
Complementary Meta-Reinforcement Learning for Fault-Adaptive Control

Ibrahim Ahmed¹, Marcos Quiñones-Grueiro², and Gautam Biswas³

^{1, 2, 3} *Vanderbilt University, Nashville TN, 37209 USA*

ibrahim.ahmed@vanderbilt.edu

gautam.biswas@vanderbilt.edu

marcos.quinones@vanderbilt.edu

ABSTRACT

Faults are endemic to all systems. Adaptive fault-tolerant control maintains degraded performance when faults occur as opposed to unsafe conditions or catastrophic events. In systems with abrupt faults and strict time constraints, it is imperative for control to adapt quickly to system changes to maintain system operations. We present a meta-reinforcement learning approach that quickly adapts its control policy to changing conditions. The approach builds upon model-agnostic meta learning (MAML). The controller maintains a complement of prior policies learned under system faults. This “library” is evaluated on a system after a new fault to initialize the new policy. This contrasts with MAML, where the controller samples new policies from a distribution of similar systems at each update step to achieve the new policy. Our approach improves sample efficiency of the reinforcement learning process. We evaluate our approach on an aircraft fuel transfer system under abrupt faults.

1. INTRODUCTION

No physical system is immune to degradation, changing environments, and faults. Since such situations can occur during operation, it is important the system respond to these changes in a way that the system continues to operate, be it in a degraded manner. This ensures safety and cost-effectiveness through less down-time. Fault-tolerant control (FTC) (Blanke, Kinnaert, Lunze, Staroswiecki, & Schröder, 2006) seeks to keep a faulty system operating, but within an acceptable margin of sub-optimal performance. This relaxes the constraints on the designers to make a system completely fail-safe and allows for a tradeoff between design and operating costs.

Data-driven approaches to FTC (MacGregor & Cinar, 2012; Hongm, Tian-You, Jin-Liang, & Brown, 2009) exploit the preponderance of data collected from system operations. They generate models that avoid the need for time-consuming and accurate physics-based simulations of system dynamics

to analyze and respond to different situations that may occur in the system. However, such methods depend on the data to span the breadth of operating conditions, and the model has to contain sufficient detail to capture multiple operating modes. This represents another compromise between design and operating costs.

In many cases, systems are complex, the number of possible faults are large, and faults that have not been seen before can occur during operations. There may not be precedent in the data to model such behaviors. A data-driven control approach will not then have “ground truth” to learn from and recall a sufficiently optimal control policy. Reinforcement learning (RL) is a semi-supervised approach to machine learning. It forfeits dependence on labelled ground truth and instead relies on accumulated feedback (i.e. experience gained) from a sequence of actions to determine a globally optimal policy. This ability to learn during operations alleviates design time effort and costs.

RL relies on gathered experience to accurately evaluate actions. This can be represented as a dynamic programming problem (Bellman, 1966) that typically has a closed-form solution, but for large systems, suffers from the curse of dimensionality. Advancements to RL have used function approximations of values to overcome the computational intractability of the problem (Boyan & Moore, 1995; Baird, 1995). However, the dependence on data to learn such approximations limits how fast and how accurately a RL-based controller can accommodate faults.

In our past work (Ahmed, Quiñones-Grueiro, & Biswas, 2020), we developed data-driven models to supplement experience with the real environment and simulate faults. In this work, we employ meta-RL for faster adaption of the RL algorithm to collected data samples. Our approach is not dependent on the time-consuming step of a data-driven model being learned first, however one can be used. The next section provides a background on RL and meta-RL. Section 3 describes our approach, and section 4 evaluates it on a simu-

lation of a fuel-transfer system. Finally, section 5 places our work in the context of extant research.

2. PRELIMINARIES

This section briefly introduces the RL approach and then discusses Model Agnostic Meta Learning in the context of RL-based control.

2.1. Reinforcement Learning

Reinforcement Learning (RL) is a semi-supervised approach to machine learning. A RL problem consists of a controller interacting with its environment. The environment can be modeled as a single Markov Decision Process (MDP) sampled from a population of available processes, $p \in P$. At a time t the controller perceives the environment’s state $x_t \in X$, and uses its policy $\pi : X \rightarrow U$ to take an action $u_t \in U$. The environment goes into a new state x_{t+1} governed by its transition function $T : X \times U \rightarrow X$ and emits a reward signal $r_t \in \mathbb{R}$, defined by the function $R : X \times U \times X \rightarrow \mathbb{R}$. The combination of (X, U, T, R) constitutes a MDP, $p \in P$.

The goal of a RL is to maximize the return, $J_\pi(x, u)$, which represents the total discounted cumulative reward for an action from each state when a policy, π is followed. A discount factor $\gamma \in [0, 1]$ is used to weigh immediate rewards over delayed rewards and to ensure convergence of the discounted reward function. The maximum future discounted reward for an action from a state is its value $V : X \times U \rightarrow \mathbb{R}$:

$$\begin{aligned} V(x_T, u_T) &= \max_{\pi} J_{\pi}(x_T, u_T) \\ &= \max_{\pi} \sum_{t=T}^{\infty} \gamma^{t-T} \cdot r_t \\ &= r_T + \gamma \cdot \max_{u_{T+1}} V(x_{T+1}, u_{T+1}) \end{aligned} \quad (1)$$

Policy gradient algorithms (Sutton, McAllester, Singh, & Mansour, 2000) in RL parametrize π with parameters θ , i.e. π_{θ} . The parameters θ are be the weights of a model representing the policy, for e.g. neural network. During training, they directly learn π_{θ} by implicitly optimizing for V using gradient ascent on the gain function $G \leftarrow \mathbb{E}[J_{\pi}(x, u)]$. Gradient ascent produces iterative updates to θ the size of which is determined by the learning rate $\alpha \in [0, 1]$.

Parameter updates at each iteration are dependent on experienced rewards under the latest policy. This is known as *on-policy* RL. This approach is sample inefficient because new trajectories of interactions need to be obtained at each iteration. A way around this is to use *importance sampling* in the gain function. By modeling the policy as a stochastic function over actions, $\pi_{\theta}(x | u)$, the relative probabilities, known as importance ratios, of the same trajectory under different

policies can be obtained. Thus, the gain function can reuse the same batch to update θ by weighing cumulative rewards.

$$\begin{aligned} G &= \mathbb{E}_{x_0 \sim X} \left(\prod_{t=0}^{\infty} \frac{\pi_{\theta}(x_t | u_t)}{\pi_{\theta_k}(x_t | u_t)} \right) J_{\pi_{\theta_k}}(x_0, u_0) \\ \theta_{k+1} &= \theta_k + \alpha \cdot \nabla_{\theta} G \end{aligned} \quad (2)$$

Large gradient updates may cause the next iteration of π_{θ} to overshoot, thus missing the optimum, causing the learning process to diverge altogether. Proximal Policy Optimization (PPO) (Schulman, Wolski, Dhariwal, Radford, & Klimov, 2017) clips the size of gradient updates by restricting the importance ratios between iterations. Thus a policy does not drastically change between updates. We use PPO in this work to learn the control policy under fault conditions.

2.2. Model-Agnostic Meta Learning

Meta-learning seeks to speed up a machine learning process through introspection. Essentially, it learns how to learn. In a RL context, meta-learning seeks to quickly adapt a policy trained on one process to another.

Model-agnostic Meta Learning (MAML) (Finn, Abbeel, & Levine, 2017) speeds up the optimization of any model learned through gradient updates. It does so by running an inner *introspective* loop within each iteration of a gradient update to the model’s parameters, which is designated as the outer loop. In the inner loop, variants of the process are sampled as $p_i \sim P$. The current model parameters θ are then optimized by training for several interactions on each p_i using gradient ascent to yield θ^i . At the end of the inner loop, gradients on a test set of interactions are computed. In the outer loop, the update to θ is a weighed aggregate of the test gradients from the inner loop. That is, the training step for the outer loop is based on the test step of the inner loop.

3. COMPLEMENTARY META-REINFORCEMENT LEARNING

3.1. Problem Formulation

The problem of the controller is thus: to exploit its past experiences with different processes, and sparse interactions under new process dynamics to quickly converge to a locally optimal policy. The proposed approach for adaptive control operates under the framework depicted in figure 1. The adaptation pipeline can either be preempted by fault detection, or happen periodically.



Figure 1

Algorithm 1: Model-agnostic meta-learning

Input: parameters θ_k , MDPs P , learning rates $\alpha_{in}, \alpha_{out}$, iterations K_{in}, K_{out}

```
1 begin
2   Set  $\theta'_1 \leftarrow \theta_k$ ;
3   for  $k_{out} = 1$  to  $K_{out}$  do
4     Sample MDPs  $p_i \sim P$ ;
5     for all  $p_i$  do
6       Set  $\theta_1^i \leftarrow \theta$ ;
7       for  $k_{in} = 1$  to  $K_{in}$  do
8         Sample training trajectories  $\mathcal{M}_i$  from  $p_i$ ;
9         Calculate gain function from  $\mathcal{M}_i$ ;
10        Update  $\theta_{k_{in}+1}^i \leftarrow \theta_{k_{in}}^i + \alpha_{in} \cdot \nabla_{\theta} G_i$ ;
11        Sample test trajectories from  $p_i$ ;
12        Calculate test gain  $G_i$  on sample;
13      Update  $\theta'_{k_{out}+1} \leftarrow \theta'_{k_{out}} + \alpha_{out} \cdot \sum_i \nabla_{\theta} G_i$ ;
14  return  $\theta_{K_{out}}$ 
```

The adaption step begins with a *fault*. The fault is abrupt, causing a discontinuous change in process dynamics. The MDP representing the system has changed. In the aftermath of a fault, a controller continues to interact with, and record, states, actions, and rewards in some memory \mathcal{M} . Once sufficient interactions t_{update} have been buffered, the controller attempts to learn from its memory. Once learning is complete, the controller consolidates the newly learned policy with its prior policies. Thus, when a new fault occurs, it is able to exploit its past experience and adapt faster.

The *learning* phase consists of two stages: the meta-update, followed by iterations of any choice of a gradient based learning algorithm. During the meta-update, the controller uses its consolidated prior experience to initialize new policy parameters. The controller can also generate a data-driven model of the system to supplement sample inefficiency of RL. After that, the parameters are iteratively updated by the RL algorithm through interactions with the actual system.

Consolidation of knowledge happens via maintaining a complement of prior policies $\mathcal{C} = \{\theta \mid \pi_{\theta}\}$. The set of policies is periodically pruned to ensure that they capture diverse behavior but are small enough to evaluate within time constraints.

3.2. Policy meta-update

Our approach mirrors MAML in that there is an outer update loop for the main policy parameters. It depends on the gradients of the test error on the inner loop. We diverge in our formulation of the inner loop. Whereas in MAML the inner loop samples alternative processes and learns parameters at that instant, we instead exploit the history of the controller's experience.

Prior to the meta-update, a memory \mathcal{M} of interactions under the new process is buffered. The meta-update step assumes

a complement $\mathcal{C} = \{\theta \mid \pi_{\theta}\}$ of prior policies trained on the system under different faults. This foregoes the need of sampling an altogether new set of processes for the meta-update. The complement of policies is then optionally trained for a few steps K_{in} to yield an updated set of meta-parameters. Finally, the test error on the meta-parameters is used to update the outer loop's policy parameters.

Evaluating policies from \mathcal{C} necessitates interactions with the changed system. This can be achieved by learning a data-driven model p_m of the process using \mathcal{M} . This introduces an additional computational load on the meta-update step. An alternative approach is to forego a model altogether and instead use importance sampling (2) on \mathcal{M} using each policy.

Algorithm 2: Complementary meta-RL

Input: parameters θ , memory \mathcal{M} , learning rates $\alpha_{in}, \alpha_{out}$, iterations K_{in}, K_{out} **Optional:** policy complement $\mathcal{C} = \{\emptyset\}$, process model $p_m = \emptyset$

```
1 begin
2   Update  $p_m$  from  $\mathcal{M}_0$ ;
3   Set  $\theta'_1 \leftarrow \theta_k$ ;
4   for  $k_{out} = 1$  to  $K_{out}$  do
5     Calculate gain function  $G_0$  from  $\mathcal{M}$ ;
6     Update baseline  $\theta_{k_{out}+1}^b \leftarrow \theta_{k_{out}}^b + \alpha_{out} \cdot \nabla_{\theta} G_0$ ;
7     for all  $\theta_i$  in  $\mathcal{C}$  do
8       Set  $\theta_1^i \leftarrow \theta$ ;
9       for  $k_{in} = 1$  to  $K_{in}$  do
10        if  $p_m$  provided then
11          Clear  $\mathcal{M}$ ;
12          Sample training trajectories into  $\mathcal{M}$ 
13          from  $p_m$  using  $\pi_{\theta_{k_{in}}}$ ;
14          Calculate gain function  $G_i$  from  $\mathcal{M}$ ;
15          Update  $\theta_{k_{in}+1}^i \leftarrow \theta_{k_{in}}^i + \alpha_{in} \cdot \nabla_{\theta} G_i$ ;
16          Sample test trajectories  $\mathcal{M}_i$  from  $p_i$ ;
17          Calculate test gain  $G_i$  on sample;
18        Update  $\theta'_{k_{out}+1} \leftarrow \theta_{k_{out}} + \alpha_{out} \cdot \sum_i \nabla_{\theta} G_i$ ;
19      Calculate  $J_{\pi_{\theta_{K_{out}}}}, J_{\pi_{\theta_{K_{out}}}^b}$  from  $p_m$ ;
20      if  $J_{\pi_{\theta_{K_{out}}}} < J_{\pi_{\theta_{K_{out}}}^b}$  then return  $\theta_{K_{out}}^b$ ;
21  return  $\theta_{K_{out}}$ 
```

3.3. Population of complement

The complement of policies should be populated with policies such that it maximally spans the parameter space. Policies should be different enough so that the meta-update has a greater likelihood of adapting to novel faults. The difference between policies is evaluated on the memory of interactions collected by the controller. Each policy in \mathcal{C} generates a probability of actions taken in \mathcal{M} . KL-divergence between the probabilities is used as a metric of difference. The total divergence of each policy from the rest of the complement becomes a score of a policy's uniqueness. Given a comple-

ment size $s \leftarrow |\mathcal{C}|$, the s most unique policies are returned as members of \mathcal{C} .

Algorithm 3: Populating complement of policies

Input: policy complement \mathcal{C} , library size s , memory \mathcal{M}

```

1 begin
2   Initialize divergence matrix  $D = [0]^{|\mathcal{C}| \times |\mathcal{C}|}$ ;
3   for  $\theta_1, \theta_2$  in  $Permute(\mathcal{C})$  do
4     Action probabilities  $p_1, p_2 = \pi_{\theta_1}(\mathcal{M}), \pi_{\theta_2}(\mathcal{M})$ ;
5     KL-Divergence  $d = \sum p_1 \cdot \log(p_2/p_1)$ ;
6      $D[\theta_1, \theta_2] = d$ ;
7   Sum each row of  $D$  for total divergence  $D_T^{|\mathcal{C}| \times 1}$ ;
8   Most divergent parameters  $\mathcal{C} \leftarrow \text{Sort}(\mathcal{C})$  by  $D_T$ ;
9   return First  $s$  parameters from  $\mathcal{C}$ 

```

4. EXPERIMENTS

The algorithm was evaluated on a simulation of a fuel transfer system of an aircraft. The system is defined in greater detail in (Ahmed et al., 2020). The objective is to maintain center of gravity, variance in fuel distribution, and closed valves to avoid unnecessary mass transfer. Faults can include increased valve resistances leading to low flow rates, and increased fuel consumption due engine faults.

A controller was first trained for 50,000 steps on the nominal system. At the beginning of a trial, a random fault occurred and the controller accumulated experience in memory \mathcal{M} . The controller then employed the meta-update step in algorithm 2 to initialize new policy parameters. Following that, the RL algorithm continued to learn on the new system. As a baseline, an RL controller was trained for $|\mathcal{C}| \times K_{in} \times K_{out}$ iterations, for the case where p_m was provided, followed by learning on the new system.




Figure 2. Simplified fuel system schematics. The controller manages valves and can observe fuel tank levels. Net outflow to engines via pumps is controlled independently. Pumps drain tanks innermost first.

First, the controller was tested on random faults with an

empty complement of policies. A complement of 3 policies under simulated faults on the system was generated. The controller was then tested on the system under random faults. The controller was allowed to adapt solely from buffered experiences after a fault, without learning a new environment model. Figure 3 shows performance with $\mathcal{C} = \{\}$. Episodic rewards start off lower than but comparable to the baseline. They quickly recover and match baseline throughout. Of note is the low variance in episode rewards compared to the baseline. Figure 4 shows performance with a complement of 3 policies. Both controllers have similar initial performance, but the meta-updated controller quickly pulls ahead and converges to an optimum.




Figure 3. Episodic rewards after abrupt fault when there is no complement of policies available to the controller.

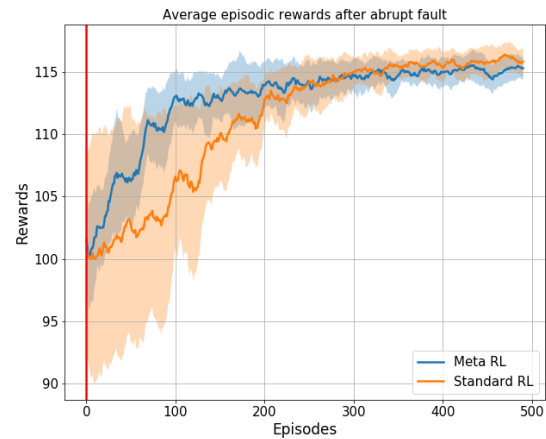


Figure 4. Episodic rewards after abrupt fault when there is a complement of 3 policies available to the controller.

5. RELATED WORK

Reinforcement learning has been explored for control systems. (Lewis, Vrabie, & Vamvoudakis, 2012) surveys RL approaches for feedback control. (Liu, Wang, & Zhang, 2016) attempts to speed up learning of neural network policies for controlling systems by manipulating the parameter update rule.

Approaches besides RL are prevalent in the field of FTC. (Jiang & Zhang, 2006; Zhang & Jiang, 2003) use performance degraded reference models to generate a library of the system under various conditions. Control is transferred to the policy learned for the most similar model in the library.

Meta RL for FTC is a nascent field. Recently, (Nagabandi et al., 2018) used model-based RL for quickly adapting control to changed system dynamics. They used MAML and a recurrent network as two approaches to develop a meta-update rule for the environment model parameters. Alternatively, (Sæmundsson, Hofmann, & Deisenroth, 2018) train a model to predict a latent representation of the environment. The latent variable is fed to the agent as a conditioning variable to represent changed dynamics. (Wang et al., 2016) use a recurrent neural network to train a controller on a population of related environments. The controller, being recurrent, has memory of this experience, and therefore learns an internal function to transition between environments as they change.

6. CONCLUSION

We have proposed a meta-RL algorithm, which exploits a controller's past experience under faults to initialize parameters for a new policy under a novel abrupt fault. The meta-update can optionally use a data-driven model to mitigate sample inefficiency, or it can fall back to using importance sampling on buffered experiences to evaluate the complement under current conditions. The newly derived parameters are added to the complement if they are divergent enough from the members of the set, thus ensuring a diverse library of behaviors for faster adaption to new faults.

MAML can be sensitive to choice of model architecture, task, and hyperparameters (Antoniou, Edwards, & Storkey, 2018). This merits further investigation on guarantees of convergence and optimality under faults. MAML can be further incorporated in our approach by using meta-learning to update the data-driven model itself. This should further reduce time taken to learn an updated model and the dependence on the size of the buffered data.

REFERENCES

- Ahmed, I., Quiñones-Grueiro, M., & Biswas, G. (2020). Fault-tolerant control of degrading systems with on-policy reinforcement learning. In *IFAC-PapersOnLine*.
- Antoniou, A., Edwards, H., & Storkey, A. (2018). How to train your maml. *arXiv preprint arXiv:1810.09502*.
- Baird, L. (1995). Residual algorithms: Reinforcement learning with function approximation. In *Machine learning proceedings 1995* (pp. 30–37). Elsevier.
- Bellman, R. (1966). Dynamic programming. *Science*, 153(3731), 34–37.
- Blanke, M., Kinnaert, M., Lunze, J., Staroswiecki, M., & Schröder, J. (2006). *Diagnosis and fault-tolerant control* (Vol. 2). Springer.
- Boyan, J. A., & Moore, A. W. (1995). Generalization in reinforcement learning: Safely approximating the value function. In *Advances in neural information processing systems* (pp. 369–376).
- Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th international conference on machine learning-volume 70* (pp. 1126–1135).
- Hongm, W., Tian-You, C., Jin-Liang, D., & Brown, M. (2009). Data driven fault diagnosis and fault tolerant control: some advances and possible new directions. *Acta Automatica Sinica*, 35(6), 739–747.
- Jiang, J., & Zhang, Y. (2006). Accepting performance degradation in fault-tolerant control system design. *IEEE transactions on control systems technology*, 14(2), 284–292.
- Lewis, F. L., Vrabie, D., & Vamvoudakis, K. G. (2012). Reinforcement learning and feedback control: Using natural decision methods to design optimal adaptive controllers. *IEEE Control Systems Magazine*, 32(6), 76–105.
- Liu, L., Wang, Z., & Zhang, H. (2016). Adaptive fault-tolerant tracking control for mimo discrete-time systems via reinforcement learning algorithm with less learning parameters. *IEEE Transactions on Automation Science and Engineering*, 14(1), 299–313.
- MacGregor, J., & Cinar, A. (2012). Monitoring, fault diagnosis, fault-tolerant control and optimization: Data driven methods. *Computers & Chemical Engineering*, 47, 111–120.
- Nagabandi, A., Clavera, I., Liu, S., Fearing, R. S., Abbeel, P., Levine, S., & Finn, C. (2018). Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. *arXiv preprint arXiv:1803.11347*.
- Sæmundsson, S., Hofmann, K., & Deisenroth, M. P. (2018). Meta reinforcement learning with latent variable gaussian processes. *arXiv preprint arXiv:1803.07551*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Sutton, R. S., McAllester, D. A., Singh, S. P., & Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances*

in neural information processing systems (pp. 1057–1063).

Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., ... Botvinick, M. (2016). Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.

Zhang, Y., & Jiang, J. (2003). Fault tolerant control system design with explicit consideration of performance degradation. *IEEE Transactions on Aerospace and Electronic Systems*, 39(3), 838–848.

APPENDIX

6.1. Hyperparameters

Unless otherwise specified, the following parameters in table 1 were used.

Table 1. Meta-update parameters


Parameter	Value
$ \mathcal{M} $	2000
α_{in}	0.001
α_{out}	0.001
K_{in}	2
K_{out}	4
s	3

Table 2 documents parameters used by our implementation of PPO algorithm.


Table 2. Parameters used by the PPO algorithm.

Parameter	Value
Optimizer	Adam
α	0.02
β	(0.9, 0.999)
Epochs	5
t_{update}	2000
Value network	(64, tanh, 64, tanh, linear, 1)
Action network	(64, tanh, 64, tanh, 6, linear, sigmoid)
γ	0.99
ϵ	0.2

6.2. Additional Figures




(a) $|\mathcal{C}| = 0, \alpha_{in} = 0.001, \alpha_{out} = 0.001, K_{in} = 4, K_{out} = 1, p_m = \emptyset$. Fault in tank 4, engine 2.




(b) $|\mathcal{C}| = 3, \alpha_{in} = 0.001, \alpha_{out} = 0.001, K_{in} = 4, K_{out} = 1, p_m = \emptyset$. Fault in tank 4, engine 2.

Figure 5. Even with a single meta-update step, $K_{out} = 1$, there is noticeable increase in performance.



(a) $|\mathcal{C}| = 0, \alpha_{in} = 0.001, \alpha_{out} = 0.01, K_{in} = 4, K_{out} = 2, p_m = \emptyset$. Fault in tank 4, engine 2.



(b) $|\mathcal{C}| = 3, \alpha_{in} = 0.001, \alpha_{out} = 0.01, K_{in} = 4, K_{out} = 2, p_m = \emptyset$. Fault in tank 4, engine 2.

Figure 6. Due to a higher outer learning rate α_{out} and iteration number K_{out} , the meta-update shows a larger change in performance. With the help of a full complement, the parameter updated is moderated in a direction such that performance variance remains low and shows a higher rate of change.