

504 A IMPLEMENTATION DETAILS

505 A.1 DT-MEM NETWORK ARCHITECTURE

506 Table 3 summarizes the different model configurations used for evaluation. In this section, we
 507 describe these model configurations in detail. While Table 3 provides a summary, we will also
 508 provide additional information here. DT-Mem, PDT and HDT are all share the same transformer
 509 architectures. However, for task-adaptation, HDT utilizes a pre-trained 2.3M hyper-network, while
 510 DT-Mem introduces 147K LoRA parameters. To compare with MDT, we use the same parameter
 511 size as reported in Lee et al. (2022).

Model	Layers	Hidden size (d)	Heads	Params	Memory Size	Memory Module Params
HDT	4	512	8	13M	N.A.	N.A.
MDT-200M	10	1280	20	200M	N.A.	N.A.
DT-Mem	4	512	8	13M	559K	7M

Table 3: Detailed Model Sizes

512 A.2 HYPER-PARAMETERS

513 In this section, we will delve into the specifics of the model parameters. Understanding these
 514 parameters is key to understanding the workings of the model. It is worth noting that the
 515 source code for this model is publicly available at <https://anonymous.4open.science/r/DT-Mem-Submission277/README.md>.
 516 This allows for a deeper understanding of the model’s
 517 inner workings and may facilitate the replication of its results.

Hyperparameters	Value
K (length of context)	28
dropout rate	0.1
maximum epochs	1000
steps for each epoch	1000
optimizer learning rate	1e-4
weight decay	1e-4
gradient norm clip	1.
data points for each dataset	500,000
batch size	64
memory slots	1290
activation	GELU
optimizer	AdamW
scheduler	LambdaLR

Table 4: Hyperparameters for DT-Mem training

518 A.3 TRAINING AND FINE-TUNING ALGORITHM

519 In this section, we present the pre-training DT-Mem in Appendix A.3 and fine-tuning DT-Mem with
 520 LoRA in Appendix 5.5.

521 We pre-train DT-Mem on multiple offline datasets. Each gradient update of the DT-Mem model
 522 considers information from each training task.

523 We fine-tune the memory module to adapt to each downstream task. To achieve this, we fix the
 524 pre-trained DT-Mem model parameters and add additional LoRA parameters for the memory module
 525 feed-forward neural networks. The fine-tuning dataset is used to update these LoRA parameters only.

Algorithm 1 Pre-train DT-Mem

- 1: **for** T episodes **do**
 - 2: **for** Task $\mathcal{T}_i \in T^{train}$ **do**
 - 3: Sample trajectories $\tau = (s_0, a_0, r_0, \dots, s_H, a_H, r_H)$ from the dataset \mathcal{D}_i .
 - 4: Split trajectories into different segments with length K and calculate return-to-go in the input sequence.
 - 5: Given $\hat{\tau}_{t+1:t+K}$, compute the sequence embedding e_{seq} .
 - 6: Update the memory module and retrieve the relative information as E_{out}
 - 7: Given E_{out} , predict actions \tilde{a}_t , reward \tilde{r}_t , and return-to-go \tilde{R}_t .
 - 8: Compute the loss according to Eqn. [1](#).
 - 9: Update all module parameters.
 - 10: **end for**
 - 11: **end for**
-

Algorithm 2 Fine-tuning DT-Mem

Require: Fine-tuning dataset $\mathcal{T}^i \in T^{test}$ dataset \mathcal{D}^i for \mathcal{T}^i . Initialize LoRA parameters $\hat{B}^q, \hat{B}^k, \hat{B}^v, \hat{A}^q, \hat{A}^k, \hat{A}^v, B^q, A^q, B^k, A^k$.

- 1: **for** T steps **do**
 - 2: Split trajectories into different segments with length K and calculate return-to-go in the input sequence.
 - 3: Given $\hat{\tau}_{t+1:t+K}$, compute the sequence embedding e_{seq} .
 - 4: Update memory module using $\hat{Q} = M(\hat{W}^q + \hat{B}^q A^q), \hat{K} = M(\hat{W}^k + \hat{B}^k A^k), \hat{V} = M(\hat{W}^v + \hat{B}^v A^v), Q = M(W^q + B^q A^q), K = M(W^k + B^k A^k)$
 - 5: Retrieve the relative information as E_{out}
 - 6: Given E_{out} , predict actions \tilde{a}_t , reward \tilde{r}_t , and return-to-go \tilde{R}_t .
 - 7: Compute the loss according to Eqn. [1](#).
 - 8: Update LoRA parameters only.
 - 9: **end for**
-

Algorithm 3 Memory Operations

- 1: **Step 0: Memory Module Initialization**
 - 2: Initialize memory as a random matrix M where each row $m_i \in \mathbb{R}^d$ and $i \in [0, N]$.
 - 3:
 - 4: **Step 1: Input Sequence Organizing**
 - 5: Restructure input sequence into format $\langle \hat{r}_t, s_t, a_t \rangle$.
 - 6: Define embedding functions $g_s(s) = e_s, g_a(a) = e_a, g_r(\hat{r}) = e_{\hat{r}}$.
 - 7: Concatenate embeddings to form input sequence $E = [\cdots; e_{s_t}, e_{a_t}, e_{\hat{r}_t}; \cdots]$.
 - 8:
 - 9: **Step 2: Content-based Address**
 - 10: Use attention to locate memory slot for new input.
 - 11: Calculate position address $w = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)$.
 - 12: Define $Q = MW^q$ and $K = EW^k$.
 - 13:
 - 14: **for** N Times memory operations **do do**
 - 15: **Step 3: Memory Update**
 - 16: Calculate erasing vector ϵ^e and adding vector ϵ^a .
 - 17: Define $\hat{Q} = M\hat{W}^q, \hat{K} = E\hat{W}^k, \hat{V} = E\hat{W}^v$.
 - 18: Compute writing strength $\beta = \text{softmax}\left(\frac{\hat{Q}\hat{K}^T}{\sqrt{d}}\right)$.
 - 19: Calculate $\epsilon^e = w \odot (1 - \beta)$.
 - 20: Calculate $\epsilon^a = (w \odot \beta)\hat{W}^v x$.
 - 21: Update memory $M_n = M_{n-1} \odot (1 - \epsilon^e) + \epsilon^a$.
 - 22:
 - 23: **Step 4: Memory Retrieve**
 - 24: Retrieve information from memory for decision-making.
 - 25: Compute read position vector using content-based address.
 - 26: Retrieve memory $E_{out} = w \odot M_n$.
 - 27: $E = E_{out}$
 - 28: **end for**
 - 29: output E for action decoder.
-

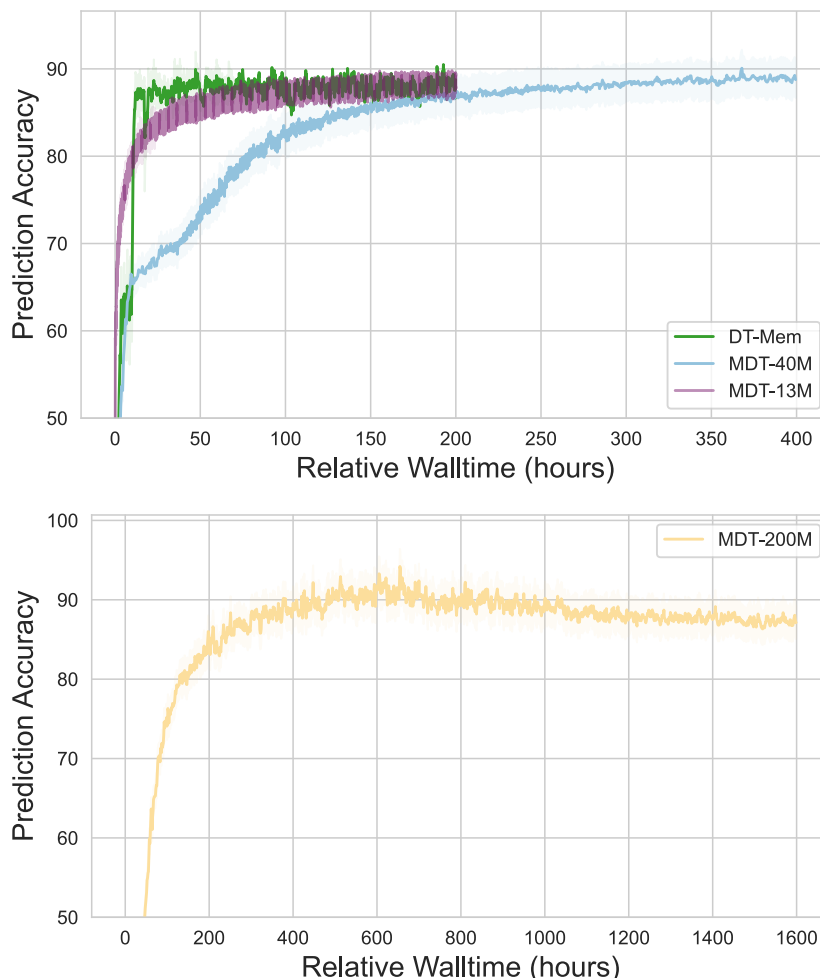


Figure 7: This graph shows the prediction accuracy during training. Each curve represents three runs with different random seeds. For better visualization, MDT-200M is displayed in a separate figure.

526 B ADDITIONAL EXPERIMENTS

527 B.1 EVALUATION PARAMETERS

528 To evaluate the performance of our model on Atari games, we randomly selected 16 different random
 529 seeds for evaluation. We chose the random seed by multiplying the number of runs by 100. For
 530 example, the random seed for run 6 is $6 \times 100 = 600$.

531 B.2 TRAINING EFFICIENCIES

532 To demonstrate training efficiency, we illustrate the model training curve in Figure 7. For the
 533 training curve, it is reasonable to report the prediction loss on the training dataset since we use a
 534 supervised loss. Here, the prediction accuracy consists of three parts: action prediction accuracy,
 535 reward prediction accuracy and return prediction accuracy. The y-axis shows the average value of
 536 these three predictions, and the x-axis is the relative walltime based on same computing resources.

537 B.3 THE ANALYSIS OF MEMORY SIZE

538 In this section, we investigate the impact of the memory module size on the performance of DT-Mem.
 539 We employ the Bayes optimization strategy to tune the parameters. It’s worth noting that the memory




Figure 8: The parameter tuning results for the number of memory slots. The blue curve shows the like from left to right over the x axis and plots the running average y value.

540 size is calculated by multiplying the number of memory slots by the size of each slot, which is fixed
 541 at 512 dimensions for the sake of evaluation simplicity. To expedite the hyper-parameter tuning
 542 process, we present the evaluation results based on 100k training steps of the StarGunner game. We
 543 assess various configurations of memory slots and calculate their corresponding average rewards
 544 over 16 runs. Figure 8 reveals several key findings: (1) Increasing the size of memory slots leads to
 545 a higher reward accumulation. Notably, there is a significant performance boost when the number
 546 exceeds 1200. (2) In summary, when the number of memory slots exceeds 1800, the performance of
 547 the system decreases. This decline occurs because there is a trade-off between the number of memory
 548 slots and the training steps. With a larger number of memory slots, it becomes necessary to allocate
 549 more training time.

550 B.4 ABLATION STUDY OF LORA ADAPTOR

	Meta-World ML45 Performances			Data size	Model	
	Train	Test (no-FT)	Test (FT)		Adap.	Per.
DT-Mem (hyper-net)	0.92 ± 0.01	0.23 ± 0.10	0.81 ± 0.15	30	5.7M	43.8%
DT-Mem	0.92 ± 0.00	0.20 ± 0.01	0.95 ± 0.10	10	147K	0.7%

Table 5: Ablation study results on Meta-World ML45 benchmarks. DT-Mem (hyper-net) denotes the variation of DT-Mem, which substitute LoRA adaptation module with hyper-networks. Adap. stands for adaptation parameters, and Per. stands for percentage of original model.

551 In this section, we conduct an ablation study of LoRA-based memory adaptor. We substitute LoRA
 552 adaptor with hyper-networks. Specifically, the parameters of the memory module are generated from
 553 hyper-networks. This approach is based on von Oswald et al. (2020), where hyper-networks take
 554 task-related information as input and generate the corresponding networks for the downstream MLP.
 555 We use the same approach and generate parameters that are conditioned on two types of inputs: the
 556 task embedding from the task encoder and the sequence embeddings from the Transformer module.

557 To generate task embeddings, we adopt the same idea from PDT (Xu et al., 2022), which demonstrates
 558 that a small part of trajectories can represent the task-related information. We further extend
 559 this idea to fully extract the task information. To achieve this goal, we use a Neural Networks
 560 (NNs) as a task encoder. Specifically, this task encoder is implemented as a transformer encoder-
 561 like structure Vaswani et al. (2017). We first formulate the first i steps of collected trajectories

562 $\tau_{0:i} = (s_0, a_0, r_0, \dots, s_i, a_i, r_i)$ as a task specific information. The task trajectory $\tau_{0:i}$ is treated
 563 as a sequence of inputs to the task encoder. The output of the task encoder is a task embedding
 564 $e_{task} \in \mathbb{R}^d$, where d is the dimension of the embedding.

565 Then, we concatenate the task embedding and sequence embedding $e = [e_{task}; e_{seq}]$ and input them
 566 to the hyper-networks. Specifically, we define the hyper-network as a function of $f_\omega(\cdot)$ parameterized
 567 by ω . The output $\Theta = f_\omega(e)$ is a set of parameters for the memory module.

568 According to the evaluation results in Table 5, the inclusion of a hyper-network in the DT-Memmodel
 569 improves generalization without the need for fine-tuning. However, it is worth noting that the
 570 hyper-network variant of DT-Mem(hyper-net) exhibits higher variance compared to DT-Mem. The
 571 primary reason for this higher variance is the uncertainty arising from the task information. In each
 572 run, different task-related sequences are collected, resulting in varying generated parameters for
 573 the memory module. Regarding the task fine-tuning results, we observe that the LoRA module
 574 outperforms other methods. This finding indicates that fine-tuning with LoRA enhances the model’s
 575 adaptability. We hypothesize that the size of the hyper-networks model plays a role in these results.
 576 Fine-tuning a large model size (5.7M) with a small step-size (100k steps in our case) becomes
 577 challenging. In an effort to improve hyper-networks fine-tuning performance, we increased the
 578 fine-tuning dataset from 10k episodes to 30k episodes. These findings suggest that LoRA-based
 579 fine-tuning demonstrates better data efficiency.

580 The motivations for using LoRA to fine-tune the model can be concluded in the following two
 581 reasons:

582 Hu et al. (2022) suggests that the LoRA method, in contrast to other adapters, maintains model quality
 583 without introducing inference latency or shortening input sequence length. Furthermore, it facilitates
 584 rapid task-switching in service deployments by sharing most model parameters. Parameter-efficient
 585 fine-tuning (PEFT) refines a limited number of model parameters, preserving most of the pre-trained
 586 LLM parameters, which reduces computational and storage demands (Hu et al., 2022). This approach
 587 also addresses catastrophic forgetting [4] and has outperformed standard fine-tuning in low-data and
 588 out-of-domain situations [5]. Besides, the results of full parameter fine-tuning vs. PEFT are shown in
 589 Table 6.

Game	PEFT	FFT-Single	FFT-All
Alien	127.4%	116.8%	113.9%
MsPacman	130.8%	122.8	77.1%
Pong	97.8%	93.7%	90.5%
SpaceInvaders	100.8%	86.8%	73.4%
StarGunner	158.3%	55.7%	40.6%

Table 6: Performance comparison of PEFT across various games

590 where PEFT stands for LoRA fine-tuning for all games together, FFT-single means full-parameter
 591 fine-tuning on a single game only, FFT-All stands for full-tine-tuning on all games together. Results
 592 are DQN-normalized score.

593 B.5 LORA HYPER-PARAMETERS TUNING

594 In this section, we explore the impact of LoRA hyper-parameters on the final fine-tuning results.
 595 LoRA employs three hyper-parameters: rank, lora_dropout, and lora_alpha. The rank parameter,
 596 denoted as m , determines the low-rank of adaptation matrices $B \in \mathbb{R}^{n \times m}$ and $A \in \mathbb{R}^{m \times d}$, as
 597 described in Section 4.4. The lora_dropout refers to the dropout rate applied to the LoRA neural
 598 networks, while lora_alpha controls the scaling factor of the LoRA outputs. Figure 9 presents the
 599 fine-tuning results, with the last column (eval/rew_mean/StarGur) specifically showcasing the
 600 fine-tuning results for the StarGunner game. To obtain the optimal set of parameters, we employ the
 601 Bayesian optimization method for parameter tuning, which suggests various parameter combinations
 602 that maximize the fine-tuning results.

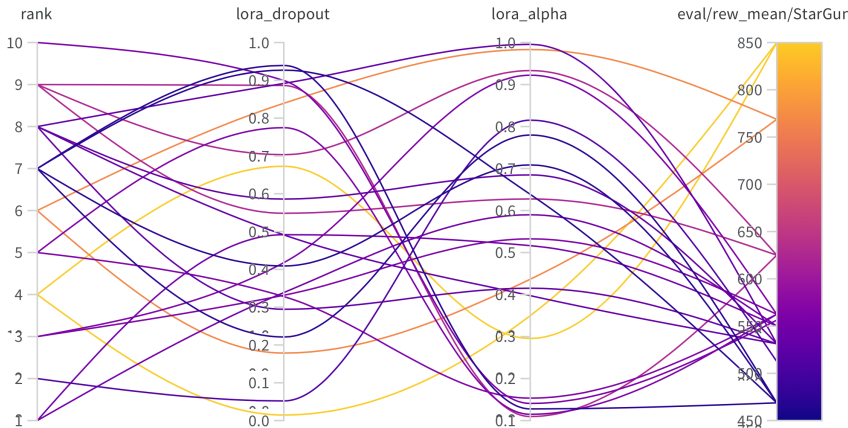


Figure 9: LoRA hyper-parameters tuning results.

Parameter	Importance score	Correlation score
rank	0.486	-0.132
lora_dropout	0.285	-0.561
lora_alpha	0.229	0.550

Table 7: Analysis of LoRA hyper-parameters

603 We further analyze these parameters and present the findings in Table 7. To gain insights, we utilize
 604 two widely used metrics in the MLOps platform Weights&Biases [1].

605 Regarding the **importance score**, we train a random forest model with the hyper-parameters as inputs
 606 and the metric as the target output. We report the feature importance values derived from the random
 607 forest. This hyper-parameter importance panel disentangles complex interactions among highly
 608 correlated hyper-parameters. It facilitates fine-tuning of hyper-parameter searches by highlighting the
 609 hyper-parameters that significantly impact the prediction of model performance.

610 The **correlation score** represents the linear correlation between each hyper-parameter and the chosen
 611 metric (in this case, val_loss). A high correlation indicates that when the hyper-parameter has a
 612 higher value, the metric also tends to have higher values, and vice versa. Correlation is a useful
 613 metric, but it does not capture second-order interactions between inputs and can be challenging to
 614 compare when inputs have widely different ranges.

615 As shown in Table 7, rank emerges as the most important hyper-parameter that requires careful
 616 tuning. The correlation score of rank is -0.132, indicating that a smaller rank number leads to better
 617 fine-tuning results. Based on our findings, a rank value of 4 yields the best outcome. Lora_dropout
 618 and lora_alpha exhibit similar importance scores, suggesting that these two parameters can be treated
 619 equally. The correlation score reveals that a smaller lora_dropout value and a larger lora_alpha
 620 result in improved performance.

621 B.6 ABLATION STUDIES ON DIFFERENT INPUT SEQUENCE ORGANIZING CHOICES

622 We examine two distinct approaches to input organization. The first approach is adopted from
 623 the trajectory transformer as outlined in (Janner et al., 2021), which organizes the inputs as
 624 $(s_1, \dots, s_t, a_1, \dots, a_t, r_1, \dots, r_t)$, grouping states, actions, and rewards accordingly. The second
 625 approach is derived from the decision transformer as described in (Chen et al., 2021), and is the
 626 method utilized in this study.

¹For better understanding, please refer to https://docs.wandb.ai/guides/app/features/panels/parameter-importance?_gl=1*4s7cuj*_ga*MTQxNjYxODU0OC4xNjgzNjY4Nzgz*_ga_JH1SJHJQXJ*MTY4NDc5NDkzNS40MS4xLjE2ODQ3OTQ5NDIuNTMuMC4w

Game	Choice one	Choice two (Ours)
Alien	211.9	239.6
MsPacman	637.1	713.4
Pong	19.0	19.1
SpaceInvaders	165.7	171.2
StarGunner	620.7	709.3

Table 8: Ablation studies on different choices of organizing. Each value represents raw scores in Atari games.

627 From the table above, we observe minor differences between the two sets of inputs. However, the
628 variance in outcomes between the two methodologies is not significant. Therefore, in this paper, we
629 empirically adopt the second approach for our design.

630 B.7 ABLATION STUDIES WITH DT

	DT-Mem (Ave)	DT-Mem FT (Ave)	DT-20M (Ave)
10k	-	-	10.1%
20k	-	-	9.8%
30k	-	-	15.3%
40k	-	-	22.6%
50k	51.0%	127.4%	41.8%
100k	-	-	83.1%
200k	-	-	120.3%
500k	-	-	170.7%

Table 9: Comparison with DT in different fine-tuning datasets

631 As shown in Table 9, the left-most column represents the size of the dataset used for training. As
632 seen in the table above, the generalized agent DT-Mem outperforms when compared to training on
633 the DT-20M 50k datasets. Fine-tuning DT-Mem on 50k datasets yields better results than training
634 DT-20M on 200k datasets.

635 B.8 FULL FINE-TUNING VS. LORA

636 ****Full Fine-tuning (FFT) vs. LoRA****: To assess whether the use of LoRA adversely affects
637 performance, we conducted experiments contrasting Full Fine-Tuning (FFT) of memory parameters
638 with LoRA. In this context, FFT-single refers to fine-tuning all parameters exclusively on a single
639 game, whereas FFT-All represents fine-tuning on the entire set of games simultaneously. Results are
DQN-normalized score. Based on above results, we conclude the following observations:

Game	PEFT	FFT-Single	FFT-All
Alien	127.4%	116.8%	113.9%
MsPacman	130.8%	122.8	77.1%
Pong	0%	0%	0%
SpaceInvaders	100.8%	86.8%	73.4%
StarGunner	158.3%	55.7%	40.6%

640

641 - LoRA appears to be the most consistently effective strategy across the games provided. - While
642 ****FFT-Single**** occasionally outperforms PEFT (like in Alien), ****FFT-All**** consistently trails
643 behind the other two.

644 The reason full fine-tuning is not comparable to PEFT comes from the following parts: 1. Fine-tuning
645 dataset size. Note that we only use 50k data in LoRA and full fine-tuning compares on 500k used in
646 MDT paper 2. The benefits of LoRA is: "This approach also addresses catastrophic forgetting and
647 has outperformed standard fine-tuning in low-data and out-of-domain situations"

648 B.9 ANALYZE OF INPUT MISLEADING

649 we conducted an experiment to assess the robustness of the proposed method against input distortion.
 650 This involved adding Gaussian noise to the input frames of Atari games. Specifically, we set the
 651 mean to 0 and experimented with various standard deviation values. The results are detailed in the
 652 table below:

	Alien	MsPacman	SpaceInvaders	StarGunner
MDT	3.8%	13.2%	8.6%	2.3%
DT-Mem	51.0%	69.3%	53.6%	62.2%
DT-Mem (std=0.5)	55.3%	67.6%	53.0%	57.8%
DT-Mem (std=1)	35.6%	56.1%	40.0%	34.6%
DT-Mem (std=2)	25.9%	35.6%	30.5%	21.1%

653 From the results above, we conclude that the proposed DT-Mem demonstrates greater robustness to
 654 noisy inputs compared to the MDT method. This is evident as the DT-Mem consistently outperforms
 655 MDT under various levels of Gaussian noise. Notably, the performance with a standard deviation of
 656 0.5 shows minimal difference compared to the no-noise scenario, illustrating DT-Mem’s effectiveness
 657 in mitigating the impact of varying input distortions.

658 C MEMORY MODULE VISUALIZATION

659 Figure 10 illustrates the visualization of the memory module. Since memory operations are trained in
 660 conjunction with the transformer module, we select a later training episode at random to mitigate
 661 uncertainties regarding operational parameters. Due to time constraints, we trained on only two
 662 games simultaneously. In the revised version of the paper, we intend to provide visualizations for all
 663 games. For clearer visualization, we opted for a memory module of a smaller size, containing 128
 664 memory slots.

665 Let’s first discuss how memory modules update within the same game. As observed in the figure, for
 666 the Amidar game, the actively updated memory slots concentrate around rows 18, 84, and 117. This
 667 pattern is consistent across episodes, albeit with reduced activity. Such a trend indicates that during
 668 each training iteration, the transformer agent tends to overwrite the same memory slot contents. We
 669 note a similar observation in the Assault game. Furthermore, we observe that the memory module’s
 670 activity diminishes in later episodes. For instance, in the Assault game, the active memory slot in
 671 row 12 during episode 200k becomes less active by episode 201k. We hypothesize that as training
 672 progresses, the accumulated knowledge becomes sufficiently robust for retrieval, reducing the need
 673 for updates.

674 Moving on, when comparing the activity of memory slots across different games, there are intriguing
 675 overlaps. For instance, comparing Amidar 200k and Assault 200k reveals that memory slots around
 676 row 120 are active in both games. We surmise that this region retains cross-task knowledge shared
 677 between games. Additionally, the varying attention across other memory slots demonstrates how
 678 these slots assist the agent in decision-making across diverse games.

679 D LIMITATIONS AND SOCIETAL IMPACT

680 **Limitations** The first limitation of our work is the sample efficiency of memory fine-tuning. The
 681 10% fine-tuning dataset is still sizeable, and we plan to explore more sample-efficient methods in the
 682 future. We could, for instance, consider a setting with more tasks, each one with less data, so that the
 683 inter-task generalization would be even more crucial to its performance. Additionally, this work does
 684 not propose a control strategy for collecting data on a new task. For future work, we plan to investigate
 685 online data collection methods, which include the design and learning of exploration strategies for an
 686 efficient fine-tuning on new tasks. Finally, the approach has been intuitively motivated, but it would
 687 be valuable to have a theoretical grounding that would show the structural limits of large models and
 688 how equipping them with a memory component overcomes them.

689 **Societal Impact** We do not foresee any significant societal impact resulting from our proposed
 690 method. The current algorithm is not designed to interact with humans or any realistic environment

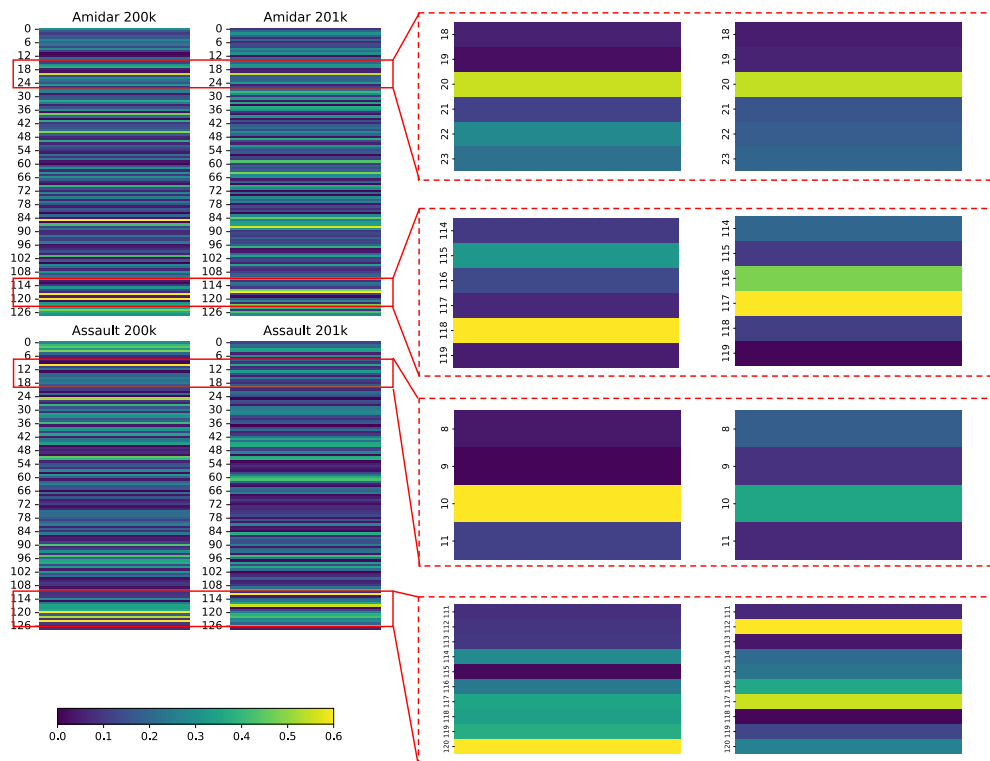


Figure 10: This visualization represents the memory module. In the figure, each row is derived from the mean of a vector that signifies a memory slot. Each depiction calculates the variation between two write operations in a single episode for each memory slot. Lighter shades indicate memory slots that have been actively updated post-write operations. The encircled areas highlight the comparison of active memory slots across different episodes.

691 yet. If one chooses to extend our methods to such situations, caution should be exercised to ensure
692 that any safety and ethical concerns are appropriately addressed. As our work is categorized in the
693 offline-RL domain, it is feasible to supplement its training with a dataset that aligns with human
694 intents and values. However, one must be wary that the way our architecture generalizes across tasks
695 is still not well understood, and as a consequence, we cannot guarantee the generalization of its
696 desirable features: performance, robustness, fairness, etc. By working towards methods that improve
697 the computational efficiency of large models, we contribute to increasing their access and reducing
698 their ecological impact.

699 E COMPARISON OF DT-MEM AND NEURAL EPISODIC CONTROL (NEC) IN 700 WRITING AND READING MEMORY

701 MEMORY MECHANISM

- 702 • **NEC:** Utilizes a Differentiable Neural Dictionary (DND) for storing experiences with
703 separate memories for each action, containing state representations (keys) and value function
704 estimates (values).
- 705 • **DT-Mem:** Integrates an internal memory module within a transformer framework, focusing
706 on storing, blending, and retrieving information for improving training efficiency and
707 generalization.

708 WRITING TO MEMORY

- 709 • **NEC:** Continuously adds new experiences and rapidly updates value function estimates in
710 memory.
- 711 • **DT-Mem:** Modifies or replaces existing information in the memory matrix using an attention
712 mechanism to calculate correlations and update memory with the attended weight of the
713 input sequence.

714 READING FROM MEMORY

- 715 • **NEC:** Implements context-based lookups in the DND to retrieve values, outputting a
716 weighted sum based on the similarity between the current state and stored keys.
- 717 • **DT-Mem:** Employs content-based addressing for memory retrieval, using attention mecha-
718 nisms to read from the updated memory and inform decision-making.

719 DISTINCTIVE FEATURES AND ADVANTAGES

- 720 • **NEC:** Designed for rapid assimilation and action upon new experiences with specialized
721 and swift updates for each action.
- 722 • **DT-Mem:** Aims to enhance generalization across tasks and reduce catastrophic forgetting
723 by integrating memory with the transformer’s sequential data handling capabilities.