

378 **A Relationship to the Laplace Transform**

379 We often think of RL in terms of discrete timesteps, but the sum in Equation 11 could be replaced
 380 with an integral for continuous-time domains, where \tilde{x} is a function of time:

$$S(n) = \int_0^n \tilde{x}(\tau) \odot \exp(-\tau\alpha) \exp(-\tau i\omega)^\top d\tau. \quad (23)$$

381 If we look at just a single dimension of \tilde{x} , we have

$$S(n) = \int_0^n \tilde{x}(\tau) \exp(-\tau(\alpha + i\omega)) d\tau, \quad (24)$$

382 which is equivalent to the Laplace Transform L of a function $\tilde{x}(\tau)$. One way to interpret the connection
 383 between our aggregator and the Laplace transform is that the aggregator transforms the memory
 384 task into a pole-finding task in the S-plane, a well-known problem in control theory. When we train
 385 FFM for an RL task, we are attempting to find the poles and zeros that minimize the actor/critic loss
 386 for a batch of sequences.

387 **B Weight Initialization**

388 We find that our model is sensitive to initialization of the α, ω values. We compute the upper limit
 389 for α such that memory will retain $\beta = 0.01 = 1\%$ of its value after some elapsed number of
 390 timesteps we call t_e . This can be computed mathematically via

$$\frac{\log \beta}{t_e}. \quad (25)$$

391 The lower limit is set based on the maximum value a double precision float can represent, minus a
 392 small ϵ

$$\frac{\log 1.79 \times 10^{308}}{t_e} - \epsilon. \quad (26)$$

393 Note that if needed, we can choose a smaller lower bound, but the input must be chunked into
 394 sequences of length t_e and run in minibatches. Since we can compute minibatches recurrently,
 395 the gradient spans across all minibatches rather than being truncated like a quadratic transformer.
 396 Nonetheless, we did not need to do this in any of our experiments. We set α to be linearly spaced
 397 between the computed limits

$$\alpha = [\alpha_1, \dots, \alpha_j, \dots, \alpha_n]^\top \quad (27)$$

$$\alpha_j = \frac{j}{m} \frac{\log \beta}{t_e} + (1 - \frac{j}{m}) \left(\frac{\log 1.79 \times 10^{308}}{t_e} - \epsilon \right) \quad (28)$$

398 We initialize the ω terms using a similar approach, with linearly spaced denominators between one
 399 and t_e

$$\omega = 2\pi / [\omega_1, \omega_2, \dots, \omega_n]^\top \quad (29)$$

$$\omega_j = \frac{j}{c} + (1 - \frac{j}{c})t_e \quad (30)$$

400 All other parameters are initialized using the default Pytorch initialization.

401 **C Computing States in Parallel with Linear Space Complexity**

402 Here, we show how the recurrent formulation can be computed in parallel. A naive implementation
 403 would use $O(n^2)$ space, but using a trick, we can accomplish this in $O(n)$ space.

404 Assume we have already computed hidden state S_{k-1} for some sequence. Our job is now to compute
 405 the next $n-k+1$ recurrent states S_k, S_{k+1}, \dots, S_n in parallel. We can rewrite Equation 11 in matrix

406 form:

$$S_{k:n} = \begin{bmatrix} S_k \\ S_{k+1} \\ S_{k+2} \\ \vdots \\ S_n \end{bmatrix} = \begin{bmatrix} \gamma^0 \odot (\mathbf{x}_k \mathbf{1}_c^\top) + \gamma^1 \odot S_{k-1} \\ \gamma^0 \odot (\mathbf{x}_{k+1} \mathbf{1}_c^\top) + \gamma^1 \odot (\mathbf{x}_k \mathbf{1}_c^\top) + \gamma^2 \odot S_{k-1} \\ \gamma^0 \odot (\mathbf{x}_{k+2} \mathbf{1}_c^\top) + \gamma^1 \odot (\mathbf{x}_{k+1} \mathbf{1}_c^\top) + \gamma^2 \odot (\mathbf{x}_k \mathbf{1}_c^\top) + \gamma^3 \odot S_{k-1} \\ \vdots \\ \left(\sum_{j=k}^n \gamma^j \odot (\mathbf{x}_{k+j} \mathbf{1}_c^\top) \right) + \gamma^{n+1} \odot S_{k-1} \end{bmatrix} \quad (31)$$

407 We can write the closed form for the p th row of the matrix, where $0 \geq p \geq n - k$

$$S_{k+p} = \left(\sum_{j=0}^p \gamma^{p-j} \odot (\mathbf{x}_{k+j} \mathbf{1}_c^\top) \right) + \gamma^{p+1} \odot S_{k-1} \quad (32)$$

408 Unfortunately, it appears we will need to materialize $\frac{(n-k+1)^2}{2}$ terms:

$$\gamma^0 \odot (\mathbf{x}_k \mathbf{1}_c^\top) \quad (33)$$

$$\gamma^1 \odot (\mathbf{x}_k \mathbf{1}_c^\top), \gamma^0 \odot (\mathbf{x}_{k+1} \mathbf{1}_c^\top) \quad (34)$$

$$\vdots \quad (35)$$

$$\gamma^{n-k} \odot (\mathbf{x}_k \mathbf{1}_c^\top), \gamma^{n-k+1} \odot (\mathbf{x}_{k+1} \mathbf{1}_c^\top), \dots, \gamma^0 \odot (\mathbf{x}_n \mathbf{1}_c^\top) \quad (36)$$

409 However, we can factor out γ^{-p} from Equation 32 via a combination of the distributive property and
410 the product of exponentials

$$S_{k+p} = \gamma^p \odot \left(\sum_{j=0}^p \gamma^{-j} \odot (\mathbf{x}_{k+j} \mathbf{1}_c^\top) \right) + \gamma^{p+1} \odot S_{k-1} \quad (37)$$

411 Now, each γ^j is associated with a single \mathbf{x}_j , requiring just $n - k + 1$ terms:

$$\gamma^0 \odot (\mathbf{x}_k \mathbf{1}_c^\top) \quad (38)$$

$$\gamma^0 \odot (\mathbf{x}_k \mathbf{1}_c^\top), \gamma^1 \odot (\mathbf{x}_{k+1} \mathbf{1}_c^\top) \quad (39)$$

$$\vdots \quad (40)$$

$$\gamma^0 \odot (\mathbf{x}_k \mathbf{1}_c^\top), \gamma^1 \odot (\mathbf{x}_{k+1} \mathbf{1}_c^\top), \dots, \gamma^n \odot (\mathbf{x}_n \mathbf{1}_c^\top) \quad (41)$$

412 We can represent each of these rows as a slice of a single $n - k + 1$ length tensor, for a space
413 complexity of $O(n - k)$ or for $k = 1, O(n)$.

414 Finally, we want to swap the exponent signs because it provides better precision when working with
415 floating point numbers. Computing small values, then big values is more numerically stable than
416 the other way around. We want to compute the inner sum using small numbers (γ^+ results in small
417 numbers while γ^- produces big numbers; $n - k$ is positive and $k - n$ is negative). We can factor
418 γ^{k-n} out of the sum and rewrite Equation 37 as

$$S_{k+p} = \gamma^{k-n+p} \odot \left(\sum_{j=0}^p \gamma^{n-k-j} \odot (\mathbf{x}_{k+j} \mathbf{1}_c^\top) \right) + \gamma^{p+1} \odot S_{k-1} \quad (42)$$

419 If we let $t = n - k$, this is equivalent to Equation 13:

$$S_{k+p} = \gamma^{p+1} \odot S_{k-1} + \gamma^{p-t} \odot \sum_{j=0}^p \gamma^{t-j} \odot (\mathbf{x}_{k+j} \mathbf{1}_c^\top), \quad S_{k+p} \in \mathbb{C}^{m \times c}. \quad (43)$$

420 We also need to compute $2(n - k + 1)$ gamma terms: $\gamma^{k-n}, \dots, \gamma^{n-k+1}$, resulting in linear space
421 complexity $O(n)$ for a sequence of length n .

422 **D Benchmark Details**

423 We utilize the episodic reward and normalized episodic reward metrics. We record episodic rewards
424 as a single mean for each epoch, and report the maximum reward over a trial.
425 We do not modify the hyperparameters for PPO, SAC, or TD3 from the original benchmark papers.
426 We direct the reader to [Morad et al., 2023] for PPO hyperparameters and [Ni et al., 2022] for SAC
427 and TD3 hyperparameters.

428 **D.1 Hardware and Efficiency Information**

429 We trained on a server with a Xeon CPU running Torch version 1.13 with CUDA version 11.7,
430 with consistent access to two 2080Ti GPUs. Wandb reports that we used roughly 161 GPU days of
431 compute to produce the POMDP-Baselines results, 10 GPU days for the POPGym results, and 45
432 GPU days for the FFM ablation using POPGym. This results in a total of 216 GPU days. Since we
433 ran four jobs per GPU, this corresponds to 54 days of wall-clock time for all experiments.

434 **D.2 PPO and POPGym Baselines**

435 Morad et al. [2023] compares models along the recurrent state size, with a recurrent state size of 256.
436 For fairness, We let $m = 32$ and $c = 4$, which results in a complex recurrent state of 128, which
437 can be represented as a 256 dimensional real vector. We initialize α, ω following Appendix B for
438 $t_e = 1024, \beta = 0.01$. We run version 0.0.2 of POPGym, and compare FFM numerically with the
439 MMER score from the paper in Table 3.

440 **D.3 SAC, TD3, and POMDP-Baselines**

441 [Ni et al., 2022] does not compare along the recurrent state size, but rather the hidden size, while
442 utilizing various hidden sizes h . In other words, the LSTM recurrent size is twice that of the GRU.
443 We let $c = h/32$ and $m = h/c$, so that $mc = h$. This produces an equivalent FFM configuration
444 to the POPGym baseline when $h = 128$, with equivalent recurrent size in bytes to the LSTM (or
445 equivalent in dimensions to the GRU). The paper truncates episodes into segments of length 32 or 64
446 depending on the task, so we let $t_e = 128$ to ensure that information can persist between segments.
447 Thanks to the determinism provided in the paper, readers should be able to reproduce our exact
448 results using the random seeds 0, 1, 2, 3, 4. We utilize separate memory modules for the actor and
449 critic, as done in the paper. We base our experiments off of the most recent commit at the time of
450 writing, 4d9cbf1.

| Model | MMER |
|------------|--------------|
| MLP | 0.067 |
| PosMLP | 0.064 |
| FWP | 0.112 |
| FART | 0.138 |
| S4D | -0.180 |
| TCN | 0.233 |
| Fr.Stack | 0.190 |
| LMU | 0.229 |
| IndRNN | 0.259 |
| Elman | 0.249 |
| GRU | 0.349 |
| LSTM | 0.255 |
| DNC | 0.065 |
| FFM | 0.400 |

Table 3: MMER score comparison from the POPGym paper

451 **E Benchmark Lineplots**

452 We provide cumulative max reward lineplots in Figure 9, Figure 7, and Figure 8 for all the experiments we ran.

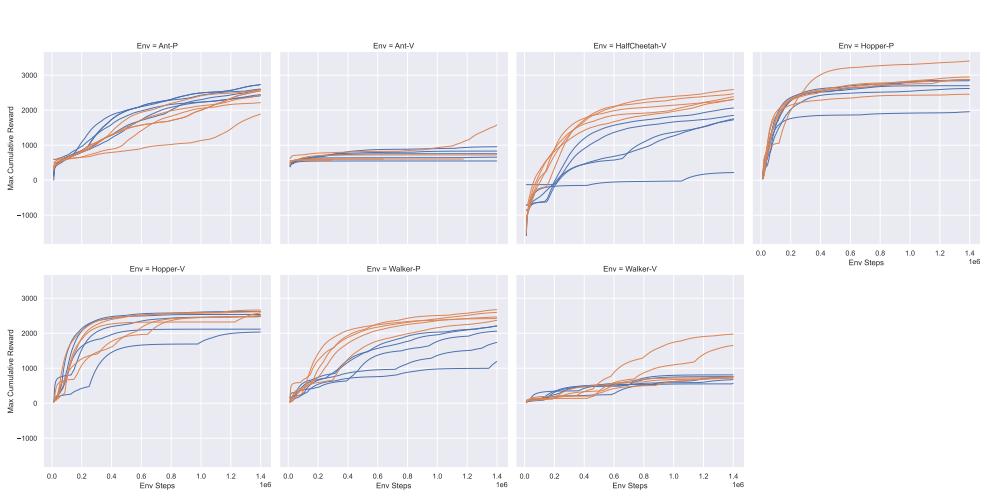


Figure 7: SAC lineplots for POMDP-Baselines, where each trial is plotted separately.

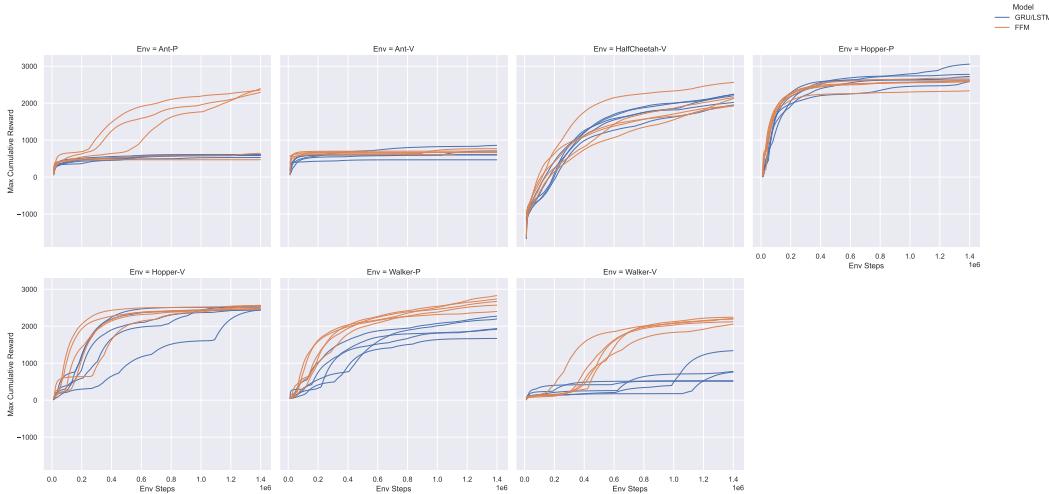


Figure 8: TD3 lineplots for POMDP-Baselines, where each trial is plotted separately.

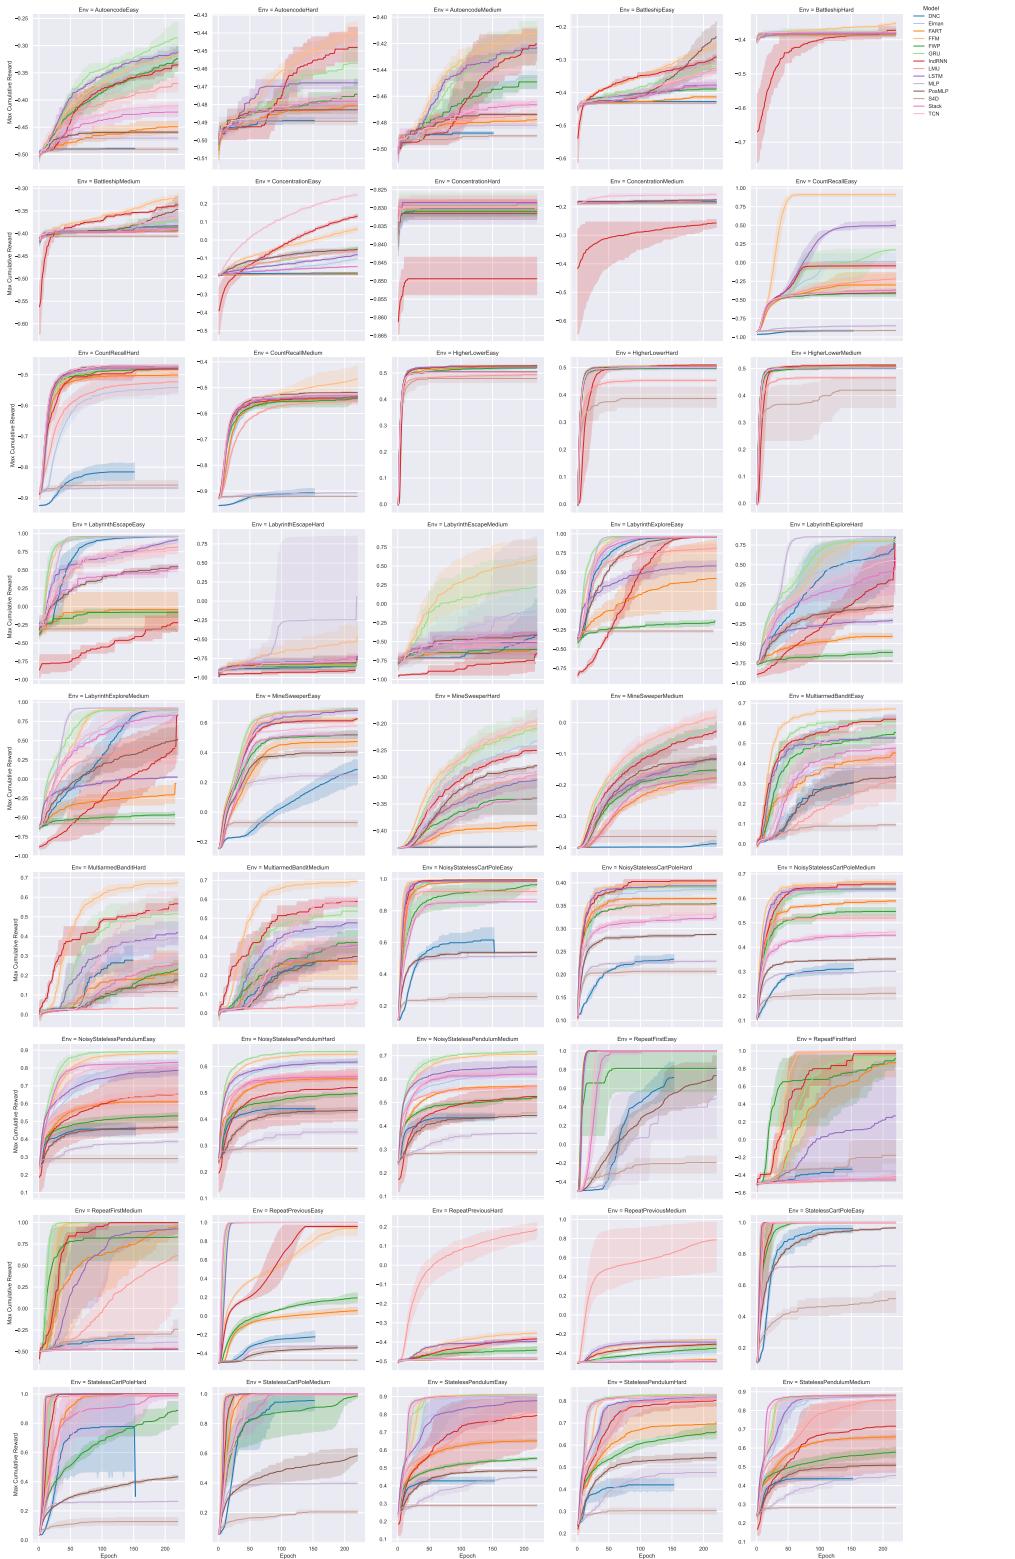


Figure 9: Lineplots for POPGym, where the shaded region represents the 95% bootstrapped confidence interval.

454 F POPGym Comparisons by Model

455 We provide Figure 10, Figure 11, Figure 12 showing the relative FFM return compared to the other
 456 12 POPGym models, including temporal convolution, linear transformers, and more.

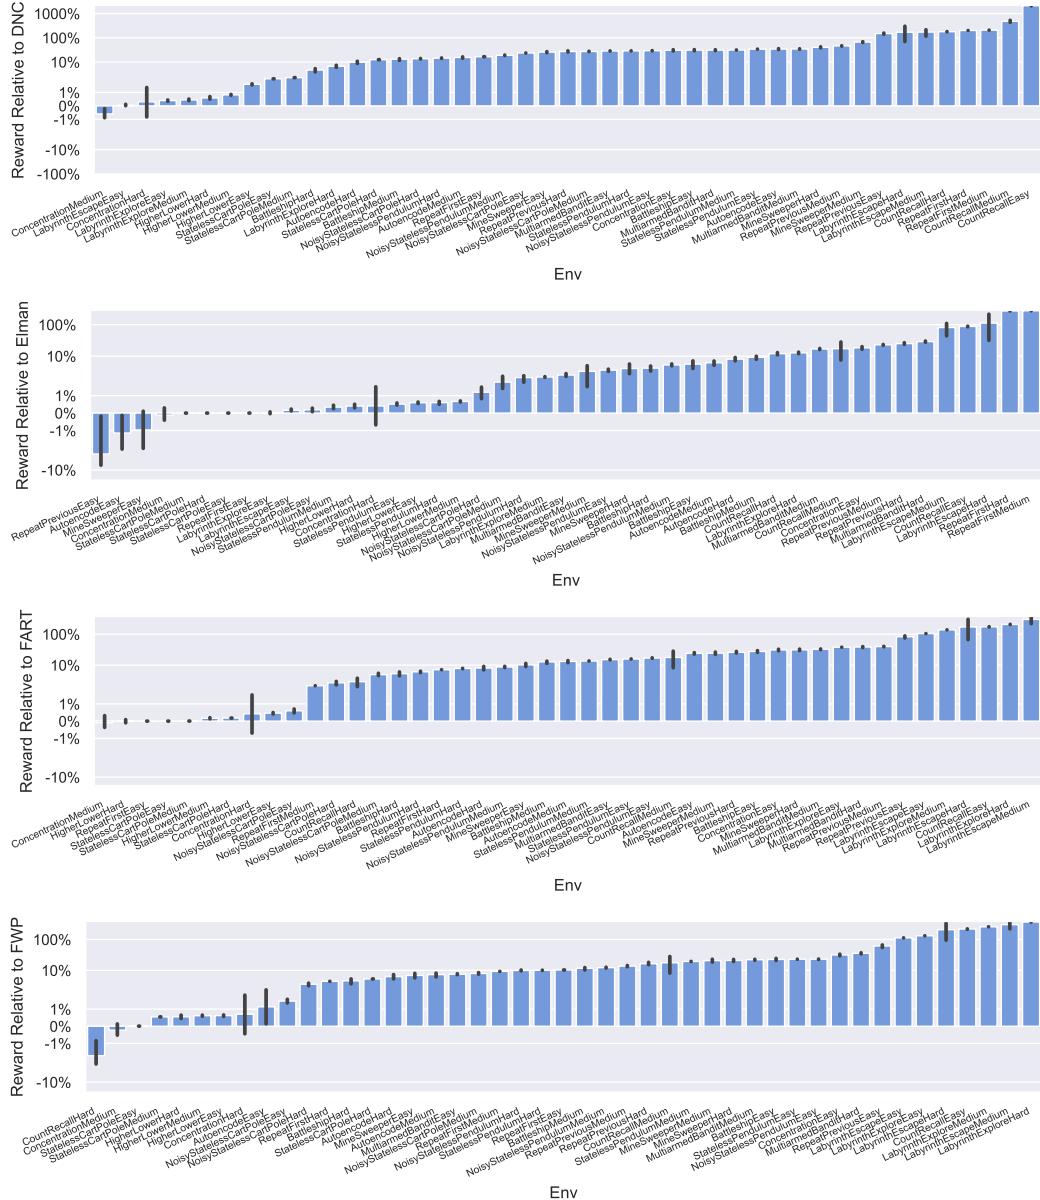


Figure 10: Relative POPGym returns compared to FFM.

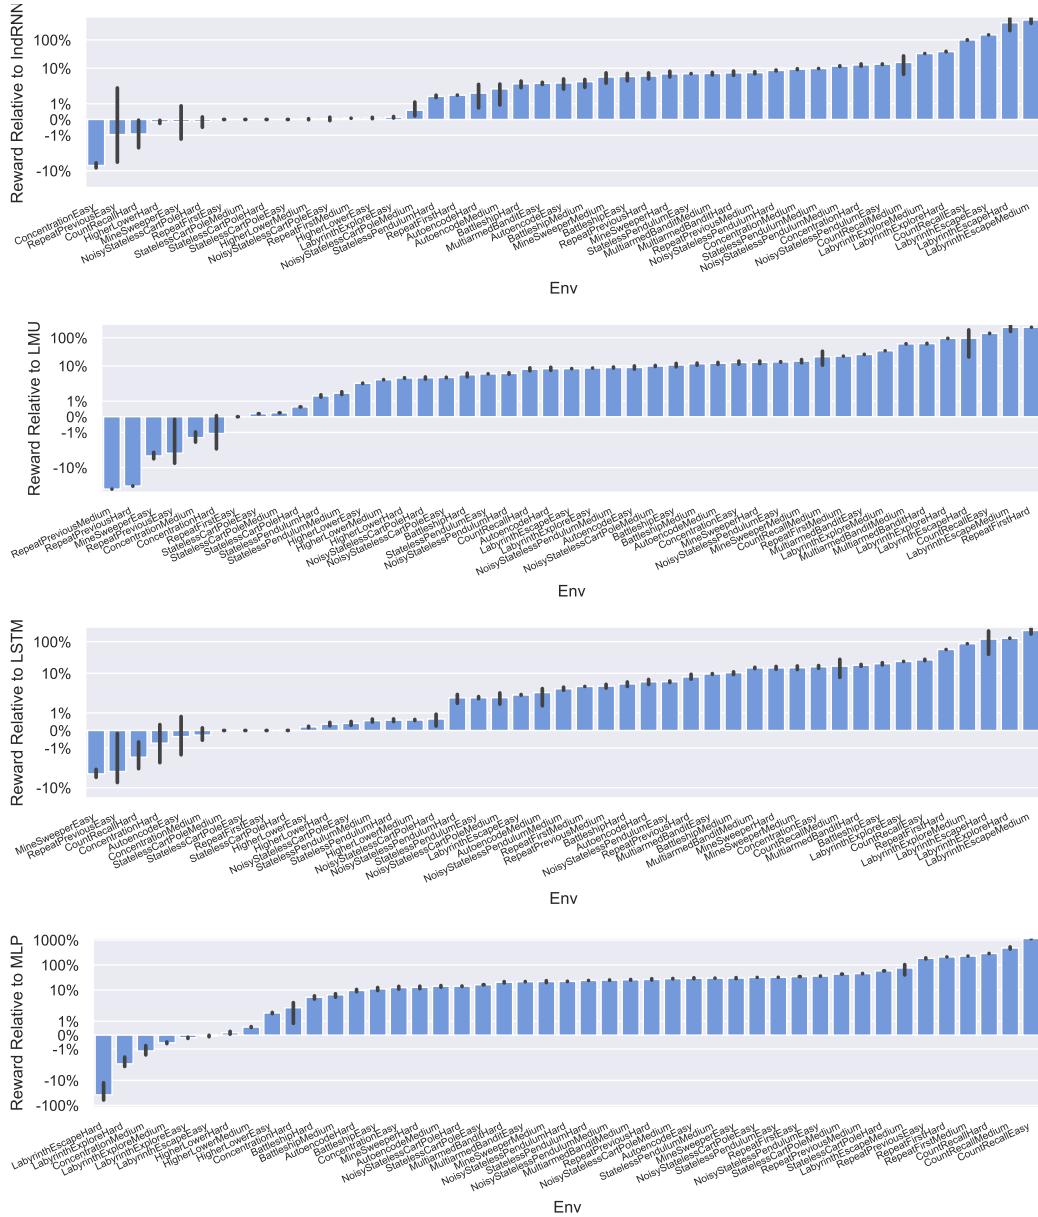


Figure 11: Relative POPGym returns compared to FFM.

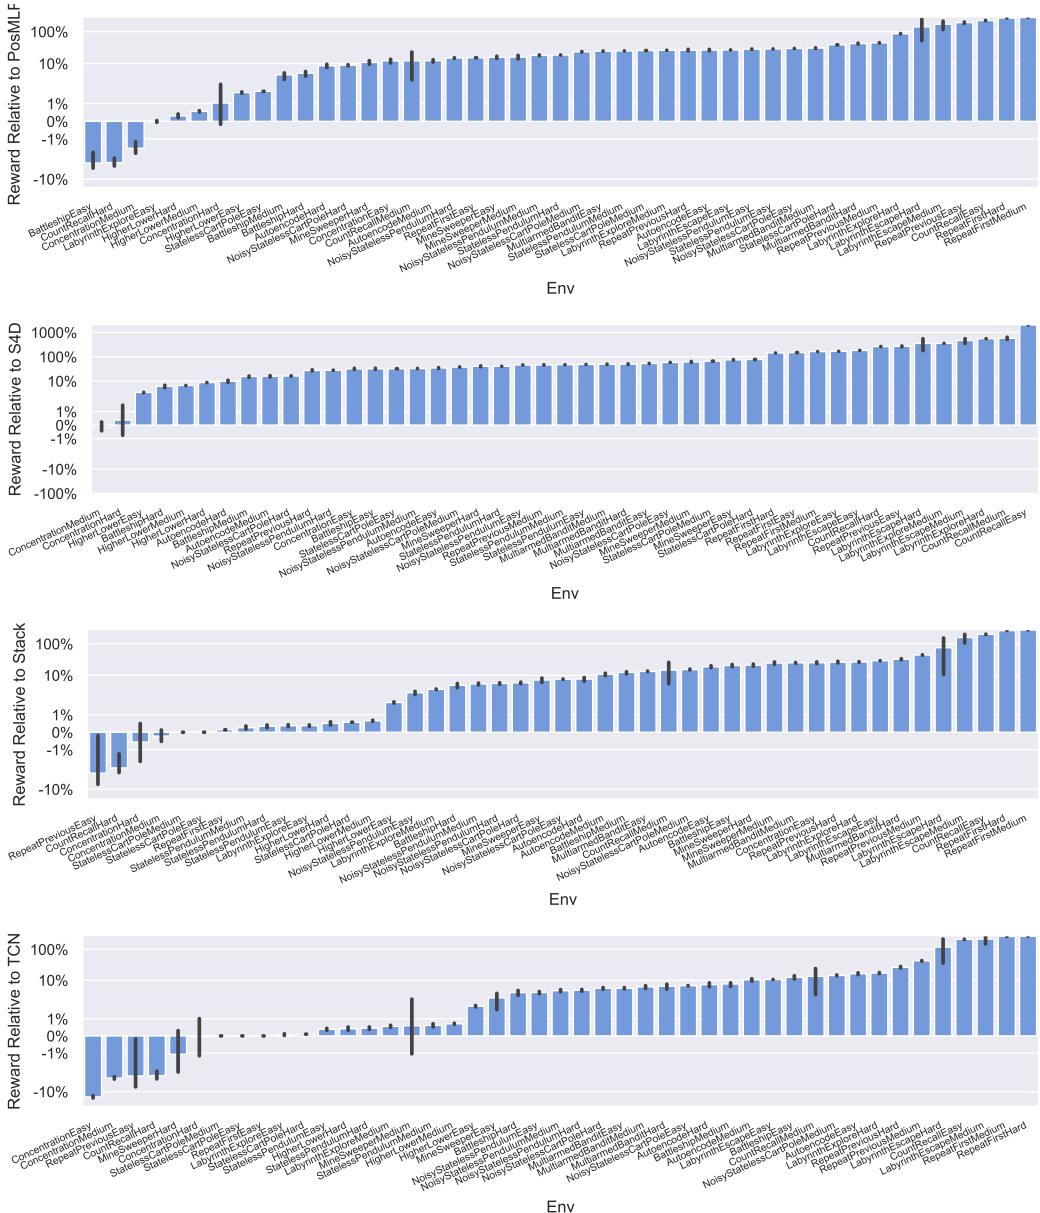


Figure 12: Relative POPGym returns compared to FFM.

457 **G Efficiency Statistic Details**

- 458 This section explains how we computed the bar plots showing efficiency statistics. We construct the
459 POPGym models and perform one PPO training epoch. The train time metric corresponds to the
460 time spent doing forward and backward passes over the data. One epoch corresponds to the epoch
461 defined in POPGym: 30 minibatches of 65,336 transitions each, with an episode length of 1024, a
462 hidden dimension of 128, and a recurrent dimension of 256. We do this 10 times and compute the
463 mean and confidence interval. Torch GRU, LSTM, and Elman networks have specialized CUDA
464 kernels, making them artificially faster than LMU and IndRNN which are written in Torch and
465 require the use of for loops. We utilize the pure python implementation of these models, wrapping
466 them in a for loop instead of utilizing their hand-designed CUDA kernels. We consider this a fair
467 comparison since this makes the GRU, Elman, and LSTM networks still run slightly faster than
468 IndRNN and LMU (Torch-native RNNs). FFM is also written in Torch and does not have access
469 to specialized CUDA kernels. CUDA programmers more skilled than us could implement a custom
470 FFM kernel that would see a speedup similar to the GRU/LSTM/Elman kernels.
- 471 To compute inference latency, we turn off Autograd and run inference for 1024 timesteps, computing
472 the time for each forward pass. We do this 10 times and compute the mean and 95% confidence
473 interval.
- 474 To compute memory usage, we utilize Torch’s GPU memory tools and record the maximum memory
475 usage at any point during the training statistic. Memory usage is constant between trials.
- 476 For reward, we take the mean reward over all environments, split by model and trial. We then report
477 the mean and 95% confidence interval over trials and models.