
Global Convergence of MAML for LQR

Igor Molybog
University of California at Berkeley
igormolybog@berkeley.edu

Javad Lavaei
University of California at Berkeley
lavaei@berkeley.edu

Abstract

The paper studies the performance of the Model-Agnostic Meta-Learning (MAML) algorithm as an optimization method. The goal is to determine the global convergence of MAML on sequential decision-making tasks possessing a common structure. We prove that the benign landscape of a single task leads to the global convergence of MAML in the single-task scenario and in the scenario of multiple structurally connected tasks. We also show that there is a two-task scenario that does not possess this global convergence property even for identical tasks. We analyze the landscape of the MAML objective on LQR tasks to determine what type of similarities in their structures enables the algorithm to converge to the globally optimal solution.

1 Introduction

Meta-learning, along with transfer learning, is a rapidly developing research area in machine learning which aims to design algorithms that gain computational advantages out of inherent similarities between learning problem instances, otherwise referred to as tasks. In this work, we study one of the most popular meta-learning algorithms, called *Model-Agnostic Meta-Learning (MAML)*, which has been developed by Finn et al. [2017]. We focus our attention on its application in the domain of reinforcement learning (RL), and aim to theoretically study the global convergence properties of the algorithm on sequential decision-making tasks.

For clarity of explanation, we investigate discrete stationary infinite-horizon decision problems in this paper, and note that the generalization of our results to non-stationary and finite-horizon cases is straightforward. A stationary discrete dynamical system is described as

$$s_{t+1} \sim T(s_t, a_t),$$

where T is a probability distribution over the next state $s_{t+1} \in \mathbb{R}^d$ given the current state $s_t \in \mathbb{R}^d$, and action $a_t \in \mathbb{R}^r$. The initial state s_0 is assumed to follow the distribution T_0 . The objective of the infinite-horizon problem is to find a control input a_t minimizing the total discounted cost (the negation of the reward)

$$\begin{aligned} \text{minimize} \quad & \mathbb{E}_{s_0 \sim T_0} \sum_{t=0}^{\infty} \gamma^t c(s_t, a_t) \\ \text{subject to} \quad & s_{t+1} \sim T(s_t, a_t) \quad t = 0, 1, \dots \end{aligned}$$

where $\gamma \in (0, 1]$ is the discount factor that is assumed to be 1 for the LQR problem.

1.1 Linear-Quadratic Regulator (LQR)

One of the important examples of decision-making problems is related to the control of linear dynamical systems with a quadratic objective, referred to as *linear-quadratic regulator (LQR)*. LQR and iLQR [Todorov and Li, 2005] are fundamental tools for model-based reinforcement learning.

At the same time, linear-quadratic systems enjoy a rich theoretical foundation with a large number of provable guarantees, which makes them the perfect benchmark for the mathematical analysis of novel learning algorithms.

We consider the following infinite-horizon exact LQR problem:

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{x_0} \left[\sum_{t=0}^{\infty} (s_t^\top Q s_t + a_t^\top R a_t) \right] \\ & \text{subject to} && s_{t+1} = A s_t + B a_t, \end{aligned} \tag{1}$$

Assuming that matrices A and B are such that the optimal cost is finite, it is a classic result that shows that the optimal control policy is deterministic and linear in the state, i.e.,

$$a_t = -W^* s_t$$

where $W^* \in \mathbb{R}^{r \times d}$ [Bertsekas, 2017]. Moreover, the matrix W^* can be found from the model parameters by solving the Algebraic Riccati Equation (ARE)

$$P = A^\top P A + Q - A^\top P B (B^\top P B + R)^{-1} B^\top P A,$$

and substituting the positive-definite root P into

$$W^* = -(B^\top P B + R)^{-1} B^\top P A.$$

This implies that in order to find the solution of LQR, it suffices to only search over deterministic policies of the form $a = -W s$ parameterized with a matrix $W \in \mathbb{R}^{r \times d}$.

In an effort to build a bridge between practical RL algorithms and the optimal control theory, Fazel et al. [2018] shows that W^* can be found by applying the policy gradient algorithm to a reformulated cost function. The LQR cost of a linear deterministic policy with respect to W can be defined as

$$C(W) := \mathbb{E}_{s_0 \sim T_0} \left[\sum_{t=0}^{\infty} (s_t^\top Q s_t + a_t^\top R a_t) \right]$$

where $a_t = -W s_t$ and $s_{t+1} = (A - BW) s_t$. This can be reformulated as

$$C(W) = \mathbb{E}_{x_0 \sim T_0} s_0^\top P_W s_0$$

where P_W is the solution of $P_W = Q + W^\top R W + (A - BW)^\top P_W (A - BW)$, from which its gradient with respect to the policy can be written explicitly:

$$\nabla C(W) = 2 \left((R + B^\top P_W B) W - B^\top P_W A \right) \Sigma_W$$

where

$$\Sigma_W = \mathbb{E}_{x_0 \sim T_0} \sum_{t=0}^{\infty} s_t s_t^\top.$$

We only consider the cost of stable policies, assuming it to be infinite for unstable ones.

1.2 Model-Agnostic Meta-Learning (MAML)

Given a set of tasks \mathcal{T} , each represented by an objective function \mathcal{L}_τ , and a probability distribution $\mathbb{P}_\mathcal{T}$ over the tasks, Finn et al. [2017] proposes an algorithm for finding an initialization of the policy gradient method that allows a fast adaptation to a task through just several gradient updates. In case the task consists in regression, classification, or clusterization, \mathcal{L}_τ is a risk or an empirical risk. In case of a reinforcement learning task, \mathcal{L}_τ is the cost (the negated return) of a policy.

If the space of considered policies is parameterized via $w \in \mathcal{W}$, then a single-shot MAML (which uses just one gradient update) aims to minimize the objective

$$\mathbb{E}_{\tau \in \mathcal{T}} f_\tau(w - \eta \nabla g_\tau(w)) \tag{2}$$

where f_τ and g_τ can be two different approximations of the objective function of the task τ . For example, if \mathcal{L}_τ is the risk of a learning problem, then f_τ can be the empirical risk conditioned on a large dataset, while g_τ is the empirical risk conditioned on a smaller dataset. The generality of this

Algorithm 1 Model-Agnostic Meta-Learning (MAML)

Require: $p(\mathcal{T})$: Probabilistic task generator

Require: η, β : Step size hyperparameters

- 1: Randomly initialize θ
 - 2: **while** not done **do**
 - 3: Sample batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$
 - 4: **for all** \mathcal{T}_i **do**
 - 5: Evaluate $\nabla g_{\mathcal{T}_i}(w)$
 - 6: Compute adapted parameters with gradient descent: $w'_i = w - \eta \nabla_w g_{\mathcal{T}_i}(w)$
 - 7: **end for**
 - 8: Update $w \leftarrow w - \beta \nabla_w \sum_{\mathcal{T}_i \sim p(\mathcal{T})} f_{\mathcal{T}_i}(w'_i)$
 - 9: **end while**
-

formulation will be used in Section 2, but for the study of LQR we will assume that the functions \mathcal{L}_τ , f_τ and g_τ all coincide and are equal to $C(W)$. When it is clear from the context which task is being discussed, we will omit the subscript.

Algorithm 1 is a basic version of MAML although other versions have been developed in the literature, e.g. FO-MAML, HF-MAML [Fallah et al., 2019], iMAML [Rajeswaran et al., 2019], Reptile [Nichol et al., 2018] and FTML [Finn et al., 2019]. We do not directly address the other formulations in the paper, but the conclusions of this work are applicable to them as well since they are created to minimize essentially the same objective function (2) and the convergence to at least a first-order stationary point has been proven for the majority of these variants.

As meta-learning seeks to improve learning performance by exploiting similarities between tasks, it is important to understand what type of similarities a particular meta-learning algorithm can take advantage of. A highly desirable feature of a meta-learning algorithm is an acceptable meta-test performance at least on the tasks it has been meta-trained on. In other words, if the set of tasks \mathcal{T} is finite and the meta-training procedure has access to all of them, then it should succeed on the meta-testing stage. For MAML, this translates into a requirement of successful minimization of the objective (2). Finn et al. [2019] shows that the global minimum is achieved in case all the functions $f_\tau = g_\tau$ are smooth and strongly convex. The LQR objective (1) is not convex in general. Tasks with non-convex objectives are substantially harder to analyze, and thus Fallah et al. [2019] only shows convergence of MAML to a first-order stationary point of (2) if f_τ are non-convex. We aim to study the global convergence properties of MAML applied to tasks with non-convex objectives since in meta-RL the objective functions of the tasks are likely non-convex, but may possess benign landscape.

The landscape of (1) has been studied by Fazel et al. [2018]. They note that there exist instances of LQR that are not convex, quasi-convex, or star-convex, which means that none of the existing results on global convergence of MAML can be applied even to such a basic decision problem as LQR. However, (1) possess benign landscape, and, as our study shows that this property can be transferred to (2).

This work studies global convergence of a popular practical algorithm (MAML) for sequential decision-making problems by utilizing optimal control theory and tools from non-convex mathematical optimization.

Notation

Given a set $\mathcal{Z} \subseteq \mathbb{R}^n$ and a point $\bar{z} \in \mathcal{Z}$, define the *open neighborhood* $\mathcal{B}_\delta(\bar{z}) = \{z \in \mathcal{Z} \mid \|z - \bar{z}\|_2 < \delta\}$. Given a function $\ell : \mathcal{Z} \rightarrow \mathbb{R}$, we call $\bar{z} \in \mathcal{Z}$ a *local minimizer* of the function ℓ if there exists $\delta > 0$ such that $\ell(z) \geq \ell(\bar{z})$ for all $z \in \mathcal{B}_\delta(\bar{z})$. The value $\ell(\bar{z})$ in this case is called a *local minimum*. A *global minimizer* of ℓ is a point \bar{z} such that $\ell(z) \geq \ell(\bar{z})$ for all $z \in \mathcal{Z}$. The *global minimum* is a value $\min_{z \in \mathcal{Z}} \ell(z)$ such that $\ell(z) \geq \min_{z \in \mathcal{Z}} \ell(z)$ for all $z \in \mathcal{Z}$. A *spurious local minimum* is a local minimum that is not a global minimum. Given a differentiable function $\ell : \mathcal{Z} \rightarrow \mathbb{R}$, a *first-order stationary point* $\bar{z} \in \mathcal{Z}$ is such that $\nabla \ell(\bar{z}) = 0$. Given a twice differentiable function ℓ , a first-order stationary point \bar{z} such that $\nabla \ell(\bar{z}) \succeq 0$ is called a *second-order stationary point*. Given a matrix




Figure 1: MAML objective (2) on a single LQR task

$A \in \mathbb{R}^{n \times m}$, we denote its transpose as A^\top and its operator norm in the space ℓ_2 as $\|A\|_2$. For a vector v from an ℓ_2 -space, its norm is denoted with $\|v\|_2$. Cardinality of the set \mathcal{T} is denoted with $|\mathcal{T}|$.

2 Main results

In this section, we study the MAML algorithm under four different scenarios. Firstly, we apply MAML to a single LQR task. Secondly, we consider several identical tasks. After that, we introduce a metric between tasks and extend the study to a number of similar tasks, and, finally, we study the convergence of MAML on a large number of distant LQR tasks.

We provide theoretical results for general multi-dimensional systems, while all of the presented examples and counter-examples are on one-dimensional LQR tasks since they are easy to visualize. Note that all of the examples are extendable to multi-dimensional systems as well. The details on the exact tasks used for the computations are provided in the Appendix.

2.1 Single task

We begin by analyzing MAML applied to a singleton task set \mathcal{T} . If MAML fails under this scenario, then there is little hope on its global convergence for the multiple-task scenario. For the single task, we rewrite the MAML objective (2) as

$$h(w) = f(w - \eta \nabla g(w))$$

where f is assumed to be a continuously differentiable function and g is assumed to be twice continuously differentiable. As noticed earlier, the existing results on global convergence of MAML are not applicable to LQR. Figure 1 demonstrates an example of the MAML objective (2) applied to a single LQR. It is non-convex and has three distinct strict local minima, but all of them are also global minima. This implies that Algorithm 1 would converge to its global minimizer from almost any initial point. The minimizer in the middle corresponds to W^* of the task, while the rightmost and leftmost minimizers are some points W such that $W - \eta \nabla C(W) = W^*$ and $\nabla C(W) \neq 0$. The minimizers on the both sides rely on the rapid adaptation during the meta-testing stage, which has been assumed by the creators of the algorithm.

This example gives rise to a hypothesis that the MAML objective for a single LQR possesses some sort of benign landscape and, more generally, that benign landscape of the cost of the underlying task results in benign landscape of the resulting MAML objective. Following Jozs et al. [2018], we formalize the notion of benign landscape by defining the global function:

Definition 1. A continuous function $\ell : \mathcal{Z} \rightarrow \mathbb{R}$ is called *global* if every local minimizer of ℓ is a global minimizer.

This property is also referred as having no spurious local minima and generalizes the notions of convexity, quasiconvexity, and star-convexity. As a relaxation of this property, we define ε -global function:

Definition 2. A continuous function $\ell : \mathcal{Z} \rightarrow \mathbb{R}$ is called ε -global if for every local minimizer \bar{z} of ℓ it holds that

$$\ell(\bar{z}) - \min_{z \in \mathcal{Z}} \ell(z) \leq \varepsilon$$

Being global is equivalent to being 0-global. This property is more likely to be satisfied than a perfect no-spurious property for cost functions coming from real-world applications, and it still implies that the landscape is benign. These two notions of benign landscape characterize when a coercive function is easy to optimize using local optimization methods based on greedy descent. The following theorem shows that the benign landscape of the cost of the underlying task does indeed lead to a benign landscape for the resulting MAML objective.

Theorem 1. *Let $f : \mathcal{W} \rightarrow \mathbb{R}$ be global and $g : \mathcal{W} \rightarrow \mathbb{R}$ be twice continuously differentiable with $\|\nabla^2 g(w)\| \leq M < \infty$ for all $w \in \mathcal{W}$. If Algorithm 1 with the parameter η chosen to be smaller than $\frac{1}{M}$ converges to a local minimizer $w^* \in \mathcal{W}$ of h , then w^* is a global minimizer of h .*

We will use the properties of the composition operator with open maps to the above theorem. First, we need to introduce some preliminary results. Consider an open subset \mathcal{W} of a finite-dimensional vector space and a continuous mapping $\mathcal{F} : \mathcal{W} \rightarrow \mathcal{Z}$ where $\mathcal{Z} = \text{range}(\mathcal{F})$.

Definition 3. *A mapping $\mathcal{F} : \mathcal{W} \rightarrow \mathcal{Z}$ with $\mathcal{Z} = \text{range}(\mathcal{F})$ is said to be locally open at w if for every $\epsilon > 0$, there exists $\delta > 0$ such that $\mathcal{B}_\delta(\mathcal{F}(w)) \subseteq \mathcal{F}(\mathcal{B}_\epsilon(w))$.*

Definition 4. *A mapping $\mathcal{F} : \mathcal{W} \rightarrow \mathcal{Z}$ with $\mathcal{Z} = \text{range}(\mathcal{F})$ is said to be open, if for every open set $U \in \mathcal{W}$, $\mathcal{F}(U)$ is (relatively) open in \mathcal{Z} .*

A mapping \mathcal{F} is open if and only if it is locally open at every point of its domain. The following lemma allows us to analyze the local landscape of the composition of a function with a locally open map.

Lemma 1 (Observation 1 of Nouiehed and Razaviyayn [2018]). *Suppose that continuous map $\mathcal{F} : \mathcal{W} \rightarrow \mathcal{Z}$ with $\mathcal{Z} = \text{range}(\mathcal{F})$ is locally open at \bar{w} . If for some $\ell : \mathcal{Z} \rightarrow \mathbb{R}$, the point \bar{w} is a local minimum of $\ell(\mathcal{F}(w))$ over \mathcal{W} , then $\bar{z} = \mathcal{F}(\bar{w})$ is a local minimum of $\ell(z)$ over \mathcal{Z} .*

The above observation applied to an ϵ -global function leads to the following corollary:

Lemma 2. *Let $\ell : \mathcal{Z} \rightarrow \mathbb{R}$ be an ϵ -global function, and consider a continuous map $\mathcal{F} : \mathcal{W} \rightarrow \mathcal{Z}$ with $\mathcal{Z} = \text{range}(\mathcal{F})$ is locally open at a local minimizer \bar{w} of $\ell \circ \mathcal{F}$. Then, it holds that*

$$\ell(\mathcal{F}(\bar{w})) - \min_{w \in \mathcal{W}} \ell(\mathcal{F}(w)) \leq \epsilon$$

Proof. By Lemma 1, $\bar{z} = \mathcal{F}(\bar{w})$ is a local minimizer of $\ell(z)$. Since ℓ is ϵ -global, $\ell(\bar{z}) - \min_{z \in \mathcal{Z}} \ell(z) \leq \epsilon$. Due to $\text{range}(\mathcal{F}) = \mathcal{Z}$, it follows that $\ell(\mathcal{F}(\bar{w})) - \min_{w \in \mathcal{W}} \ell(\mathcal{F}(w)) \leq \epsilon$. \square

The Inverse Function Theorem provides us with the following corollary, which states that mappings with a nonsingular Jacobian are open.

Lemma 3 (Theorem 9.25 in Rudin [1976]). *Let \mathcal{F} be a continuously differentiable mapping of an open set $\mathcal{W} \subset \mathbb{R}^n$ into \mathbb{R}^n . Suppose that the Jacobian $\nabla \mathcal{F}(x)$ is a nonsingular matrix for all $x \in \mathcal{W}$. Then, for every open subset V of \mathcal{W} the set $\mathcal{F}(V)$ is open. In other words, $\mathcal{F} : \mathcal{W} \rightarrow \text{range}(\mathcal{F})$ is an open map.*

Now, we have defined all the necessary machinery for the proof of Theorem 1.

Proof of Theorem 1. The mapping $\mathcal{F}(w) = w - \eta \nabla g(w)$ is continuously differentiable over \mathcal{W} . The Jacobian of this mapping is $\nabla \mathcal{F}(w) = \mathcal{I} - \eta \nabla^2 g(w)$, where \mathcal{I} is the identity matrix. By assumption, we have $\eta < \frac{1}{M}$, and consequently the following inequality holds for all $v \in \mathbb{R}^n \setminus \{0\}$:

$$v^\top [\mathcal{I} - \eta \nabla^2 g(w)] v \geq (1 - \eta \|\nabla^2 g(w)\|) v^\top v > 0$$

this means that $\nabla \mathcal{F}(w)$ is positive definite for all $w \in \mathcal{W}$, and hence is non-singular. By Lemma 3, \mathcal{F} is an open mapping.

The function f is global by the assumption, and by Lemma 1 the composition of a global function with an open map is global. Thus, h is global as a composition of f and \mathcal{F} , and consequently each of its local minimizers is a global minimizer. \square

The by-product of Theorem 1 for the LQR problem will be stated below.

Theorem 2. Let $w^* \in \mathbb{R}^{r \times d}$ be the limit point of the sequence produced by Algorithm 1 (MAML) with $\eta < \frac{1}{\|\nabla^2 C(w^*)\|_2}$ applied to a single LQR task, meaning that $h(W) = C(W - \eta \nabla C(W))$. Then, w^* is the global minimizer of h , which implies that $C(W - \eta \nabla C(W))$ is a global function.

Proof. It is shown by Fallah et al. [2019] that MAML converges to a first-order stationary point of a smooth nonconvex loss. Therefore, if the limit point w^* exists, it must be a stationary point of $C(W - \eta \nabla C(W))$. Consider the first-order stationarity condition for w^* :

$$0 = \nabla C(w^* - \eta \nabla C(w^*)) = [\mathcal{I} - \eta \nabla^2 C(w^*)]^\top \nabla C(W) \Big|_{W=w^* - \eta \nabla C(w^*)}$$

Similarly to the proof of Theorem 1, $\mathcal{I} - \eta \nabla^2 C(w^*)$ is a full-rank matrix, meaning that w^* is a first-order stationary point for the MAML objective if and only if $\nabla C(W) \Big|_{W=w^* - \eta \nabla C(w^*)} = 0$.

Theorem 7 of Fazel et al. [2018] states that the Gradient descent algorithm finds an ε -approximation of the global optimum of $C(W)$ in polynomial time for any initial point with a finite value. This directly implies that all first-order stationary points of $C(W)$ are the global minimizers of $C(W)$ because otherwise, we could initialize the Gradient descent algorithm at a stationary point and since it would converge to the point of initialization, it would lead to a contradiction. Being a first-order stationary point is a sufficient condition of local minimality, and hence guaranteeing $C(W)$ being a global function. By Theorem 1, the MAML objective is global for a sufficiently small η .

Since $w^* - \eta \nabla C(w^*)$ is a first-order stationary point of C , it is a global minimizer of C and since $\min_W C(W) \leq \min_W C(W - \eta \nabla C(W))$, the point w^* is the global minimizer of the MAML objective. \square

Theorem 1 has implications far beyond the study of LQR. It is applicable in case f is not a reinforcement learning objective but an objective of a regression or a classification task. It also has a straightforward extension to the multi-step MAML and other variants of the MAML algorithm. Theorem 1 can also be viewed as a practical guideline for proving global convergence post-factum. If the function f is global and one runs the Algorithm 1 with a parameter η that converges to a point w^* such that $\nabla h(w^*) = 0$, then one can check whether $\nabla^2 h(w^*) \succ 0$ and $\eta < \frac{1}{\|\nabla^2 g(w^*)\|_2}$ and if so, w^* is guaranteed to be a global minimizer of h . In case f is not global but its landscape has some benign properties, then the following generalization takes place:

Proposition 1. If f is ε -global for some $\varepsilon > 0$ and $w^* \in \mathcal{W}$ is such that $\nabla h(w^*) = 0$, $\nabla^2 h(w^*) \succ 0$ and $\eta < \frac{1}{\|\nabla^2 g(w^*)\|_2}$, then

$$h(\bar{w}) - \min_{w \in \mathcal{W}} h(w) \leq \varepsilon$$

Proof. The mapping $\mathcal{F}(w) = w - \eta \nabla g(w)$ is continuously differentiable over \mathcal{W} . The Jacobian of this mapping is $\nabla \mathcal{F}(w) = \mathcal{I} - \eta \nabla^2 g(w)$. By assumption, $\eta < \frac{1}{\|\nabla^2 g(w^*)\|_2}$ and consequently there exists δ such that $\eta < \frac{1}{\|\nabla^2 g(w)\|_2}$ for all $w \in \mathcal{B}_\delta(w^*)$. Similarly to the proof of Theorem 1, $\nabla \mathcal{F}(w)$ is positive definite for all $w \in \mathcal{B}_\delta(w^*)$, and therefore by Lemma 3, $\mathcal{F}|_{\mathcal{B}_\delta(w^*)}$ is an open mapping and hence locally open at w^* .

The second-order sufficient condition of local optimality is satisfied at w^* and by Lemma 2 it means that $h(\bar{w}) - \min_{w \in \mathcal{W}} h(w) \leq \varepsilon$. \square

In conclusion of this part of the study, we have shown that the benign landscape properties of MAML applied to a singleton \mathcal{T} are inherited from the benign landscape of the objective of the task to which MAML has been applied. In particular, this holds true for the LQR tasks.

2.2 Several identical tasks

Results obtained for the single-task scenario give rise to a hypothesis that the benign landscape of every individual task would help with the convergence of MAML in a multi-task setting as well, provided that all of the tasks have a similar structure. Thus, in this part we study multiple-task learning for which the MAML has originally been designed for. Starting with some tasks that are the most similar to each other, we consider LQR tasks that coincide up to multiplication of the cost by a positive scalar. This is the highest degree of similarity one can hope to obtain between




Figure 2: Two MAML objective functions (2) for identical LQR tasks (dashed lines) and the MAML objective for the uniform distribution among them (solid line). 2b demonstrates spurious local minima of the solid line.

sequential decision-making problems. More precisely, the landscape features of the cost function (local and global minimizers, maximizers and saddle points) are preserved under this transformation, and therefore we refer to tasks like this as identical henceforth.

We consider a finite set of LQR tasks with a uniform distribution among them, which allows us to reduce the MAML objective (2) to the form

$$\frac{1}{|\mathcal{T}|} \sum_{\tau \in \mathcal{T}} C_{\tau}(W - \eta \nabla C_{\tau}(W))$$

Figure 2 demonstrates an example of the MAML objective applied to two identical LQR tasks. Although the MAML objective of each individual task is global, only one global minimizer coincides among all of them. This is the minimizer that corresponds to W^* . The two minimizers on the sides are shifted and after interference they produce spurious local minima for the total objective function. From this we can conclude that MAML fails to catch this kind of similarity between the tasks. One practical lesson to learn from this picture is that keeping the cost (or the reward) function of the considered tasks normalized may improve the quality of the solution provided by MAML. Another lesson is that the design of meta-learning algorithms may benefit from considering the analysis of identical tasks as the simplest form of common structure in the tasks.

2.3 Several similar tasks

Moving forward, it is desirable to consider a different type of similarity between different tasks in the MAML setting. Therefore, we study the landscape of MAML on those tasks that are similar to each other in terms of the norm of the difference between parameters. For LQR, the parameters are the matrices A, B, Q and R . The result states that the benign landscape of the underlying tasks carries over to the MAML objective if the tasks are sufficiently close to each other. For LQR, it can be put in the form of the following informal proposition:

Proposition 2. *For every $\varepsilon > 0$ and $k \in \mathbb{N}$, there exists $\delta > 0$ such that the MAML objective (2) is ε -global for almost any set \mathcal{T} of k LQR tasks defined through the parameters $\{(A_i, B_i, Q_i, R_i)\}_{i=1}^k$ such that*

$$\|A_i - A_j\| + \|B_i - B_j\| + \|Q_i - Q_j\| + \|R_i - R_j\| \leq \delta \quad \forall i, j \in \{1 \dots k\}$$

Proof. The formal statement along with the proof of the result and additional discussions are provided in the Appendix.

Intuitively, as long as the dependence of the cost of a single task on the parameters is continuous, for a set of tasks that are close to each other, the multi-task landscape of MAML remains close to the landscape of MAML for just one of them. \square

Proposition 2 goes beyond the uniform distribution holds for any distribution on a finite number of tasks. An example of the MAML objective for five similar LQR tasks is demonstrated in Figure 3a. In the end, we conclude that MAML is able to capture the common structure among different tasks given through similar values of the parameters.

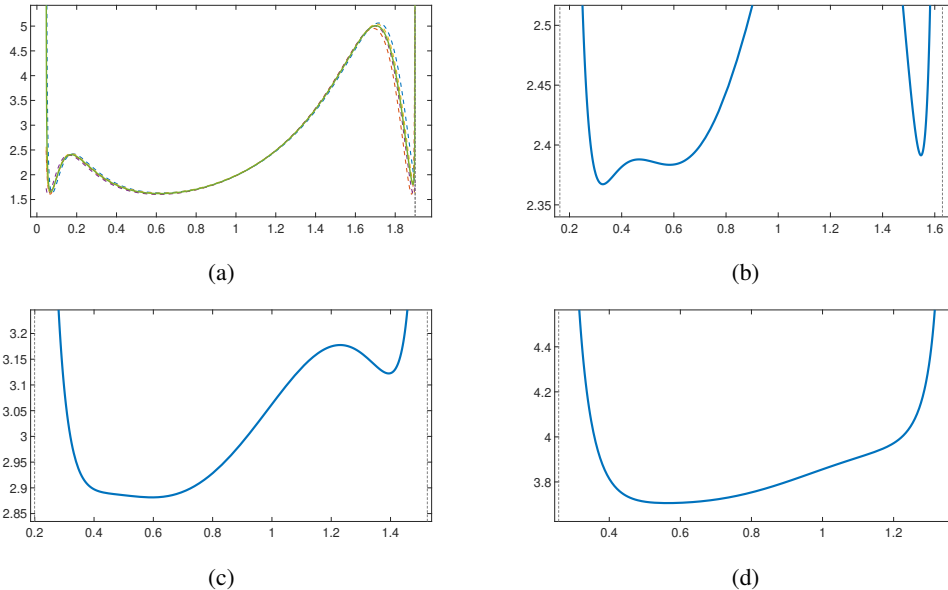


Figure 3: Five MAML objective functions (2) for LQR tasks with similar values of A , B , Q and R (dashed lines) and the MAML objective for the uniform distribution among them (solid line) (**plot 3a**). MAML objective (2) for the uniform distribution among two (**plot 3b**), five (**plot 3c**) and eleven (**plot 3d**) different LQR tasks that share the same dynamics A and B but have different cost matrices Q and R .

2.4 A number of tasks with common dynamics

In general, distant tasks may produce a wide variety of landscapes for the MAML objective. However, achieving benign landscape for several non-identical tasks with different norms is still possible. Figure 3 demonstrates the MAML objective for several LQR tasks that share the same dynamics but have different cost functions. Increasing the number of considered tasks improves the features of the landscape of the total MAML objective. In the eleven-task scenario, MAML seems to learn the mean of the optimal policies W^* for the tasks, which means that instead of learning to rapidly adapt, MAML would learn the average policy by effectively finding the minimum of $\frac{1}{|\mathcal{T}|} \sum_{\tau \in \mathcal{T}} C_{\tau}(W)$. However, for the two-task scenario, the global minimizer of the MAML objective appears to correspond to a policy W that has $\nabla C_{\tau}(W)$ far from zero for every considered τ , and thus the algorithm that converged to that point would learn to adapt to a task during the meta-testing phase.

In general, that the average of a large number of functions that each has a global minimizer in a small region of the domain would end up being almost global itself with the minimizer in the same region. Therefore, in practice, increasing the number of tasks for meta-training may improve the properties of the MAML landscape and assure convergence to a better local minimum.

3 Conclusion

In the paper, we studied the objective of the MAML algorithm 2 constructed for different numbers of tasks with some common structure. We have shown that the MAML objective for LQR defined by similar parameters possesses benign structure, and that this desirable property fails to hold for LQR tasks that coincide up to a multiplication of the objective by a positive value. The study of MAML on a larger subset of LQR tasks has led to the conclusion that a larger number of meta-training tasks is able to improve the final landscape of the meta-learning problem.

During this study of the MAML on LQR, a number of general results on benign landscapes were revealed, including that the benign landscape is transferable from an individual task to the MAML objective. Moreover, the notion of ε -global and its properties with respect to convex combination is described in the Appendix.

References

- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- Emanuel Todorov and Weiwei Li. A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *American Control Conference*, pages 300–306. IEEE, 2005.
- Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2017.
- Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *Proceedings of the 35th International Conference on Machine Learning*. PMLR, 2018.
- Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. *arXiv preprint arXiv:1908.10400*, 2019.
- Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Advances in Neural Information Processing Systems*, pages 113–124, 2019.
- Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- Chelsea Finn, Aravind Rajeswaran, Sham Kakade, and Sergey Levine. Online meta-learning. *arXiv preprint arXiv:1902.08438*, 2019.
- Cedric Jozs, Yi Ouyang, Richard Zhang, Javad Lavaei, and Somayeh Sojoudi. A theory on the absence of spurious solutions for nonconvex and nonsmooth optimization. In *Advances in neural information processing systems*, pages 2441–2449, 2018.
- Maher Nouiehed and Meisam Razaviyayn. Learning deep models: Critical points and local openness. *arXiv preprint arXiv:1803.02968*, 2018.
- Walter Rudin. *Principles of mathematical analysis (third edition)*. McGraw-Hill Inc., 1976.

Appendix

In the Appendix, we provide details on Proposition 2 and the experiments.

Lemma 4 (Theorem 9.28 of the book by Rudin [1976] (Generalized implicit function theorem)). *Let \mathcal{F} be a continuously differentiable mapping of an open set $\mathcal{E} \subset \mathbb{R}^{n+m}$ into \mathbb{R}^n such that $(\bar{x}, \bar{t}) = 0$ for some point $(\bar{x}, \bar{t}) \in \mathcal{E}$. Assume that $\nabla_x \mathcal{F}(\bar{x}, \bar{t})$ is a non-singular matrix. Then, there exist open sets $U \subset \mathbb{R}^{n+m}$ and $V \subset \mathbb{R}^m$, with $(\bar{x}, \bar{t}) \in U$ and $\bar{t} \in V$, having the following properties*

1. For all $t \in V$ there exists a unique $x = x(t)$ such that $(x, t) \in U$ and $\mathcal{F}(x, t) = 0$.
2. The function $x(t)$ is a continuously differentiable mapping of V into \mathbb{R}^n , $x(\bar{t}) = \bar{x}$ and $\nabla_t x(t) = -[\nabla_x \mathcal{F}(x(t), t)]^{-1} \nabla_t \mathcal{F}(x(t), t)$ for all $t \in V$.

Assumption 1. *Let $\ell : \mathcal{X} \times \mathbb{R}^m \rightarrow \mathbb{R}$ with a compact set $\mathcal{X} \subset \mathbb{R}^n$ be a twice continuously differentiable function with respect to $x \in \mathcal{X}$ with $\ell, \nabla_x \ell$ and $\nabla_{xx}^2 \ell$ being continuous with respect to $t \in \mathbb{R}^m$. Assume that $\ell(\cdot, t)$ has a finite number of first-order stationary points for all $t \in \mathbb{R}^m$ and that the Hessian is non-singular ($\det[\nabla_{xx}^2 \ell(x, t)] \neq 0$) for all $t \in \mathbb{R}^m$ and all $x \in \mathcal{X}$ such that $\nabla_t \ell(x, t) = 0$.*

Theorem 3. *If for some \bar{t} and $\varepsilon > 0$ the function $\ell(\cdot, \bar{t})$ satisfying Assumption 1 is ε -global, then for any $\varepsilon' > 0$ such that $\varepsilon' > \varepsilon$ there exists $\delta > 0$ such that the function $\ell(\cdot, t)$ is ε' -global for all $t \in \mathcal{B}_\delta(\bar{t})$.*

Proof. We prove the theorem in three steps. Step 1: $\ell(\cdot, t)$ has stationary points in the neighborhoods of the stationary points of $\ell(\cdot, \bar{t})$. Step 2: the types of the stationary points coincide. Step 3: There are no stationary points outside of the considered neighborhoods.

A finite number of stationary points means that all of them are isolated. Consider a stationary point \bar{x} of $\ell(\cdot, \bar{t})$. By Lemma 4, there exist $\psi > 0$ and $\phi > 0$ such that for all $t \in \mathcal{B}_\psi(\bar{t})$ there exists a unique $x(t) \in \mathcal{B}_\phi(\bar{x})$ with the property that $\nabla_x \ell(x(t), t) = 0$. This implies that for every function $\ell(\cdot, t)$ with $t \in \mathcal{B}_\psi(\bar{t})$ there is a unique stationary point over $\mathcal{B}_\phi(\bar{x})$.

Note that for any $v \in \mathbb{R}^n$ the function $h(x, t, v) = v^\top [\nabla_{xx}^2 \ell(x, t)] v$ is continuous in both x and t . By the assumption of the theorem, $\det[\nabla_{xx}^2 \ell(\bar{x}, \bar{t})] \neq 0$, and therefore $\nabla_{xx}^2 \ell(\bar{x}, \bar{t}) \succ 0$ if \bar{x} is a local minimum, $\nabla_{xx}^2 \ell(\bar{x}, \bar{t}) \prec 0$ if it is a local maximum, and $\nabla_{xx}^2 \ell(\bar{x}, \bar{t})$ indefinite if it is a saddle point. In each of these cases, we describe how to find values δ' and ϕ' such that ℓ has a bounded value of local minima over $\mathcal{B}_{\phi'}(\bar{x}) \times \mathcal{B}_{\delta'}(\bar{t})$.

Case 1: \bar{x} is a local minimum. The value of $h(\bar{x}, \bar{t}, v)$ is positive for all $v \in \mathbb{R}^n \setminus \{0\}$. By continuity, there exist $\psi' > 0$ and $\phi' > 0$ such that $\psi' < \psi$ and $\phi' < \phi$ and $h(x, t, v) > 0$ for all $x \in \mathcal{B}_{\phi'}(\bar{x})$, $t \in \mathcal{B}_{\psi'}(\bar{t})$ and $v \in \mathbb{R}^n \setminus \{0\}$. This way, $\nabla_{xx}^2 \ell(x(t), t) \succ 0$ and therefore $x(t)$ is a local minimum of $\ell(\cdot, t)$. By continuity of $\ell(x, t)$ with respect to x and t , there exists $\delta' > 0$ such that $\delta' < \psi'$ and $|\ell(x(t), t) - \ell(\bar{x}, \bar{t})| < \frac{\varepsilon' - \varepsilon}{2}$ for all $t \in \mathcal{B}_{\delta'}(\bar{t})$.

Case 2: \bar{x} is a local maximum. The value of $h(\bar{x}, \bar{t}, v)$ is negative for all $v \in \mathbb{R}^n \setminus \{0\}$. By continuity, there exist $\psi' > 0$ and $\phi' > 0$ such that $\psi' < \psi$ and $\phi' < \phi$ and $h(x, t, v) < 0$ for all $x \in \mathcal{B}_{\phi'}(\bar{x})$, $t \in \mathcal{B}_{\psi'}(\bar{t})$ and $v \in \mathbb{R}^n \setminus \{0\}$. This way, $\nabla_{xx}^2 \ell(x(t), t) \prec 0$ and therefore $x(t)$ is a local maximum of $\ell(\cdot, t)$. In this case, we assign to the point \bar{x} the value $\delta' = \psi'$.

Case 3: \bar{x} is a saddle point. There exist $v \in \mathbb{R}^n$ and $u \in \mathbb{R}^n$ such that $h(\bar{x}, \bar{t}, v) > 0$ and $h(\bar{x}, \bar{t}, u) < 0$. By continuity, there exist $\psi' > 0$ and $\phi' > 0$ such that $\psi' < \psi$ and $\phi' < \phi$ yet $h(x, t, v) > 0$ and $h(x, t, u) < 0$ for all $x \in \mathcal{B}_{\phi'}(\bar{x})$, $t \in \mathcal{B}_{\psi'}(\bar{t})$ and $v \in \mathbb{R}^n \setminus \{0\}$. This way, $\nabla_{xx}^2 \ell(x(t), t)$ is indefinite and therefore $x(t)$ is a saddle point of $\ell(\cdot, t)$. In this case, we assign to the point \bar{x} the value $\delta' = \psi'$.

As a result, having selected a single stationary point \bar{x} of $\ell(\cdot, \bar{t})$, we can find δ' and ϕ' such that all the stationary points of $\ell(\cdot, t)$ for $t \in \mathcal{B}_{\delta'}(\bar{t})$ are of the same type as \bar{x} , and in case they are local minimizers, they have a value that is not too different from $\ell(\bar{x}, \bar{t})$. One can repeat this argument for all the stationary points and therefore form sets of numbers $\{\delta'_i\}_{i=1}^N$ and $\{\phi'_i\}_{i=1}^N$, where i corresponds to the index of each of the N stationary points \bar{x}_i of $\ell(\cdot, \bar{t})$.

Consider the set $\mathcal{Y} = \mathcal{X} \setminus \cup_{i \in \{1 \dots N\}} \mathcal{B}_{\phi'_i}(\bar{x}_i)$. It is a compact set as a compact set minus an open set and $\|\nabla_x \ell(x, \bar{t})\| > 0$ for all $x \in \mathcal{Y}$. Since $\|\nabla_x \ell(x, \bar{t})\|$ is continuous in x over a compact, it is uniformly continuous and it reaches its lower bound, meaning that there exists $\xi > 0$ such that

$\|\nabla_x \ell(x, \bar{t})\| \geq \xi$ for all $x \in \mathcal{Y}$. By continuity of $\|\nabla_x \ell(x, t)\|$ with respect to t , there exists δ'' such that $\|\nabla_x \ell(x, t)\| > \frac{\xi}{2}$ over $\mathcal{Y} \times \mathcal{B}_{\delta''}(\bar{t})$ and therefore there are no stationary points of ℓ over $\mathcal{Y} \times \mathcal{B}_{\delta''}(\bar{t})$.

We select $\delta = \min[\{\delta'_i | i \in \{1 \dots N\}\} \cup \{\delta''\}]$ and observe that for all $t \in \mathcal{B}_\delta(\bar{t})$ the only local minimizers of $\ell(\cdot, t)$ are those close to the local minimizers of $\ell(\cdot, \bar{t})$. In *Case 1*, the corresponding δ' was selected such that, given a local minimizer $x(t)$ of $\ell(\cdot, t)$ that is neighboring a local minimizer \bar{x} of $\ell(\cdot, \bar{t})$ and a global minimizer $x'(t)$ of $\ell(\cdot, t)$ that is neighboring a local minimizer \bar{x}' of $\ell(\cdot, \bar{t})$, it holds that

$$\begin{aligned} \ell(x(t), t) - \min_x \ell(x, t) &= \ell(x(t), t) - \ell(x'(t), t) = \\ &= \ell(x(t), t) - \ell(\bar{x}, \bar{t}) + \ell(\bar{x}, \bar{t}) - \ell(\bar{x}', \bar{t}) + \ell(\bar{x}', \bar{t}) - \ell(x'(t), t) \leq \\ &= |\ell(x(t), t) - \ell(\bar{x}, \bar{t})| + |\ell(\bar{x}, \bar{t}) - \ell(\bar{x}', \bar{t})| + |\ell(\bar{x}', \bar{t}) - \ell(x'(t), t)| < \\ &= \frac{\varepsilon' - \varepsilon}{2} + \varepsilon + \frac{\varepsilon' - \varepsilon}{2} = \varepsilon' \end{aligned}$$

□

Lemma 5. *Given $\lambda_i > 0$ and $\sum_{i=1}^k \lambda_i = 1$, if for some \bar{t} and $\varepsilon > 0$ the function $\ell(\cdot, \bar{t})$ satisfying Assumption 1 is ε -global, then for any $\varepsilon' > 0$ such that $\varepsilon' > \varepsilon$ there exists $\delta > 0$ for which the convex combination $\lambda_1 \ell(\cdot, t_1) + \dots + \lambda_k \ell(\cdot, t_k)$ is ε' -global for all $t_1, \dots, t_k \in \mathcal{B}_\delta(\bar{t})$.*

Proof. We provide the proof for $k = 2$, but the argument holds true for other finite values of k . By Theorem 3, $\ell(\cdot, t_1)$ can be assumed to be ε'' -global for $\varepsilon' > \varepsilon'' > \varepsilon$. Therefore, without loss of generality, we assume that $t_1 = \bar{t}$, and then aim to prove that for a given $\lambda \in [0, 1]$ there exists $\delta > 0$ such that $\lambda \ell(x, \bar{t}) + (1 - \lambda) \ell(x, t)$ is ε' -global.

We introduce $\mathfrak{r}(x, t) = \lambda \ell(x, \bar{t}) + (1 - \lambda) \ell(x, t)$ and note that $\mathfrak{r}(x, \bar{t}) = \ell(x, \bar{t})$, which means that $\mathfrak{r}(\cdot, \bar{t})$ is ε -global and satisfies Assumption 1. Theorem 3 applied to \mathfrak{r} , yields that there exists $\delta > 0$ such that $\lambda \ell(x, \bar{t}) + (1 - \lambda) \ell(x, t)$ is ε' -global. □

Proposition 3. *Given $\lambda_i > 0$ and $\sum_{i=1}^k \lambda_i = 1$, if for all the values of t in a compact set $\mathcal{C} \subset \mathbb{R}^m$ and some $\varepsilon > 0$ the function $\ell(\cdot, t)$ satisfying Assumption 1 (restricted from \mathbb{R}^m to \mathcal{C}) is ε -global, then for any $\varepsilon' > 0$ such that $\varepsilon' > \varepsilon$ there exists $\delta > 0$ for which any convex combination $\lambda_1 \ell(\cdot, t_1) + \dots + \lambda_m \ell(\cdot, t_m)$ is ε' -global for all t_1, \dots, t_m such that $\|t_i - t_j\| < \delta$.*

Proof. To prove by contradiction, suppose that there exists $\varepsilon' > 0$ such that $\varepsilon' > \varepsilon$ and for all $\delta > 0$ there are $t_1(\delta), \dots, t_k(\delta)$ such that $\|t_i - t_j\| < \delta$ and $\lambda_1 \ell(\cdot, t_1) + \dots + \lambda_m \ell(\cdot, t_m)$ is not ε' -global. From the sequence $t_1(\frac{1}{l})$, one can extract a converging sub-sequence $t_1(\delta_l)$ since \mathcal{C} is compact. By Lemma 5, for the point $\bar{t} = \lim_{l \rightarrow \infty} t_1(\delta_l)$ there exists $\delta' > 0$ such that for all $t_1, \dots, t_k \in \mathcal{B}_{\delta'}(\bar{t})$ the convex combination $\lambda_1 \ell(\cdot, t_1) + \dots + \lambda_k \ell(\cdot, t_k)$ is ε' -global. Therefore, for the $t_1(\delta_l), \dots, t_k(\delta_l)$ such that $\delta_l < \frac{\delta'}{2k}$ and $\|t_1(\delta_l) - \bar{t}\| < \frac{\delta'}{2k}$, this convex combination is ε' -global, which is a contradiction. □

Proposition 4 (Formal statement of Proposition 2). *Consider the instances of LQR such that their parameters A, B, Q and R belong to a compact set. Assume that none of them produces a cost function $C(W)$ with a singular Hessian at a stationary point. For any $\varepsilon > 0$ and $k \in \mathbb{N}$, there exists $\delta > 0$ such that for any set of k considered instances of LQR defined through the parameters $\{(A_i, B_i, Q_i, R_i)\}_{i=1}^k$ with*

$$\|A_i - A_j\| + \|B_i - B_j\| + \|Q_i - Q_j\| + \|R_i - R_j\| \leq \delta \quad \forall i, j \in \{1 \dots m\},$$

the MAML objective (2) is ε -global.

Proof. Take $t = (A, B, Q, R)$. By Theorem 1, every function $C(W, t)$ satisfies Assumption 1. Therefore, $C(W, t)$ satisfies all the assumptions of Proposition 3 and the proof follows immediately. □

Given a matrix \bar{W} , the system of equations

$$\begin{cases} \nabla C(W) = 0 \\ \det(\nabla^2 C(W)) = 0 \end{cases} \quad (3)$$

is satisfied by the parameters A, B, Q and R that correspond to LQRs with \bar{W} as a stationary point with singular Hessian. All the LQR tasks that satisfy this system of equations are denoted by $HS(W)$. All the systems that have at least one stationary point with a singular Hessian are contained in $\cup_W HS(W)$. Now, notice that the union occurs over an $n \times n$ dimensional space, while $HS(W)$ is defined with a system of $(n \times n) + 1$ equations. This implies that $\cup_W HS(W)$ is still a low-dimensional manifold and thus almost no LQR has a stationary point with a singular Hessian. Thus, the compact domain that is mentioned in the assumption of Proposition 4 can be a large closed set with $\cup_W HS(W)$ excluded together with its small neighborhood. Thus, Proposition 2 is an informal restatement of Proposition 4.

Details on the experiments

For the numerical experiments, we computed the MAML objective explicitly using the formulas provided in Section 1.1.

To simplify the visualization, all of the examples and counterexamples in the paper were given for one-dimensional LQR systems. As a result, the parameters A, B, Q and R were scalar values, and so were the state and the action. The initial state s_0 was chosen to be deterministic. For reproducibility, Table 3 collects the parameters used to construct each of the examples.

Figure	LQR $(A, B, Q, R, s_0 s_0^\top)$	η
Figure 1	$(1, 1, 2, 2, 1)$	0.01
Figure 2	$(1, 1, 2, 2, 1), (1, 1, 0.1, 0.1, 1)$	0.01
Figure 3a	$(1.01, 1, 1, 1, 1), (1, 1.01, 1, 1, 1), (1, 1, 1.01, 1, 1), (0.99, 1, 1, 1, 1)$	0.01
Figure 3b	$(1, 1, 1, 2, 1), (1, 1, 2, 1, 1)$	0.1
Figure 3c	$(1, 1, 1, 1, 1), (1, 1, 1, 2, 1), (1, 1, 2, 1, 1), (1, 1, 2, 3, 1), (1, 1, 3, 2, 1)$	0.1
Figure 3d	$(1, 1, 1, 1, 1), (1, 1, 1, 2, 1), (1, 1, 2, 1, 1), (1, 1, 2, 3, 1), (1, 1, 3, 2, 1), (1, 1, 3, 1, 1), (1, 1, 1, 3, 1), (1, 1, 4, 1, 1), (1, 1, 1, 4, 1), (1, 1, 5, 3, 1), (1, 1, 3, 5, 1)$	0.1