapter 2: Plasma confinement and transport. *Nuclear Fusion*, 39(12):2175–2249, 1999.
- [69] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [70] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

- [71] T Wakatsuki, T Suzuki, N Oyama, and N Hayashi. Ion temperature gradient control using reinforcement learning technique. *Nuclear Fusion*, 61(4):046036, 2021.
- [72] ML Walker, DA Humphreys, JA Leuer, JR Ferron, and BG Penaflo. Implementation of model-based multivariable control on diii-d. *GA-A23468*, 2000.
- [73] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.
- [74] Junhang Wang, Jianing Zhu, Chao Zou, Linlin Ou, and Xinyi Yu. Robust adaptive pid control based on reinforcement learning for mimo nonlinear six-joint manipulator. In *2021 China Automation Congress (CAC)*, pages 2633–2638. IEEE, 2021.
- [75] Zhihan Yang and Hai Huu Nguyen. Recurrent off-policy baselines for memory-based continuous control. In *Deep RL Workshop NeurIPS 2021*, 2021.
- [76] Wenhao Yu, Jie Tan, C Karen Liu, and Greg Turk. Preparing for the unknown: Learning a universal policy with online system identification. *arXiv preprint arXiv:1702.02453*, 2017.
- [77] Luisa Zintgraf, Kyriacos Shiarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. *arXiv preprint arXiv:1910.08348*, 2019.

Appendices

Contents

1	Introduction	1
2	Preliminaries	2
3	Methodology	3
4	Related Work	5
5	Experiments	6
5.1	Tracking Problems	6
5.2	PyBullet Locomotion Tasks	9
5.3	GPIDE Ablations	9
6	Discussion	10
7	Acknowledgements	10
	Appendices	16
A	Additional GPIDE Details	17
A.1	Forming the P, I, and D Features with GPIDE	17
B	Implementation Details	17
B.1	GPIDE Implementation Details	18
B.2	Recurrent and Transformer Baseline Details	18
B.3	PID Baseline	18
C	Hyperparameters	19
C.1	Hyperparameters for Tracking Tasks	19
C.2	Hyperparameters for PyBullet Task	19
C.3	Hyperparameters for Ablations	19
D	Computation Details	20
E	Environment Descriptions	20
E.1	Mass Spring Damper	20
E.2	Navigation Environment	21
E.3	Tokamak Control Environment	21
F	Additional Experiments	23
F.1	Experiments Using VIB + GRU	23

F.2	Lookback Size Ablations	24
G	Further Results	24
G.1	Evaluation Procedure	25
G.2	MSD and DMSD Results	25
G.3	Navigation Results	28
G.4	Tokamak Control Results	30
G.5	PyBullet Results	33
G.6	Attention Scheme Visualizations	36

A Additional GPIDE Details

A.1 Forming the P, I, and D Features with GPIDE

As mentioned in Section 3, the P, I, and D features associated with a PID controller can be easily formed with GPIDE, and in this Appendix we show this concretely. Assume that the observations take the form $o_t = (x_t^{(1)}, \dots, x_t^{(M)}, \sigma_t^{(1)}, \dots, \sigma_t^{(M)})$, and assume that the actions and rewards are not given to GPIDE for simplicity. Then the linear projection for each head, f_θ^h , takes in $x_{t-1}^{(1)}, \dots, x_{t-1}^{(M)}, \sigma_{t-1}^{(1)}, \dots, \sigma_{t-1}^{(M)}, (x_t^{(1)} - x_{t-1}^{(1)}), \dots, (x_t^{(M)} - x_{t-1}^{(M)}), (\sigma_t^{(1)} - \sigma_{t-1}^{(1)}), \dots, (\sigma_t^{(M)} - \sigma_{t-1}^{(M)})$ at each time step t . We use a GPIDE architecture with three heads and where g_θ is the identity function. The linear projections for each head is as follows:

$$f_\theta^1(o_{t-1}, o_t - o_{t-1}) = f_\theta^2(o_{t-1}, o_t - o_{t-1}) = \begin{bmatrix} (x_t^{(1)} - x_{t-1}^{(1)}) + x_{t-1}^{(1)} - (\sigma_t^{(1)} - \sigma_{t-1}^{(1)}) + \sigma_{t-1}^{(1)} \\ \vdots \\ (x_t^{(M)} - x_{t-1}^{(M)}) + x_{t-1}^{(M)} - (\sigma_t^{(M)} - \sigma_{t-1}^{(M)}) + \sigma_{t-1}^{(M)} \end{bmatrix}$$

$$f_\theta^3(o_{t-1}, o_t - o_{t-1}) = \begin{bmatrix} (x_t^{(1)} - x_{t-1}^{(1)}) - (\sigma_t^{(1)} - \sigma_{t-1}^{(1)}) \\ \vdots \\ (x_t^{(M)} - x_{t-1}^{(M)}) - (\sigma_t^{(M)} - \sigma_{t-1}^{(M)}) \end{bmatrix}$$

Note that f_θ^1 and f_θ^2 form the current error at time t (i.e. the P term), and f_θ^3 forms the change in error (i.e. the D term). For accumulation strategies, using exponential smoothing with $\alpha = 1$ for the first and third heads and a summation head for the second head will recover the P, I, and D terms for heads 1, 2, and 3, respectively. Note that the above assumes $dt = 1$, but the linear projections can be adjusted to take different dt values into account.

B Implementation Details

Code Release All code for implementations are provided in the supplemental material along with instructions for how to run experiments. The code can also be found at <https://github.com/IanChar/GPIDE>. The only experiment that cannot be run are the “real” cases for tokamak control.

Architecture We use the same general architecture for each of the RL methods in this paper (see Figure 5). Each input to the history encoders, policy functions, and Q -value functions have corresponding encoders. This setup closely follows what was done in Ni et al. [50]. The encoders are simply linear projections; however, in the case of our GRU history encoder we do linear projections followed by a ReLU activation (as done in Ni et al. [50]). Although hypothetically the policy only needs to take in history encoding, z_t , since int includes the current observation, we found it essential for the current observation to be passed in independently and have its own encoder.

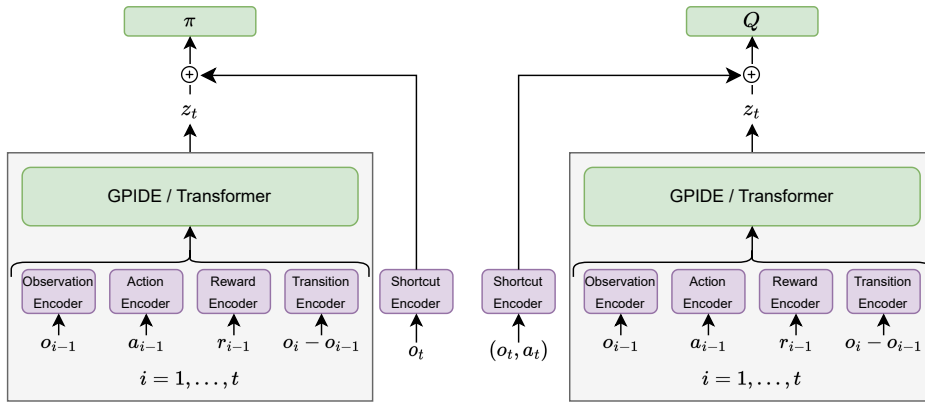


Figure 5: **General Policy and Q Function Architectures.** This architecture is heavily inspired by Ni et al. [50]. The gray box shows the history encoder modules, and this is the only thing that changes between baseline methods in the tracking problems. Note that there are two encoders: one for the policy function and one for the Q value function. The purple boxes show the input encoders, and hyperparameters for these can be found in Table 6. We found the shortcut encoders to be essential to good performance. The architecture when using GRU is nearly identical; however, there is no “Transition Encoder” since Ni et al. [50] encodes (o_i, a_{i-1}, r_{i-1}) for each time step instead.

B.1 GPIDE Implementation Details

In addition to what is mentioned in Section 3, we found that there were several choices that helped with training. First, there may be some scaling issues because $o_t - o_{t-1}$ may be small or the result of summation type heads may result in large encodings. To account for this, we use batch normalization layers [34] before each input encoding and after each ℓ^h .

There are very few nonlinear components of GPIDE. The only one that remains constant across all experiments is that a tanh activation is used for the final output of the encoder. For tracking tasks, the decoder g_θ has 1 hidden layer with 64 units and uses a ReLU activation function. For PyBullet tasks, g_θ is a linear function.

B.2 Recurrent and Transformer Baseline Details

Recurrent Encoder. For the recurrent encoder, we tried to match as many details as Ni et al. [50] as possible. We double checked our implementation against theirs and confirmed that it achieves similar performance.

Transformer Encoder. We follow the GPT2 architecture [58] for inspiration, and particularly the code provided in Karpathy [38]. In particular, we use a number of multi-headed self-attention blocks in sequence with residual connections. We use layer normalization [6] before multi-headed attention and out projections; however, we do not use dropout. The out projection for each multi-headed self-attention block has one hidden layer with four times the number of units as the embedding dimension. Although Melo [46] suggests using T-Fixup weight initialization, we found that more reliably high performance was achieved with the weight initialization of Radford et al. [58]. Lastly, we used the same representation for the history as GPIDE, i.e. $(o_{t-1}, a_{t-1}, r_{t-1}, o_t - o_{t-1})$, since it results in better performance.

B.3 PID Baseline

To tune our PID baseline, we used Bayesian Optimization over the three (for SISO) or six (for MIMO) dimensional space. Specifically we use the library provided by Nogueira [51]. The output of the blackbox optimization is the average over 100 different settings (independent from the 100 settings used for testing). We allow the optimization procedure to collect as many samples as the RL methods. The final performance reported uses the PID controller with the best gains found during the

optimization procedure. The bounds for each of the tracking tasks were eyeballed to be appropriate, which potentially preferably skews performance.

C Hyperparameters

Because of resource restrictions, we were unable to do full hyperparameter tuning for each benchmark presented in this paper. Instead, we focused on ensuring that all history encoding methods were roughly comparable, e.g. dimension of encoding, number of parameters, etc. Tables 5 show selected hyperparameters, and the following subsections describe how an important subset of these hyperparameters were picked. Any tuning that was done was over three seeds using 100 fixed settings (different from the 100 settings used for testing).

Task Type	Learning Rate	Batch Size	Discount Factor	Policy Network	Q Network	Path Length Encoding
Tracking	$3e^{-4}$	32 (256 for PIDE)	0.95	[24]	[256, 256]	100
PyBullet	$3e^{-4}$	32 (256 for PIDE)	0.99	[256, 256]	[256, 256]	64

Table 5: **SAC Hyperparameters.** The “Path Length Encoding” is the amount of history each encoder gets to observe besides PIDE which, because of the nature of it, uses the entire episode.

	Observation	Action	Reward	Transition	Policy Shortcut	Q Shortcut	History Encoding
GPIDE (Tracking)	8	N/A	N/A	8	8	64	64
GRU (Tracking)	8	N/A	N/A	N/A	8	64	64
Transformer (Tracking)	16	N/A	N/A	16	8	64	64
GPIDE (PyBullet)	32	16	16	64	8	64	128
Transformer (PyBullet)	48	16	16	48	8	64	128

Table 6: **Dimension for the Input Encoders and Final History Encoding.** The input encoders correspond to the output dimensions of the purple boxes in Figure 5. By “History Encoding” size we mean the dimension of z_t .

Task Type	D	g_θ	Hidden Size
Tracking	16		[64]
PyBullet	32		[]

Table 7: **GPIDE Specific Hyperparameters.** Recall that D corresponds to the output dimension of f_θ . Empty brackets for the hidden size means that g_θ is a linear function.

C.1 Hyperparameters for Tracking Tasks

For tracking tasks, we tried using a history encoding size of 32 and 64 for GRU, and we found that performance was better with 64. This is surprising since PIDE can perform well in these environments even though its history encoding is much smaller (3 or 6 dimensional). To make it a fair comparison, we set the history encoding dimension for GPIDE and transformer to be 64 as well. We use one layer for GRU. For the transformer-specific hyperparameters we choose half of what appears in the PyBullet tasks.

C.2 Hyperparameters for PyBullet Task

For the PyBullet tasks, we simply tried to emulate most of the hyperparameters found in Ni et al. [50]. For the transformer, we choose to use similar hyperparameters to those found in Melo [46]. However, we found that, unlike the tracking tasks, positional encoding hurts performance. As such, we do not include it for PyBullet experiments.

C.3 Hyperparameters for Ablations

For the ablations of GPIDE, we use $\alpha = 0.01, 0.1, 0.25, 0.5, 0.9, 1.0$ for the smoothing parameters when only exponential smoothing is used. When using exponential smoothing and summation, the $\alpha = 0.01$ head is replaced with a summation head. The attention version of GPIDE replaces all six of these heads with attention.

Task Type	Number of Layers	Number of Heads	Embedding Size per Head
Tracking	2	4	8
PyBullet	4	8	16

Table 8: **Transformer Specific Hyperparameters**

Encoder	SISO Tracking	MIMO Tracking (2D)	PyBullet
Transformer	25,542	25,644	793,868-795,026
GRU	14,240	14,264	74,816-75,248
GPIDE	13,228	13,288	75,296-76,486
GPIDE-ES	12,204	12,264	50,720-51,910
GPIDE-ESS	12,204	12,264	50,720-51,910
GPIDE-Attention	15,276	15,336	99,872-101,062

Table 9: **Number of Parameters in History Encoder Modules.** The number of parameters corresponds to the gray boxes in Figure 5. The difference in SISO vs MIMO and the PyBullet tasks is due to the different observation and action space dimensionalities.

D Computation Details

We used an internal cluster of machines to run these experiments. We mostly leveraged Nvidia Titan X GPUs for this, but also used a few Nvidia GTX 1080s. It is difficult to get an accurate estimate of run time since job loads vary drastically on our cluster from other users. However, to train a single policy on DMSD to completion (1 million transitions collected, or 1,000 epochs) using PIDE takes roughly 4.5 hours, using GPIDE takes roughly 17.25 hours, using a GRU takes roughly 14.5 hours, and using a transformer takes roughly 21 hours. This is similar for other tracking tasks. For PyBullet tasks, using GPIDE took roughly 43.2 hours and using a transformer took roughly 64.2 hours. We note that our implementation of GPIDE is somewhat naive and could be vastly improved. In particular, for exponential smoothing and summation heads, w_t can be cached to save on compute, which is not being done currently. This is a big advantage in efficiency that GPIDE (especially one without attention heads) has over transformers.

E Environment Descriptions

E.1 Mass Spring Damper




Figure 6: **Diagram of the Mass Spring Damper Environments.** The diagram on the left the Mass Spring Damper (MSD) environment, and the diagram on the right shows the Double Mass Spring Damper (DMSD) environment. In the diagram, we have labelled the **system parameters** and the **parts of the observation**. The dotted line shows where the center of the mass is located with no force applied, and the current position of the mass is measured with respect to this point.

For both MSD and DMSD, the observations include the current mass position(s), the target reference position(s), and the last action played. Each episode lasts for 100 time steps. For all RL methods, the action is a difference in force applied to the mass, but for the PID the action is simply the force

System Parameter	Fixed	Small	Large
Damping Constant	$\mathcal{U}(4.0, 4.0)$	$\mathcal{U}(3.5, 5.5)$	$\mathcal{U}(2.0, 10.0)$
Spring Constant	$\mathcal{U}(2.0, 2.0)$	$\mathcal{U}(1.75, 3.0)$	$\mathcal{U}(0.5, 6.0)$
Mass	$\mathcal{U}(20.0, 20.0)$	$\mathcal{U}(17.5, 40.0)$	$\mathcal{U}(10.0, 100.0)$

Table 10: **MSD and DMSD System Parameter Distributions.** Each episode system parameters are uniformly at random drawn from these bounds.

System Parameter	No Friction	Friction
Total Friction	$\mathcal{U}(0.0, 0.0)$	$\mathcal{U}(0.05, 0.25)$
Static Friction (Proportion)	$\mathcal{U}(0.0, 0.0)$	$\mathcal{U}(0.25, 0.75)$
Mass	$\mathcal{U}(15.0, 25.0)$	$\mathcal{U}(5.0, 35.0)$

Table 11: **Navigation System Parameter Distributions.** Each episode system parameters are uniformly at random drawn from these bounds. The static friction parameter drawn is the proportion of the total friction that is static friction.

to be applied to the mass at that time. The force is bounded between -10 and 10 N for MSD and -30 and 30 N for DMSD. Each episode, system parameters are drawn from a uniform distribution with bounds shown in Table 10 (they are the same for both MSD and DMSD). Targets are drawn to uniformly at random to be -1.5 to 1.5 m offset from the masses' resting positions.

E.2 Navigation Environment

Like the MSD and DMSD environments, the navigation experiment lasts 100 time steps each episode. Additionally, the observation includes position signal, target locations, and the last action. For all methods we set the action to be the change in force, and the total amount of force is bounded between -10 and 10 N . The penalty on the reward is equal to 0.01 times the magnitude of the change in force. In addition, the maximum magnitude of the velocity for the agent is bounded by $1.0m/s$. The agent always starts at the location $(0, 0)$, and the target is picked uniformly at random to be within a box of length 10 centered around the origin.

Every episode, the mass, kinetic friction coefficient, and static friction coefficient is sampled, The friction is sampled by first sampling the total amount of friction in the system, and then sampling what proportion of the total friction is static friction. All distributions for the system parameters are uniform, and we show the bounds in Table 11.

E.3 Tokamak Control Environment

Simulator Our simulator version of the tokamak control is inspired by equations used by Boyer et al. [9], Scoville et al. [63]. In particular, we use the following relations for stored energy, E , and rotation, v_{rot} :

$$\begin{aligned}\dot{E} &= P - \frac{E}{\tau_E} \\ \dot{v}_{\text{rot}} &= C_{\text{rot}}T - \frac{v_{\text{rot}}}{\tau_m}\end{aligned}$$

where P is the total power, T is the total torque, τ_E is the energy confinement time, τ_m is the momentum confinement time, and C_{rot} is a quantity relying on the ion density and major radius of the plasma. We treat τ_m and C_{rot} as constants with values of 0.1 and 80.0 respectively.

We base the energy confinement time off of the ITERH-98 scaling [68]. This uses many measurements of the plasma, but we focus on a subset of these and treat the rest as constants. In particular,

$$\tau_E = C_E I^{0.95} B^{0.15} P^{-0.69}$$

where C_E is a constant value we set to be 200, I is the plasma current, and B is the toroidal magnetic field. To relate the stored energy to β_N we use the rough approximation

$$\beta_N = C_\beta \left(\frac{aB}{I} \right) E$$

Minor Radius (m)	Plasma Current (MA)	Toroidal Magnetic Field (T)
$\mathcal{N}(0.589, 0.02)$	$\mathcal{N}(1e6, 1e5)$	$\mathcal{N}(2.75, 0.1)$

Table 12: Tokamak Control Simulator Distributions.

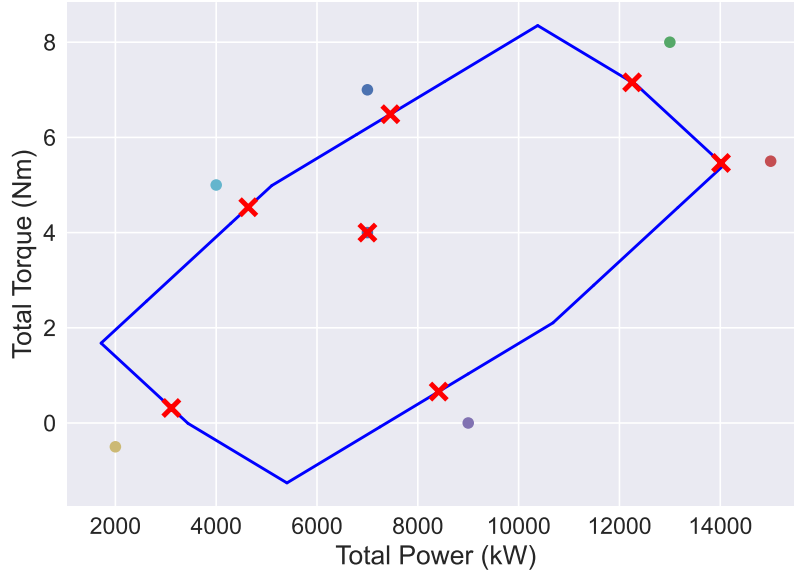


Figure 7: **Power and Torque Bounds.** The region outlined in blue shows the possible power-torque configurations. The dots show possible requests, and the corresponding red X marks show the actual achieved power-torque setting.

where C_β is a constant we set to be 5, and a is the minor radius of the plasma. For a , I , and B , we sample these from the distribution described in Table 12 for each episode. Lastly, we add momentum to the stored energy. That is, the stored energy derivative at time t , \dot{E}_t , is

$$\dot{E}_t = 0.5 \left(P_t - \frac{E_t}{\tau_E} \right) + 0.5 \dot{E}_{t-1}$$

The actions for all control methods is the amount of change for the power and torque. Because the total amount of power and torque injected rely on the beams, they are not totally disentangled. In Figure 7, we show the bounds for the action space. Furthermore, we bound the amount that power and torque can be changed by roughly $40MW/s$ and $35Nm/s$, respectively. Each step is 0.025 seconds.

Each episode lasts for 100 increments of 0.025 seconds. The observations are the current β_N and rotation values, their reference values, and the current power and torque settings. We make the initial β_N and rotation relatively small in order to simulate the plasma ramping up. We let the β_N and rotation targets be distributed as $\mathcal{U}(1.75, 2.75)$ and $\mathcal{U}(25.0, 50.0)$ rad/s , respectively.

“Real” For the real versions of the tokamak control experiments, most of the previous (such as action bounds and target distributions) stays the same. This data-driven simulator is based on the one from Char et al. [14], and we refer the reader there for more details. That being said, there are some differences between the architecture presented there and the one used in this work. Our network is a recurrent network that uses a GRU, has four hidden layers with 512 units each, and outputs the mean and log variance of a normal distribution describing how β_N and rotation will change. In addition to power and torque, it takes in measurements for the plasma current, the toroidal magnetic field, nrms (a measurement related to the plasma’ stability), and 13 other actuator requests for gas control and plasma shaping. In addition to sampling from the normal distribution outputted by the network, we train an ensemble of ten networks, and an ensemble member is selected every episode. We use five of

these models during training and the other five during testing. Along with an ensemble member being sampled each episode, we also sample a historical run, which determines the starting conditions of the plasma and how the other inputs to the neural network which are not modelled evolve over time. Recall that 100 fixed settings are used to evaluate the policy every epoch of training. In this case, a setting consists of targets, an ensemble member, and a historical run.

F Additional Experiments

F.1 Experiments Using VIB + GRU

As shown in this work, using a GRU for a history encoder often results in a policy that is ill-equipped to handle changes in the dynamics not seen at train time. One may wonder whether using other robust RL techniques is able to mask this inadequacy of GRU. To test this, we look at adding Variational Information Bottlenecking (VIB) to our GRU baseline [4]. Previous works applying this concept to RL usually do not consider the same class of POMDPs as us [42, 33]; however, Eysenbach et al. [21] does have a baseline that uses VIB with a recurrent policy.

To use VIB with RL, we alter the policy network so that it encodes input to a latent random variable, and the decodes into an action. Following the notation of Lu et al. [42], let this latent random variable be Z and the random variable representing the input of the network be S . The goal is to learn a policy that maximizes $J(\pi)$ subject to $I(Z, S) \leq I_C$, where $I(Z, S)$ is the mutual information between Z and S , and I_C is some given threshold. In practice, we optimize the Lagrangian. Where β is a Lagrangian multiplier, $p(Z|S)$ is the conditional density of Z outputted by the encoder, and $q(Z)$ is the prior, the penalizer is $-\beta \mathbb{E}_S [D_{\text{KL}}(p(Z|S)||q(Z))]$. Like other works, we assume that $q(Z)$ is a standard multivariate normal.

We alter our GRU baseline for tracking tasks so that the policy uses VIB. This is not entirely straightforward since our policy network is already quite small. We choose to keep as close to original policy architecture as possible and set the dimension of the latent variable, Z , to be 24. Note that this change has no effect on the history encoder; this only affects the policy network. For our experiments, we set $\beta = 0.1$, but we note that we may be able to achieve better performance through more careful tuning or annealing of β .

In any case, we do see that VIB helps with robustness in many instances (see Table 13). However, the cases where there are improvements are instances where the GRU policy already did a good job at generalizing to the test environment. These are primarily the MSD and DMSD environments where the system parameters drawn during training time are simply a subset of those drawn during testing time (interestingly, this notion of dynamics generalization matches the set up of the experiments presented in Lu et al. [42]). Surprisingly, in the navigation and tokamak control experiments, where there are more complex differences between the train and test environments, VIB can sometimes hurt the final performance.

	PID Controller	GRU	GRU+VIB	Transformer	PIDE	GPIDE
MSD Fixed / Fixed	-6.14 ± 0.02	-5.76 ± 0.02	-5.73 ± 0.01	-5.75 ± 0.01	-5.69 ± 0.00	-5.76 ± 0.01
MSD Fixed / Large	-11.39 ± 0.09	-12.52 ± 0.11	-12.50 ± 0.14	-10.87 ± 0.05	-11.44 ± 0.03	-11.61 ± 0.07
MSD Small / Small	-7.49 ± 0.03	-7.02 ± 0.01	-7.01 ± 0.01	-7.15 ± 0.02	-7.14 ± 0.01	-7.12 ± 0.04
MSD Small / Large	-11.18 ± 0.09	-9.82 ± 0.07	-9.57 ± 0.03	-10.01 ± 0.03	-10.88 ± 0.04	-10.43 ± 0.14
DMSD Fixed / Fixed	-15.33 ± 0.14	-16.20 ± 0.31	-15.83 ± 0.28	-15.41 ± 0.13	-12.64 ± 0.04	-13.49 ± 0.22
DMSD Fixed / Large	-27.59 ± 0.44	-37.21 ± 0.35	-35.34 ± 0.28	-28.16 ± 0.17	-25.29 ± 0.18	-27.54 ± 0.33
DMSD Small / Small	-21.78 ± 0.14	-22.49 ± 0.34	-22.51 ± 0.24	-20.56 ± 0.16	-18.09 ± 0.04	-18.67 ± 0.17
DMSD Small / Large	-26.57 ± 0.22	-31.27 ± 0.36	-30.93 ± 0.34	-26.04 ± 0.24	-23.82 ± 0.13	-23.65 ± 0.20
Nav Sim / Sim	-17.23 ± 0.18	-13.82 ± 0.01	-14.69 ± 0.02	-13.68 ± 0.01	-13.74 ± 0.00	-13.65 ± 0.00
Nav Sim / Real	-23.87 ± 0.29	-29.85 ± 0.55	-39.57 ± 0.24	-22.84 ± 0.11	-20.37 ± 0.08	-21.23 ± 0.12
β_N Sim / Sim	-8.09 ± 0.00	-7.19 ± 0.00	-7.24 ± 0.01	-7.22 ± 0.00	-8.71 ± 0.02	-7.21 ± 0.01
β_N Sim / Real	-16.41 ± 0.30	-31.21 ± 1.67	-32.19 ± 1.19	-31.49 ± 3.66	-43.78 ± 6.46	-26.83 ± 1.36
β_N -Rotation Sim / Sim	-27.56 ± 0.08	-18.53 ± 0.02	-18.61 ± 0.12	-18.79 ± 0.09	-21.36 ± 0.05	-18.45 ± 0.03
β_N -Rotation Sim / Real	-30.08 ± 0.95	-45.91 ± 2.10	-44.24 ± 1.33	-48.23 ± 0.25	-60.23 ± 3.20	-41.86 ± 0.69
Average	-18.33	-20.12	-21.14	-18.71	-19.58	-17.51

Table 13: **Tracking Experiments with GRU+VIB.** We use green and red text to highlight significant improvements and deteriorations in performance over vanilla GRU. We only highlight a subset of configurations since we are focused on the robustness properties. This table shows average (unnormalized) returns.

F.2 Lookback Size Ablations

To better understand the role of the maximum lookback size (i.e. the amount of history used to form the encoding) of GPIDE, we repeat the PyBullet experiments using a lookback size of 4, 16, 64, and 128 with and without attention (labelled GPIDE and GPIDE-ESS, respectively). Figures 8 and 9 show performance curves for GPIDE and GPIDE-ESS respectively. It is clear that there is a massive increase in improvement when expanding the maximum size of the lookback from 4 to 16. For the most part, this trend continues expanding the lookback size from 16 to 64; however, it seems pushing from 64 to 128 yields mixed results. For some tasks, such as HalfCheetah-P, expanding the lookback to 128 results in noticeable improvements both with and without attention. For other tasks, such as Hopper-V, this expansion yields slightly worse performance, possibly because of training stability issues.

One interesting observation is that, when attention is included, this decrease in performance from expanding lookback can occur when increasing from 16 to 64 (see HalfCheetah-V and Walker-V). At the same time, however, it appears the GPIDE can sometimes maintain good performance when expanding the lookback to 128 when GPIDE-ESS cannot (Walker-P).

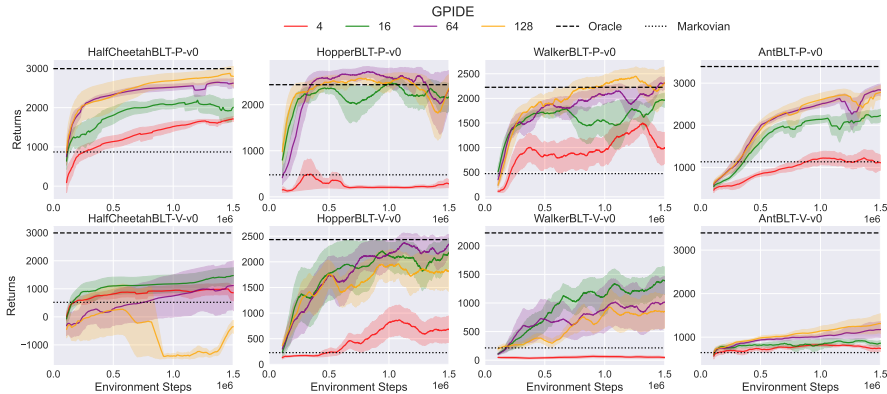


Figure 8: **GPIDE Lookback Ablation.** Each curve shows the average over four seeds and the standard error of each.

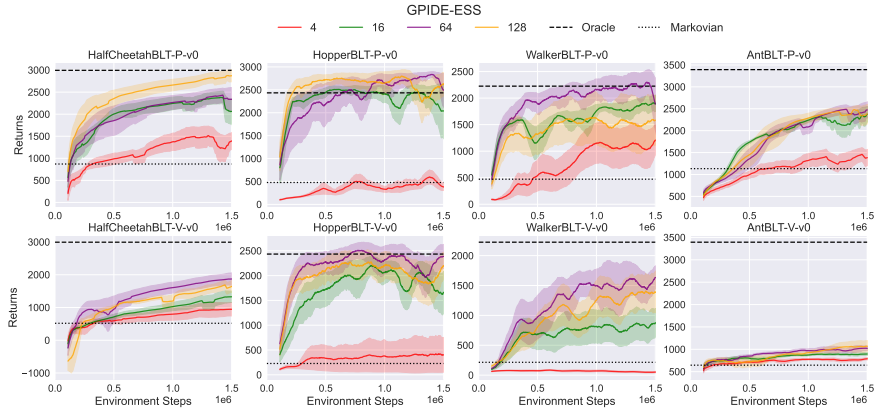


Figure 9: **GPIDE-ESS Lookback Ablation.** Each curve shows the average over four seeds and the standard error of each.

G Further Results

In this Appendix, we give further evaluation of the evaluation procedure. In addition, we give full tables of results for normalized and unnormalized scores for all methods. We also show performance

traces. Note that the percentage changes in Table 4 do not necessarily reflect tables in this section since they report all combinations of environment variants.

G.1 Evaluation Procedure

As stated in the main paper, for tracking tasks, we fix 100 settings (each comprised of targets, start state, and system parameters) that are used to evaluate the policy for every epoch of training (i.e. for every epoch the evaluation returns is the average over all 100 settings returns). We use a separate 100 settings when tuning. For the final returns, we average over the last 10% of recorded evaluations.

For the PyBullet tasks, we use ten different rollouts for evaluation following Ni et al. [50]. We also average over the last 20% of recorded evaluations like they do.

Normalized Table Scores. We now give an in-depth explanation of how the scores in the table are computed. Let $\pi_{(b,i)}$ be the policy trained with baseline method b (e.g. with GPIDE, transformer, or GRU encoder) on environment variant i (e.g. fixed, small, or large). Let $J_j(\pi_{(b,i)})$ be the evaluation of policy $\pi_{(b,i)}$ on environment variant j , i.e. the average returns over all seeds and episodes. The normalized score for policy $\pi_{(b,i)}$ on variant j is then

$$\frac{J_j(\pi_{(b,i)}) - \min_{b',i'} J_j(\pi_{(b',i')})}{\max_{b',i'} J_j(\pi_{(b',i')}) - \min_{b',i'} J_j(\pi_{(b',i')})}$$

Note that we only min and max over baseline methods presented in the table.

For PyBullet tasks, we do the same procedure but normalize by the oracle policy’s performance (sees both position and velocity and has no history encoder) and the Markovian policy’s performance (sees only position or velocity and has no history encoder). For both of these policies, we use what was reported from Ni et al. [50]. Note the our normalized scores differ slightly from those used in Ni et al. [50] since they normalize based on the best and worst returns of any policy; however, we believe our scheme gives a more intuitive picture of how any given policy is performing.

G.2 MSD and DMSD Results

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Fixed / Fixed	-6.14 ± 0.02	-5.76 ± 0.02	-5.75 ± 0.01	-5.69 ± 0.00	-5.76 ± 0.01	-5.75 ± 0.01	-5.73 ± 0.01	-5.83 ± 0.02
Fixed / Small	-7.51 ± 0.04	-7.56 ± 0.03	-7.29 ± 0.01	-7.37 ± 0.01	-7.33 ± 0.04	-7.37 ± 0.01	-7.32 ± 0.03	-7.39 ± 0.03
Fixed / Large	-11.39 ± 0.09	-12.52 ± 0.11	-10.87 ± 0.05	-11.44 ± 0.03	-11.61 ± 0.07	-11.48 ± 0.05	-12.50 ± 0.19	-11.52 ± 0.10
Small / Fixed	-6.26 ± 0.06	-5.80 ± 0.00	-5.92 ± 0.01	-5.95 ± 0.01	-5.93 ± 0.05	-5.89 ± 0.01	-5.92 ± 0.02	-5.91 ± 0.02
Small / Small	-7.49 ± 0.03	-7.02 ± 0.01	-7.15 ± 0.02	-7.14 ± 0.01	-7.12 ± 0.04	-7.09 ± 0.02	-7.15 ± 0.02	-7.12 ± 0.02
Small / Large	-11.18 ± 0.09	-9.82 ± 0.07	-10.01 ± 0.03	-10.88 ± 0.04	-10.43 ± 0.14	-10.42 ± 0.13	-10.43 ± 0.12	-10.07 ± 0.14
Large / Fixed	-6.78 ± 0.16	-6.08 ± 0.01	-6.28 ± 0.03	-6.27 ± 0.01	-6.27 ± 0.03	-6.23 ± 0.04	-6.25 ± 0.04	-6.28 ± 0.05
Large / Small	-7.78 ± 0.12	-7.25 ± 0.02	-7.44 ± 0.05	-7.43 ± 0.02	-7.45 ± 0.03	-7.44 ± 0.05	-7.44 ± 0.04	-7.48 ± 0.06
Large / Large	-11.12 ± 0.05	-9.44 ± 0.02	-9.67 ± 0.05	-10.37 ± 0.02	-9.66 ± 0.04	-9.68 ± 0.05	-9.70 ± 0.05	-9.69 ± 0.06
Average	-8.41	-7.92	-7.82	-8.06	-7.95	-7.93	-8.05	-7.92

Table 14: Unnormalized MSD Results.

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Fixed / Fixed	58.09 ± 1.66	93.18 ± 1.46	94.04 ± 0.83	100.00 ± 0.27	93.18 ± 1.20	93.77 ± 1.26	96.20 ± 1.22	87.16 ± 1.78
Fixed / Small	36.41 ± 5.36	29.74 ± 3.82	64.96 ± 1.38	54.90 ± 0.73	59.54 ± 5.78	54.35 ± 1.35	60.89 ± 3.48	51.84 ± 3.38
Fixed / Large	36.58 ± 2.86	0.00 ± 3.42	53.70 ± 1.71	34.92 ± 0.93	29.55 ± 2.32	33.62 ± 1.71	0.60 ± 6.25	32.51 ± 3.29
Small / Fixed	46.87 ± 5.88	89.05 ± 0.32	78.21 ± 1.31	75.81 ± 0.79	77.27 ± 4.66	81.41 ± 1.20	78.82 ± 1.44	79.64 ± 1.80
Small / Small	38.25 ± 3.44	100.00 ± 0.98	83.49 ± 3.07	84.88 ± 0.81	87.78 ± 5.31	90.66 ± 2.02	83.97 ± 2.40	87.57 ± 2.65
Small / Large	43.52 ± 2.82	87.63 ± 2.28	81.44 ± 0.82	53.21 ± 1.31	68.03 ± 4.43	68.09 ± 4.10	67.78 ± 3.84	79.57 ± 4.71
Large / Fixed	0.00 ± 15.12	63.36 ± 1.17	45.01 ± 3.18	46.37 ± 1.29	46.68 ± 3.06	49.86 ± 3.84	48.52 ± 3.69	45.03 ± 4.72
Large / Small	0.00 ± 15.75	70.44 ± 3.30	45.45 ± 6.93	45.73 ± 2.47	43.66 ± 4.47	44.71 ± 6.45	45.21 ± 5.42	39.64 ± 7.82
Large / Large	45.60 ± 1.71	100.00 ± 0.61	92.60 ± 1.49	69.88 ± 0.69	93.03 ± 1.27	92.36 ± 1.62	91.67 ± 1.68	91.95 ± 1.80
Average	33.92	70.38	70.99	62.86	66.53	67.65	63.74	66.10

Table 15: Normalized MSD Results.

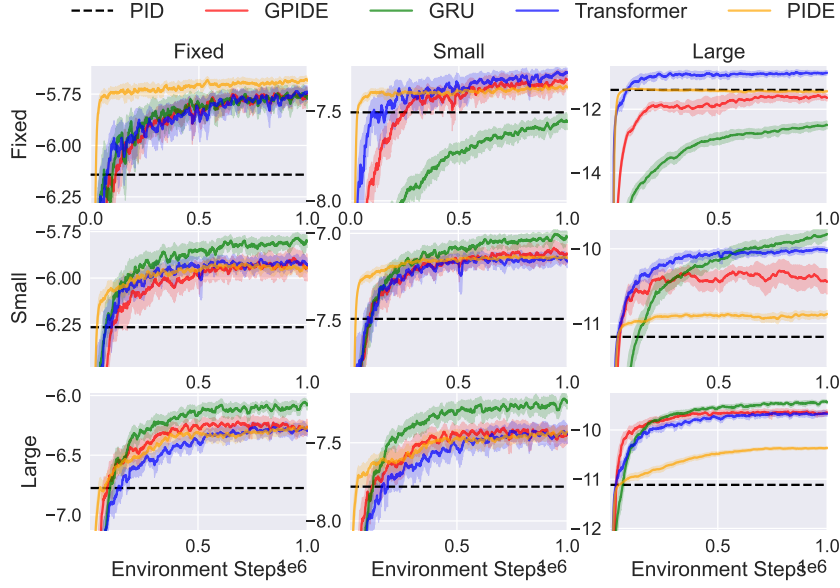


Figure 10: **MSD Performance Curves**. Each row corresponds to a training environment, and each column corresponds to a testing environment.

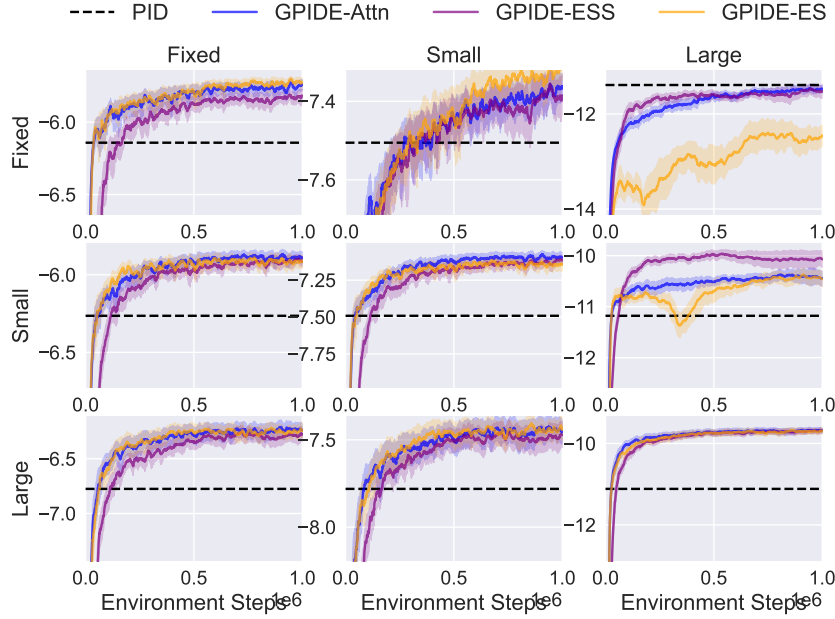


Figure 11: **MSD Performance Curve for Ablations**. Each row corresponds to a training environment, and each column corresponds to a testing environment.

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Fixed / Fixed	-15.33 ± 0.14	-16.20 ± 0.31	-15.41 ± 0.13	-12.64 ± 0.04	-13.49 ± 0.22	-13.92 ± 0.09	-13.35 ± 0.05	-16.77 ± 0.13
Fixed / Small	-21.29 ± 0.29	-25.21 ± 0.32	-21.37 ± 0.16	-18.58 ± 0.05	-19.77 ± 0.24	-21.31 ± 0.07	-20.09 ± 0.08	-23.29 ± 0.15
Fixed / Large	-27.59 ± 0.44	-37.21 ± 0.35	-28.16 ± 0.17	-25.29 ± 0.18	-27.54 ± 0.33	-31.14 ± 0.13	-28.14 ± 0.11	-31.84 ± 0.71
Small / Fixed	-18.15 ± 0.91	-17.75 ± 0.42	-15.86 ± 0.11	-13.43 ± 0.09	-14.37 ± 0.17	-14.35 ± 0.11	-13.57 ± 0.10	-16.85 ± 0.11
Small / Small	-21.78 ± 0.14	-22.49 ± 0.34	-20.56 ± 0.16	-18.09 ± 0.04	-18.67 ± 0.17	-18.93 ± 0.10	-17.97 ± 0.07	-21.77 ± 0.10
Small / Large	-26.57 ± 0.22	-31.27 ± 0.36	-26.04 ± 0.24	-23.82 ± 0.13	-23.65 ± 0.20	-23.66 ± 0.10	-22.72 ± 0.08	-28.26 ± 0.12
Large / Fixed	-21.96 ± 0.62	-22.41 ± 0.32	-18.37 ± 0.30	-14.83 ± 0.12	-15.75 ± 0.14	-16.79 ± 0.04	-15.23 ± 0.12	-18.89 ± 0.28
Large / Small	-22.30 ± 0.44	-26.63 ± 0.39	-22.00 ± 0.24	-19.46 ± 0.08	-19.99 ± 0.15	-21.14 ± 0.07	-19.71 ± 0.12	-23.19 ± 0.32
Large / Large	-25.29 ± 0.30	-29.34 ± 0.30	-24.43 ± 0.21	-24.06 ± 0.03	-22.08 ± 0.14	-23.06 ± 0.07	-21.81 ± 0.09	-25.32 ± 0.19
Average	-22.25	-25.39	-21.36	-18.91	-19.48	-20.48	-19.18	-22.91

Table 16: **Unnormalized DMSD Results**.

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Fixed / Fixed	72.45 ± 1.44	63.59 ± 3.16	71.62 ± 1.30	100.00 ± 0.39	91.35 ± 2.28	86.93 ± 0.89	92.74 ± 0.55	57.75 ± 1.31
Fixed / Small	61.66 ± 3.35	16.43 ± 3.75	60.71 ± 1.80	93.01 ± 0.60	79.26 ± 2.81	61.50 ± 0.81	75.51 ± 0.97	38.55 ± 1.77
Fixed / Large	62.47 ± 2.86	0.00 ± 2.24	58.78 ± 1.11	77.38 ± 1.14	62.76 ± 2.13	39.41 ± 0.83	58.92 ± 0.73	34.84 ± 4.61
Small / Fixed	43.59 ± 9.27	47.76 ± 4.25	67.02 ± 1.10	91.92 ± 0.90	82.32 ± 1.72	82.52 ± 1.14	90.46 ± 0.99	56.98 ± 1.16
Small / Small	56.04 ± 1.57	47.82 ± 3.96	70.07 ± 1.88	98.69 ± 0.48	91.94 ± 2.00	88.95 ± 1.18	100.00 ± 0.78	56.17 ± 1.11
Small / Large	69.11 ± 1.42	38.57 ± 2.33	72.51 ± 1.58	86.96 ± 0.82	88.08 ± 1.31	87.99 ± 0.64	94.09 ± 0.51	58.08 ± 0.80
Large / Fixed	4.64 ± 6.34	0.00 ± 3.30	41.37 ± 3.09	77.62 ± 1.24	68.16 ± 1.45	57.60 ± 0.36	73.51 ± 1.24	36.06 ± 2.85
Large / Small	50.02 ± 5.07	0.00 ± 4.56	53.45 ± 2.80	82.77 ± 0.98	76.66 ± 1.75	63.36 ± 0.85	79.93 ± 1.43	39.74 ± 3.65
Large / Large	77.38 ± 1.93	51.09 ± 1.98	82.96 ± 1.38	85.37 ± 0.18	98.23 ± 0.90	91.86 ± 0.44	100.00 ± 0.56	77.22 ± 1.21
Average	55.26	29.47	64.28	88.19	82.08	73.35	85.02	50.60

Table 17: Normalized DMSD Results.

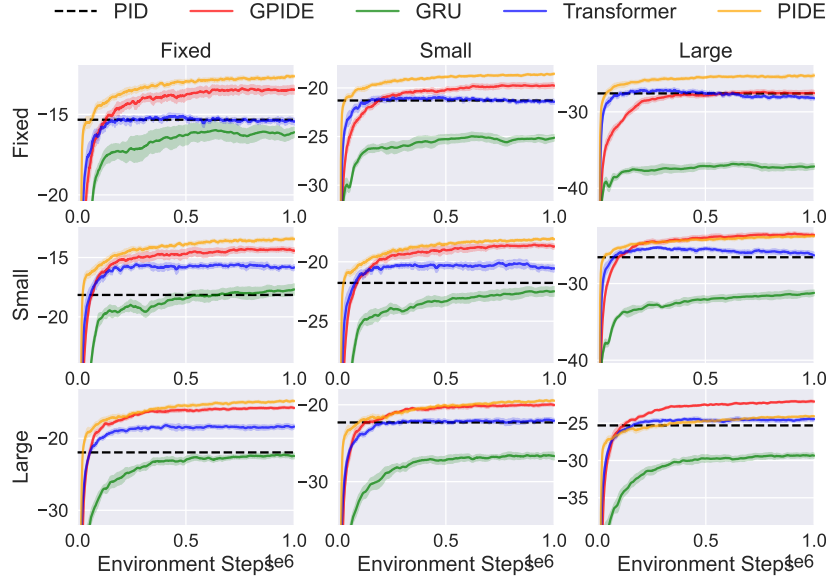


Figure 12: DMSD Performance Curves. Each row corresponds to a training environment, and each column corresponds to a testing environment.

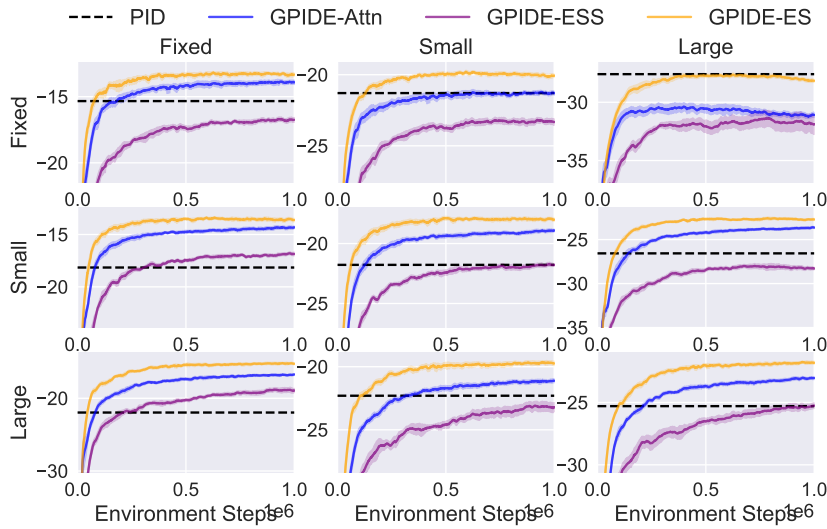


Figure 13: DMSD Performance Curve for Ablations. Each row corresponds to a training environment, and each column corresponds to a testing environment.

G.3 Navigation Results

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Sim / Sim	28.94 ± 3.63	96.76 ± 0.15	99.57 ± 0.12	98.33 ± 0.06	100.00 ± 0.07	99.64 ± 0.06	99.66 ± 0.06	99.81 ± 0.09
Sim / Real	43.12 ± 2.08	0.00 ± 3.94	50.55 ± 0.78	68.34 ± 0.57	62.16 ± 0.89	63.17 ± 0.57	59.21 ± 1.15	52.52 ± 0.50
Real / Sim	0.00 ± 4.09	57.49 ± 1.17	68.03 ± 0.40	59.54 ± 0.85	74.88 ± 0.61	72.84 ± 0.64	74.75 ± 0.68	71.13 ± 0.72
Real / Real	67.28 ± 2.05	97.29 ± 0.20	99.20 ± 0.14	95.94 ± 0.04	100.00 ± 0.21	99.19 ± 0.09	99.11 ± 0.21	99.67 ± 0.17
Average	34.83	62.89	79.34	80.54	84.26	83.71	83.18	80.78

Table 18: **Normalized Navigation Results.** Note that these results are after 1 million collected samples.

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Sim / Sim	-17.23 ± 0.18	-13.82 ± 0.01	-13.68 ± 0.01	-13.74 ± 0.00	-13.65 ± 0.00	-13.67 ± 0.00	-13.67 ± 0.00	-13.66 ± 0.00
Sim / Real	-23.87 ± 0.29	-29.85 ± 0.55	-22.84 ± 0.11	-20.37 ± 0.08	-21.23 ± 0.12	-21.09 ± 0.08	-21.64 ± 0.16	-22.57 ± 0.07
Real / Sim	-18.69 ± 0.21	-15.79 ± 0.06	-15.26 ± 0.02	-15.69 ± 0.04	-14.92 ± 0.03	-15.02 ± 0.03	-14.93 ± 0.03	-15.11 ± 0.04
Real / Real	-20.52 ± 0.28	-16.36 ± 0.03	-16.09 ± 0.02	-16.55 ± 0.01	-15.98 ± 0.03	-16.09 ± 0.01	-16.11 ± 0.03	-16.03 ± 0.02
Average	-20.08	-18.96	-16.97	-16.59	-16.45	-16.47	-16.59	-16.84

Table 19: **Unnormalized Navigation Results.** Note that these results are after 1 million collected samples.




Figure 14: **Navigation Performance Curves.** Each row corresponds to a training environment, and each column corresponds to a testing environment. Note that these runs are only done for one million transitions.




Figure 15: **Navigation Performance Curve for Ablations.** Each row corresponds to a training environment, and each column corresponds to a testing environment. Note that these runs are only done for one million transitions.

G.4 Tokamak Control Results

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Sim / Sim	90.95 ± 0.05	100.00 ± 0.03	99.63 ± 0.03	84.74 ± 0.16	99.75 ± 0.06	99.91 ± 0.02	99.90 ± 0.02	99.47 ± 0.04
Sim / Real	89.15 ± 0.99	40.96 ± 5.45	40.05 ± 11.91	0.00 ± 21.04	55.21 ± 4.44	61.56 ± 7.40	65.65 ± 5.66	35.66 ± 4.41
Real / Sim	50.62 ± 3.96	36.33 ± 3.61	35.26 ± 2.22	0.00 ± 3.48	48.40 ± 4.04	52.62 ± 1.38	56.30 ± 2.25	16.33 ± 5.98
Real / Real	98.45 ± 0.77	98.24 ± 0.38	98.74 ± 0.29	100.00 ± 0.23	99.30 ± 0.64	98.39 ± 0.33	98.55 ± 0.33	98.27 ± 0.37
Average	82.29	68.88	68.42	46.18	75.67	78.12	80.10	62.43

Table 20: Normalized β_N Tracking Results.

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Sim / Sim	-8.09 ± 0.00	-7.19 ± 0.00	-7.22 ± 0.00	-8.71 ± 0.02	-7.21 ± 0.01	-7.19 ± 0.00	-7.20 ± 0.00	-7.24 ± 0.00
Sim / Real	-16.41 ± 0.30	-31.21 ± 1.67	-31.49 ± 3.66	-43.78 ± 6.46	-26.83 ± 1.36	-24.88 ± 2.27	-23.63 ± 1.74	-32.83 ± 1.35
Real / Sim	-12.12 ± 0.40	-13.55 ± 0.36	-13.66 ± 0.22	-17.18 ± 0.35	-12.34 ± 0.40	-11.92 ± 0.14	-11.55 ± 0.22	-15.55 ± 0.60
Real / Real	-13.56 ± 0.23	-13.62 ± 0.12	-13.47 ± 0.09	-13.08 ± 0.07	-13.30 ± 0.20	-13.58 ± 0.10	-13.53 ± 0.10	-13.61 ± 0.11
Average	-12.55	-16.39	-16.46	-20.69	-14.92	-14.39	-13.98	-17.31

Table 21: Unnormalized β_N Tracking Results.




Figure 16: β_N Tracking Performance Curves. Each row corresponds to a training environment, and each column corresponds to a testing environment.




Figure 17: β_N Tracking Performance Curve for Ablations. Each row corresponds to a training environment, and each column corresponds to a testing environment.

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Sim / Sim	46.78 ± 0.44	99.50 ± 0.12	97.99 ± 0.50	82.96 ± 0.27	100.00 ± 0.15	99.64 ± 0.19	99.97 ± 0.12	96.18 ± 1.35
Sim / Real	83.48 ± 2.63	39.65 ± 5.83	33.22 ± 0.69	0.00 ± 8.87	50.86 ± 1.92	54.36 ± 2.07	52.56 ± 2.44	42.51 ± 2.97
Real / Sim	0.00 ± 8.79	21.31 ± 2.45	7.23 ± 3.86	22.49 ± 1.84	19.02 ± 3.88	22.70 ± 4.42	5.20 ± 20.06	15.35 ± 8.29
Real / Real	91.76 ± 0.84	98.07 ± 0.52	96.05 ± 0.31	97.94 ± 0.23	99.73 ± 0.46	97.62 ± 0.46	100.00 ± 0.28	96.33 ± 0.47
Average	55.51	64.63	58.62	50.85	67.40	68.58	64.43	62.59

Table 22: Normalized β_N -Rotation Tracking Results.

	PID Controller	GRU	Transformer	PIDE	GPIDE	GPIDE-ES	GPIDE-ESS	GPIDE-Attn
Sim / Sim	-27.56 ± 0.08	-18.53 ± 0.02	-18.79 ± 0.09	-21.36 ± 0.05	-18.45 ± 0.03	-18.51 ± 0.03	-18.45 ± 0.02	-19.10 ± 0.23
Sim / Real	-30.08 ± 0.95	-45.91 ± 2.10	-48.23 ± 0.25	-60.23 ± 3.20	-41.86 ± 0.69	-40.60 ± 0.75	-41.25 ± 0.88	-44.88 ± 1.07
Real / Sim	-35.57 ± 1.50	-31.92 ± 0.42	-34.33 ± 0.66	-31.72 ± 0.32	-32.31 ± 0.66	-31.68 ± 0.76	-34.68 ± 3.43	-32.94 ± 1.42
Real / Real	-27.09 ± 0.30	-24.81 ± 0.19	-25.54 ± 0.11	-24.86 ± 0.08	-24.21 ± 0.16	-24.98 ± 0.17	-24.12 ± 0.10	-25.44 ± 0.17
Average	-30.08	-30.29	-31.72	-34.54	-29.21	-28.94	-29.62	-30.59

Table 23: Unnormalized β_N -Rotation Tracking Results.

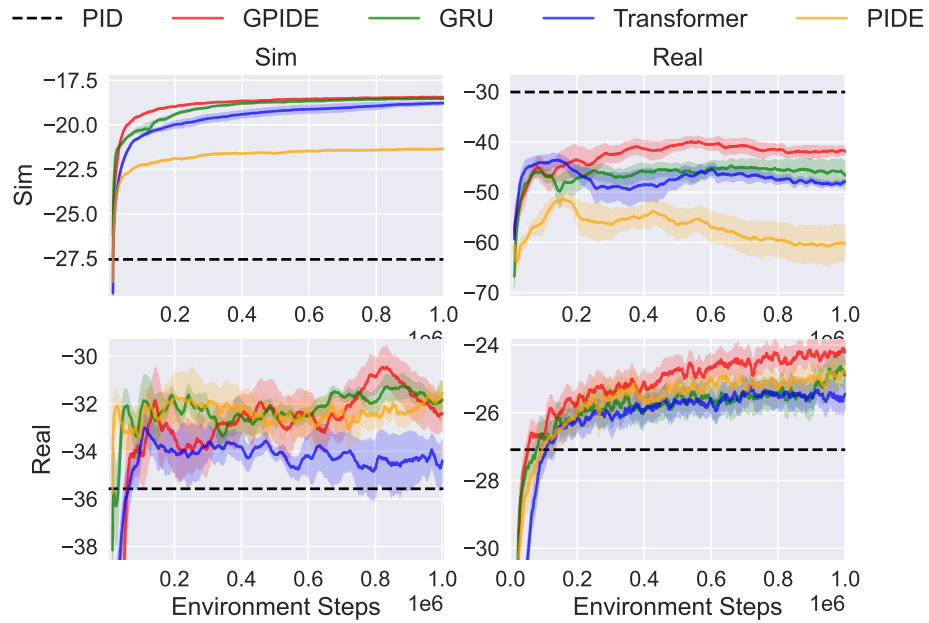


Figure 18: β_N -Rotation Tracking Performance Curves. Each row corresponds to a training environment, and each column corresponds to a testing environment.




Figure 19: β_N -Rotation Tracking Performance Curve for Ablations. Each row corresponds to a training environment, and each column corresponds to a testing environment.

G.5 PyBullet Results

For these results, SAC encodes observations, actions and rewards. TD3 encodes observations and actions since it is the best performing on average.

	PPO-GRU	A2C-GRU	SAC-LSTM	TD3-GRU	VRM	SAC-Transformer	SAC-GPDE	SAC-GPDE-ES	SAC-GPDE-ESS	SAC-GPDE-Attn
HalfCheetah-P	27.09 ± 7.85	-22.00 ± 5.13	77.06 ± 7.96	85.80 ± 5.15	-107.00 ± 1.39	37.00 ± 9.97	82.63 ± 3.46	85.45 ± 4.83	71.08 ± 6.95	77.63 ± 4.60
Hopper-P	49.00 ± 5.22	-2.29 ± 3.33	64.36 ± 9.94	84.63 ± 8.33	3.53 ± 1.63	59.54 ± 19.64	93.27 ± 13.56	111.48 ± 4.20	113.95 ± 4.67	104.87 ± 8.76
Walker-P	1.67 ± 4.39	-10.48 ± 1.73	40.92 ± 15.56	29.08 ± 9.67	-3.89 ± 1.25	24.89 ± 14.80	96.61 ± 1.60	76.58 ± 5.47	94.58 ± 11.00	71.36 ± 6.37
Ant-P	39.48 ± 3.74	-13.06 ± 6.52	60.97 ± 3.54	-36.36 ± 3.35	-36.39 ± 0.17	-10.57 ± 2.34	66.66 ± 2.94	64.73 ± 3.82	57.78 ± 3.78	63.19 ± 5.32
HalfCheetah-V	19.68 ± 11.71	-50.13 ± 9.50	18.54 ± 33.09	59.03 ± 2.88	-80.49 ± 2.97	-41.31 ± 26.15	20.39 ± 29.60	51.03 ± 13.93	53.14 ± 5.86	-54.70 ± 19.89
Hopper-V	13.86 ± 4.80	-0.60 ± 3.33	16.26 ± 12.44	57.43 ± 8.63	10.08 ± 3.51	0.28 ± 8.49	90.98 ± 4.28	72.63 ± 19.28	90.09 ± 2.50	30.73 ± 1.60
Walker-V	8.12 ± 5.43	-8.02 ± 0.57	-1.57 ± 1.88	-4.63 ± 1.30	-1.80 ± 0.70	-8.21 ± 1.31	36.90 ± 16.59	68.30 ± 4.33	67.54 ± 3.60	14.85 ± 11.26
Ant-V	1.43 ± 3.26	-13.67 ± 1.83	-16.95 ± 1.29	17.03 ± 6.55	-13.41 ± 0.12	0.81 ± 1.31	18.03 ± 5.10	4.56 ± 5.20	12.85 ± 1.67	-1.84 ± 5.76
Average	20.04	-15.03	32.45	36.50	-28.67	7.80	63.18	66.84	70.13	38.26

Table 24: Normalized PyBullet Scores.

	PPO-GRU	AZC-GRU	SAC-LSTM	TD3-GRU	VRM	SAC-Transformer	SAC-GPIDE	SAC-GPIDE-ES	SAC-GPIDE-ESS	SAC-GPIDE-Attn
HalfCheetah-P	1445.81 ± 166.79	403.35 ± 108.97	2506.88 ± 168.93	2692.53 ± 109.43	-1401.67 ± 29.62	1656.13 ± 211.75	2625.13 ± 73.49	2684.98 ± 102.57	2379.79 ± 147.67	2519.06 ± 97.72
Hopper-P	1436.43 ± 102.09	433.19 ± 65.09	1736.81 ± 194.51	2133.42 ± 162.93	546.93 ± 31.81	1642.63 ± 384.10	2302.31 ± 265.21	2658.48 ± 82.18	2706.81 ± 91.39	2529.31 ± 171.41
Walker-P	501.06 ± 76.99	288.10 ± 30.39	1189.28 ± 272.77	981.63 ± 169.46	403.60 ± 21.85	908.17 ± 259.52	2165.52 ± 28.10	1814.40 ± 95.91	2129.91 ± 192.87	1722.81 ± 115.22
Ant-P	2023.52 ± 84.58	837.57 ± 147.53	2511.54 ± 80.13	310.72 ± 75.68	310.24 ± 3.83	893.84 ± 52.83	2640.16 ± 66.46	2596.63 ± 86.26	2439.48 ± 85.37	2561.67 ± 120.21
HalfCheetah-V	1005.13 ± 289.84	-723.40 ± 235.29	977.02 ± 819.24	1979.56 ± 71.40	-1475.15 ± 73.42	-505.00 ± 647.43	1022.93 ± 732.93	1781.36 ± 344.95	1833.60 ± 145.14	-836.47 ± 492.45
Hopper-V	534.05 ± 105.85	215.22 ± 73.48	587.10 ± 274.42	1495.11 ± 190.42	450.77 ± 77.35	234.49 ± 187.36	2235.02 ± 94.45	1830.26 ± 425.16	2215.47 ± 55.16	906.05 ± 35.29
Walker-V	377.80 ± 109.11	53.25 ± 11.45	182.97 ± 37.89	121.44 ± 26.14	178.28 ± 14.09	49.32 ± 26.43	956.43 ± 333.46	1587.56 ± 87.15	1572.41 ± 72.46	513.07 ± 226.34
Ant-V	684.36 ± 89.48	269.32 ± 50.35	178.98 ± 35.57	1113.19 ± 179.93	276.33 ± 3.18	667.20 ± 35.98	1140.73 ± 140.22	770.51 ± 143.02	998.35 ± 46.04	594.54 ± 158.48
Average	1001.27	222.08	1233.82	1353.45	-88.84	693.35	1886.03	1965.52	2034.48	1313.76

Table 25: Unnormalized PyBullet Scores.

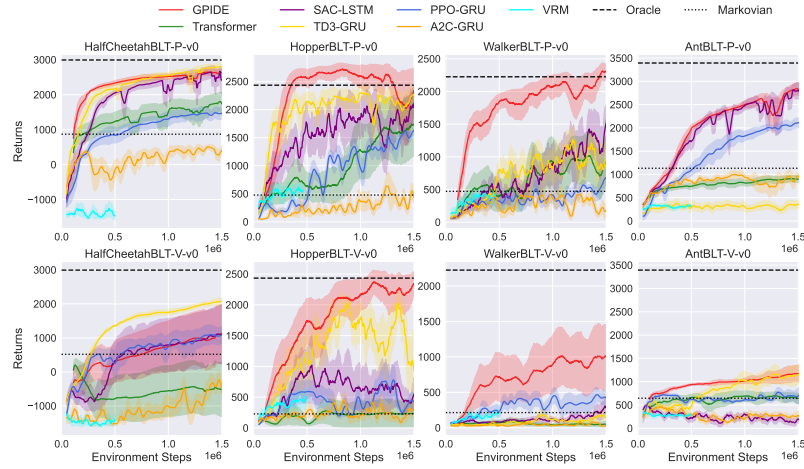


Figure 20: PyBullet Performance Curves.

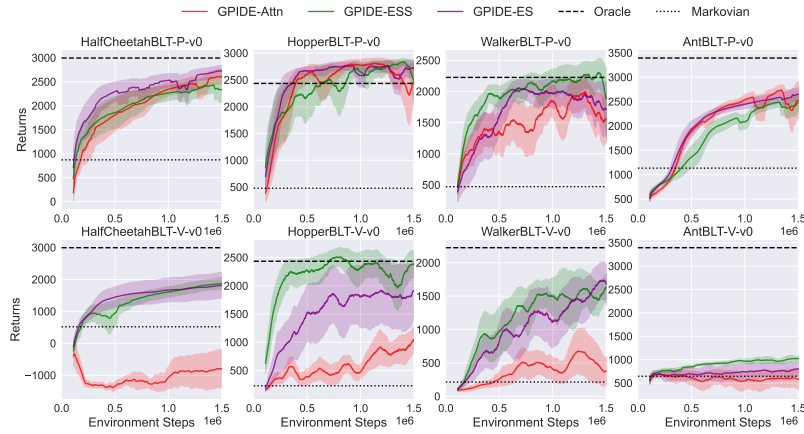


Figure 21: PyBullet Performance Curve for Ablations.

Interestingly, we found that GPIDE policies often outperform the oracle policy on Hopper-P. While the oracle performance here was taken from Ni et al. [50], we confirmed this also happens with our own implementation of an oracle policy. We hypothesize that this may be due to the fact the GPIDE policy gets to see actions and rewards and the oracle does not.

G.6 Attention Scheme Visualizations

We generate the attention visualizations (as seen in Figure 4) by doing a handful of rollouts with a GPIDE policy using only attention heads. During this rollout we collect all of the weighting schemes, i.e. $\text{softmax}\left(\frac{q_{1:t}k_{1:t}^T}{\sqrt{D}}\right)$, generated throughout the rollouts and average them together. Below, we show additional attention visualizations. In all figures, each plot shows one of the different six heads. For each of these, the policies were evaluated on the same version of the environment they were trained on.




Figure 22: MSD-Fixed Attention.




Figure 23: MSD-Small Attention.




Figure 24: MSD-Large Attention.




Figure 25: DMSD-Fixed Attention.




Figure 26: DMSD-Small Attention.




Figure 27: DMSD-Large Attention.




Figure 28: Navigation No Friction Attention.




Figure 29: Navigation Friction Attention.




Figure 30: β_N Tracking Sim Attention.




Figure 31: β_N Tracking Rotation Attention.




Figure 32: β_N -Rotation Tracking Sim Attention.




Figure 33: β_N -Rotation Tracking Rotation Attention.




Figure 34: HalfCheetah-P Attention.




Figure 35: HalfCheetah-V Attention.




Figure 36: Hopper-P Attention.




Figure 37: Hopper-V Attention.




Figure 38: Walker-P Attention.




Figure 39: Walker-V Attention.




Figure 40: Ant-P Attention.




Figure 41: **Ant-V Attention**. Note that total path length is less than 64 here since the agent falls down pretty fast.