

Information-theoretic generalization bounds for meta-learning and applications

Jose, Sharu Theresa; Simeone, Osvaldo

DOI:

[10.3390/e23010126](https://doi.org/10.3390/e23010126)

License:

Creative Commons: Attribution (CC BY)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Jose, ST & Simeone, O 2021, 'Information-theoretic generalization bounds for meta-learning and applications', *Entropy*, vol. 23, no. 1, 126. <https://doi.org/10.3390/e23010126>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Article

Information-Theoretic Generalization Bounds for Meta-Learning and Applications

Sharu Theresa Jose *  and Osvaldo Simeone 

Department of Engineering, King's College London, London WC2R 2LS, UK; osvaldo.simeone@kcl.ac.uk

* Correspondence: sharu.jose@kcl.ac.uk

Abstract: Meta-learning, or “learning to learn”, refers to techniques that infer an inductive bias from data corresponding to multiple related tasks with the goal of improving the sample efficiency for new, previously unobserved, tasks. A key performance measure for meta-learning is the meta-generalization gap, that is, the difference between the average loss measured on the meta-training data and on a new, randomly selected task. This paper presents novel information-theoretic upper bounds on the meta-generalization gap. Two broad classes of meta-learning algorithms are considered that use either separate within-task training and test sets, like model agnostic meta-learning (MAML), or joint within-task training and test sets, like reptile. Extending the existing work for conventional learning, an upper bound on the meta-generalization gap is derived for the former class that depends on the mutual information (MI) between the output of the meta-learning algorithm and its input meta-training data. For the latter, the derived bound includes an additional MI between the output of the per-task learning procedure and corresponding data set to capture within-task uncertainty. Tighter bounds are then developed for the two classes via novel individual task MI (ITMI) bounds. Applications of the derived bounds are finally discussed, including a broad class of noisy iterative algorithms for meta-learning.

Keywords: meta-learning; generalization bounds; mutual information; noisy iterative algorithms



Citation: Jose, S.T.; Simeone, O. Information-Theoretic Generalization Bounds for Meta-Learning and Applications. *Entropy* **2021**, *23*, 126. <https://doi.org/10.3390/e23010126>

Received: 17 November 2020
Accepted: 14 January 2021
Published: 19 January 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

As formalized by the “no free lunch theorem”, any effective learning procedure must be based on prior assumptions on the task of interest [1]. These include the selection of a model class and of the hyperparameters of a learning algorithm, such as weight initialization and learning rate. In conventional single-task learning, these assumptions, collectively known as *inductive bias*, are fixed *a priori* relying on domain knowledge or validation [1–3]. Fixing a suitable inductive bias can significantly reduce the sample complexity of the learning process, and is thus crucial to any learning procedure. The goal of meta-learning is to automatically infer the inductive bias, thereby learning to learn from past experiences via the observation of a number of related tasks, so as to speed up learning a new and unseen task [4–8].

In this work, we consider the meta-learning problem of inferring the hyperparameters of a learning algorithm. The learning algorithm (henceforth, called base-learning algorithm or base-learner) is defined as a stochastic mapping $P_{W|Z^m, u}$ from the input training set $Z^m = (Z_1, \dots, Z_m)$ of m samples to a model parameter $W \in \mathcal{W}$ for a fixed hyperparameter vector u . The meta-learning algorithm (or meta-learner) infers the hyperparameter vector u , which defines the inductive bias, by observing a finite number of related tasks.

For example, consider the well-studied algorithm of biased regularization for supervised learning [9,10]. Let us denote each data point $Z = (X, Y)$ as a tuple of input features $X \in \mathbb{R}^d$ and label $Y \in \mathbb{R}$. The loss function $l : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$ is given as the quadratic measure $l(w, z) = (\langle w, x \rangle - y)^2$ that quantifies the loss accrued by the inferred model parameter w on a data sample z . Corresponding to each per-task data set Z^m , the biased

regularization algorithm $P_{W|Z^m, u}$ is a Kronecker delta function centered at the minimizer of the following optimization problem

$$\frac{1}{m} \sum_{j=1}^m l(w, Z_j) + \frac{\lambda}{2} \|w - u\|^2, \quad (1)$$

which corresponds to an empirical risk minimization problem with a biased regularizer. Here, $\lambda > 0$ is a regularization constant that weighs the deviation of the model parameter w from a bias vector u . The bias vector u can be then thought of as a common “mean” among related tasks. In the context of meta-learning, the objective then is to infer the bias vector u by observing data sets from a number of similar related tasks. Different meta-learning algorithms have been developed for this problem [11,12].

In the meta-learning problem under study, we follow the standard setting of Baxter [13] and assume that the learning tasks belong to a task environment, which is defined by a probability distribution P_T on the space of learning tasks \mathcal{T} , and per-task data distributions $\{P_{Z|T=\tau}\}_{\tau \in \mathcal{T}}$. The data set Z^m for a task τ is then generated i.i.d. according to the distribution $P_{Z|T=\tau}$. The meta-learner observes the performance of the base-learner on the meta-training data from a finite number of meta-training tasks, which are sampled independently from the task environment, and infers the hyperparameter U such that it can learn a new task, drawn from the same task environment, from fewer data samples.

The quality of the inferred hyperparameter U is measured by the *meta-generalization loss*, $\mathcal{L}_g(U)$, which is the average loss incurred on the data set $Z^m \sim P_{Z^m|T}$ of a new, previously unseen task T sampled from the task distribution P_T . The notation will be formally introduced in Section 2.2. While the goal of meta-learning is to infer a hyperparameter U that minimizes the meta-generalization loss $\mathcal{L}_g(U)$, this is not computable, since the underlying task and data distributions are unknown. Instead, the meta-learner can evaluate an empirical estimate of the loss, $\mathcal{L}_t(U|Z_{1:N}^m)$, using the meta-training set $Z_{1:N}^m$ of data from N tasks, which is referred to as meta-training loss. The difference between the meta-generalization loss and the meta-training loss is the meta-generalization gap,

$$\Delta\mathcal{L}(U|Z_{1:N}^m) = \mathcal{L}_g(U) - \mathcal{L}_t(U|Z_{1:N}^m), \quad (2)$$

and measures how well the inferred hyperparameter U generalizes to a new, previously unseen task. In particular, if the meta-generalization gap is small, on average or with high probability, then the performance of the meta-learner on the meta-training set can be taken as a reliable estimate of the meta-generalization loss.

In this paper, we study information-theoretic upper bounds on the *average meta-generalization gap* $\mathbb{E}_{P_{Z_{1:N}^m} P_{U|Z_{1:N}^m}}[\Delta\mathcal{L}(U|Z_{1:N}^m)]$, where the average is with respect to the meta-training set $Z_{1:N}^m$ and the meta-learner defined by the stochastic kernel $P_{U|Z_{1:N}^m}$. Specifically, we extend the recent line of work initiated by Russo and Zhou [14], and Xu and Raginsky [15], which obtain mutual information (MI)-based bounds on the average generalization gap for conventional learning, to meta-learning. To the best of our knowledge, this is the first work that studies information-theoretic bounds for meta-learning.

The bounds on average meta-generalization gap, studied in this work, are distinct from the other well-known bounds on meta-generalization gap in literature. Broadly speaking, existing bounds on the meta-generalization gap can be grouped into two—high probability, probably approximately correct (PAC) bounds, and high probability PAC-Bayesian bounds. These upper bounds take the general form, $\mathbb{E}_{P_{U|Z_{1:N}^m}}[\Delta\mathcal{L}(U|Z_{1:N}^m)] \leq \epsilon$, that hold with probability at least $1 - \delta$, for $\delta \in (0, 1)$, over the meta-training set $Z_{1:N}^m$. In contrast, our work focuses on bounding $\mathbb{E}_{P_{Z_{1:N}^m}} \mathbb{E}_{P_{U|Z_{1:N}^m}}[\Delta\mathcal{L}(U|Z_{1:N}^m)]$ on average also over the meta-training set. Notable PAC bounds on meta-generalization gap include the bound of Baxter [13] obtained using the framework of Vapnik–Chervonenkis (VC) dimensions; and of Maurer [16], which employs the algorithmic stability [17,18] properties. In contrast, the PAC-Bayesian bounds also incorporate prior beliefs on the base-learner

and the meta-learner posteriors via an auxiliary data-independent prior distribution $Q_{W|U}$ and a hyper-prior distribution Q_U , respectively. Most notably, PAC-Bayesian bounds include that of Pentina and Lambert [19], the tighter bound of Amit and Meir [20], and most recently, the bounds of Rothfuss et al. [21]. While the high-probability bounds are agnostic to task and data distributions, our information-theoretic bounds depend explicitly on the task and per-task data distributions, on the loss function, and on the meta-training algorithm, in accordance to prior work on information-theoretic generalization bounds.

Another general property inherited from the information-theoretic approach adopted in this paper is that the bounds on the average meta-generalization gap under study are designed to hold for arbitrary base-learners and meta-learners. As such, they generally do not result in tighter bounds as compared to non-information theoretic generalization guarantees obtained for specific meta-learning problems, such as the ridge regression problem with meta-learned bias vector mentioned above [22]. In contrast, the general purpose of the bounds in this paper is to provide insights into the number of tasks, and the number of samples per task required to ensure that the training-based metrics are a good approximation to their population counterparts.

1.1. Main Contributions

The derivation of bounds on average meta-generalization gap differs from conventional learning owing to two levels of uncertainties—*environment-level* uncertainty and *within-task* uncertainty. While within-task uncertainty results from observing a finite number m of data samples per task as in conventional learning, environment-level uncertainty results from observing a finite number N of tasks from the task-environment. The relative importance of these two forms of uncertainty depend on the use made by the meta-learner of the meta-training data. In fact, depending on how the meta-training data are used by the meta-learner, we identify two main classes of meta-training algorithms—with separate within-task training and test sets, and joint within-task training and test sets. The former class includes the state-of-the-art meta-learning algorithms, such as model agnostic meta-learning (MAML) [23], that splits the training data corresponding to each task into training and test sets, with the latter reserved for within-task validation. In contrast, the second class of algorithms, such as reptile [24], use the entire per-task data both for training and testing. Our main contributions are as follows.

- In Theorem 1, we show that, for the case with separate within-task training and test sets, the average meta-generalization gap contains only the contribution of environment-level uncertainty. This is captured by a ratio of the mutual information (MI) between the output of the meta-learner U and the meta-training set $Z_{1:N}^m$, and the number of tasks N , as

$$\left| \mathbb{E}_{P_{Z_{1:N}^m} P_{U|Z_{1:N}^m}} [\Delta \mathcal{L}^{\text{sep}}(U|Z_{1:N}^m)] \right| \leq \sqrt{\frac{2\sigma^2}{N} I(U; Z_{1:N}^m)}, \quad (3)$$

where σ^2 is the sub-Gaussianity variance factor of the meta-loss function. This is a direct parallel of the MI-based bounds for single-task learning [25].

- In Theorem 3, we then show that, for the case with joint within-task training and test sets, the bound on the average meta-generalization gap also contains a contribution due to the within-task uncertainty via the ratio of the MI between the output of the base-learner and within-task training data and the per-task data sample size m . Specifically, we have the following bound

$$\left| \mathbb{E}_{P_{Z_{1:N}^m} P_{U|Z_{1:N}^m}} [\Delta \mathcal{L}^{\text{joint}}(U|Z_{1:N}^m)] \right| \leq \sqrt{\frac{2\sigma^2}{N} I(U; Z_{1:N}^m)} + \mathbb{E}_{P_T} \left[\sqrt{\frac{2\delta_T^2}{m} I(W; Z^m|T = \tau)} \right], \quad (4)$$

where δ_T^2 is the sub-Gaussianity variance factor of the loss function $l(w, z)$ for task T .

- In Theorems 2 and 4, we extend the individual sample MI (ISMI) bound of [26] to obtain novel individual task MI (ITMI)-based bounds on the meta-generalization gap for both separate and within-task training and test sets as

$$\left| \mathbb{E}_{P_{Z_{1:N}^m} P_{U|Z_{1:N}^m}} [\Delta \mathcal{L}^{\text{sep}}(U|Z_{1:N}^m)] \right| \leq \frac{1}{N} \sum_{i=1}^N \sqrt{2\sigma^2 I(U; Z_i^m)}, \quad (5)$$

and

$$\left| \mathbb{E}_{P_{Z_{1:N}^m} P_{U|Z_{1:N}^m}} [\Delta \mathcal{L}^{\text{joint}}(U|Z_{1:N}^m)] \right| \leq \frac{1}{N} \sum_{i=1}^N \sqrt{2\sigma^2 I(U; Z_i^m)} + \mathbb{E}_{P_T} \left[\frac{1}{m} \sum_{j=1}^m \sqrt{2\delta_T^2 I(W; Z_j|T = \tau)} \right]. \quad (6)$$

These bounds can be seen to be tighter than the MI-based bounds in (3) and (4), respectively.

- Finally, we study the applications of the derived bounds to two meta-learning problems. The first is a parameter estimation setup that involves one-shot meta-learning and base-learning procedures, for which a closed form expression for meta-generalization gap can be derived. The second application covers a broad range of noisy iterative meta-learning algorithms and is inspired by the work of Pensia et al. [27] for conventional learning.

1.2. Related Work

For conventional learning, there exists a rich literature on diverse frameworks for deriving upper bounds on the generalization gap, i.e., on the difference between generalization and training losses. Classical bounds from statistical learning theory quantify the generalization gap in terms of measures of complexity of the model class, most notably VC dimension [28] and Radmacher complexity [29]. This approach obtains high-probability, probably approximate correct (PAC) bounds on the generalization gap with respect to the training set. An alternate line of high-probability bounding techniques relies on the notion of algorithmic stability, which measures the sensitivity of the output of a learning algorithm to the replacement of individual samples from the training data set. The pioneering work [30] has been extended to include various notions of algorithmic stability [31–33]. As a notable example, a distributional notion of stability in terms of *differential privacy*, which quantifies the sensitivity of the distribution of algorithm's output to data set, has been studied in [34,35]. The high-probability PAC–Bayesian bounds rely on change of measure arguments and uses the Kullback–Leibler (KL) divergence between the algorithm and a data-independent prior to quantifying the algorithmic sensitivity [36–38].

Following the initial work of Russo and Zou [14], information-theoretic bounds on the average generalization gap for conventional learning have been widely investigated in recent years. Xu and Raginsky [25] showed that the MI between the output of the learning algorithm and its training data set yields an upper bound in expectation on the generalization gap. The bound has been shown to offer computable generalization guarantees for noisy iterative algorithms, including stochastic gradient Langevin dynamics (SGLD) in [27]. Various refinements of the MI-based bound have since been analyzed to obtain tighter bounds. In particular, the bounds in [39] employ chaining mutual information techniques to tighten the bounds in [25], while the bound in [26] depends on the MI between the output of the algorithm and an individual data sample. The MI between the output of the algorithm and a random subset of the data set appears in the bounds introduced in [40]. The total variation information between the joint distribution of the training data and algorithmic output and the product of marginals was shown in [41] to yield a bound on the generalization gap for any bounded loss function. Subsequent works in [42–44] consider other information-theoretic measures, such as maximum leakage and lautum information. Most recently, a conditional mutual information (CMI)-based approach has been proposed in [45] to develop generalization bounds.

1.3. Notation

Throughout this paper, upper case letters, e.g., X , denote random variables and lower case letters, e.g., x , their realizations. We use $\mathcal{P}(\cdot)$ to denote the set of all probability distributions on the argument set or vector space. For a discrete or continuous random variable X taking values in a set or vector space \mathcal{X} , $P_X \in \mathcal{P}(\mathcal{X})$ denotes its probability distribution, with $P_X(x)$ being the probability mass or density value at $x \in \mathcal{X}$. We denote as P_{X^n} the n -fold product distribution induced by P_X . The conditional distribution of a random variable X given random variable Y is similarly defined as $P_{X|Y}$, with $P_{X|Y}(x|y)$ representing the probability mass or density at $X = x$ conditioned on the event $Y = y$. We use $\|\cdot\|_2$ to denote the Euclidean norm of the argument vector, and I_d to denote a d -dimensional identity matrix. We define the Kronecker delta $\delta(x - x_0) = 1$ if $x = x_0$ and $\delta(x - x_0) = 0$ otherwise.

2. Problem Definition

In this section, we define the problem of interest by introducing the key definitions of generalization gap for conventional, or single-task, learning and for meta-learning.

2.1. Generalization Gap for Single-Task Learning

Consider first the conventional problem of learning a task $\tau \in \mathcal{T}$.

As illustrated in Figure 1, each task $\tau \in \mathcal{T}$ is associated with an underlying unknown data distribution, $P_{Z|T=\tau} \in \mathcal{P}(\mathcal{Z})$, defined in a subset or vector space \mathcal{Z} . Henceforth, we use $P_{Z|\tau}$ to denote $P_{Z|T=\tau}$ for notational convenience.




Figure 1. Directed graph representing the variables involved in the definition of generalization gap (11) for single-task learning.

The training procedure, which is referred to as the base-learner, has access to a training data set $Z^m = (Z_1, Z_2, \dots, Z_m) \sim P_{Z^m|\tau}$ of m independent and identically distributed (i.i.d.) samples drawn from distribution $P_{Z|\tau}$. The base-learner uses this data set to choose a model, or hypothesis, W from the model class \mathcal{W} by using a randomized training procedure defined by a conditional distribution $P_{W|Z^m, u}$ as

$$W \sim P_{W|Z^m, u}. \tag{7}$$

The conditional distribution $P_{W|Z^m, u}$ defines a stochastic mapping from the training data set Z^m to the model class \mathcal{W} . The training procedure (7) is parameterized by a vector $u \in \mathcal{U}$ of hyperparameters, which defines the inductive bias. As an example, the base-learner $P_{W|Z^m, u}$ may follow stochastic gradient descent (SGD) updates with hyperparameters u , including the learning rate and the initialization point.

The performance of a parameter vector $w \in \mathcal{W}$ on a data sample $z \in \mathcal{Z}$ is measured by a loss function $l : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}_+$. The generalization loss for a model parameter vector $w \in \mathcal{W}$ is the average

$$L_g(w|\tau) = \mathbb{E}_{P_{Z|\tau}}[l(w, Z)], \quad (8)$$

over a test example Z independently drawn from the data distribution $P_{Z|\tau}$. The subscript g is used to distinguish the generalization loss from the training loss defined below. The generalization loss cannot be computed by the learner, given that the data distribution $P_{Z|\tau}$ is unknown. Instead, the learner can evaluate the training loss on the data set Z^m , which is defined as the empirical average

$$L_t(w|Z^m) = \frac{1}{m} \sum_{i=1}^m l(w, Z_i). \quad (9)$$

The subscript t specifies that the loss is the empirical training loss.

The difference between generalization loss (8) and training loss (9) is known as generalization gap,

$$\Delta L(w|Z^m, \tau) = L_g(w|\tau) - L_t(w|Z^m), \quad (10)$$

and is a key metric that quantifies the level of uncertainty (This type of uncertainty is known as epistemic.) at the learner regarding the data distribution $P_{Z|\tau}$. The average generalization gap for the data distribution $P_{Z|\tau}$ and base-learner $P_{W|Z^m, u}$ is defined as

$$\mathbb{E}_{P_{Z^m, W|\tau, u}}[\Delta L(W|Z^m, \tau)], \quad (11)$$

where the expectation is taken with respect to the joint distribution $P_{Z^m, W|\tau, u} = P_{Z^m|\tau} P_{W|Z^m, u}$. A summary of the variables involved in the Definition of the generalization gap (11) can be found in Figure 1.

Intuitively, if the generalization gap is small, on average or with high probability, then the base-learner can take the performance (9) on the training set Z^m as a reliable measure of the generalization loss (8) of the trained model W . Furthermore, data-dependent bounds on the generalization gap can be used as regularization terms to avoid overfitting, yielding generalized Bayesian inference problems [46,47].

2.2. Generalization Gap for Meta-Learning

As discussed, in single-task learning, the inductive bias u , defining the hyperparameters of the training procedure, must be selected a priori, i.e., without having access to task-specific data. The inductive bias determines the training data set size m needed to ensure a small generalization loss (8), since, generally speaking, richer models require more data to be trained [1]. The sample complexity can be generally reduced if one selects a suitable inductive bias based on prior information. Such prior information is typically obtained from domain knowledge on the problem under study. In contrast, meta-learning aims at automatically inferring an effective inductive bias based on data from related tasks.

To elaborate, we follow the setting of [13], in which a meta-learner observes data from a number of tasks, known as meta-training tasks, from the same task environment. A task environment is defined by a task distribution $P_T \in \mathcal{P}(\mathcal{T})$, supported on the space \mathcal{T} of tasks, and by a per-task data distribution $P_{Z|\tau}$ for each task $\tau \in \mathcal{T}$. Using the meta-training data drawn from a randomly selected subset of tasks, the meta-learner infers a hyperparameter vector $u \in \mathcal{U}$ defining the inductive bias. This is done with the goal of ensuring that, using hyperparameter u , the base-learner $P_{W|Z^m, u}$ can efficiently learn on a new task, referred to as *meta-test task*, drawn independently from the same task distribution P_T .

To elaborate, the meta-training data consist of N data sets $Z_{1:N}^m = (Z_1^m, \dots, Z_N^m)$. Each i th data set is generated independently by first drawing a task $T_i \sim P_T$ from the task environment and then a task-specific training data set $Z_i^m \sim P_{Z^m|T_i}$. The meta-learner uses the meta-training data set $Z_{1:N}^m$ to infer a hyperparameter vector $u \in \mathcal{U}$. To this end, we consider a randomized meta-learner

$$U \sim P_{U|Z_{1:N}^m}, \tag{12}$$

where $P_{U|Z_{1:N}^m}$ is a stochastic mapping from the meta-training set $Z_{1:N}^m$ to the space \mathcal{U} of hyperparameters. We distinguish two different formulations of meta-learning that are often considered in the literature. In the first, the per-task data set Z^m is split into training, or support, and test, or query subsets [23,48]; while, in the second, the entire data set Z^m is used for both within-task training and testing [13,19,20].

2.2.1. Separate Within-Task Training and Test Sets

As seen in Figure 2, in this first approach to meta-learning, each meta-training sub data set Z_i^m is split into a training set and a test set as $Z_i^m = (Z_i^{m_{tr}}, Z_i^{m_{te}})$, where $Z_i^{m_{tr}}$ contains m_{tr} i.i.d. training examples and $Z_i^{m_{te}}$ contains m_{te} i.i.d. test examples, with $m = m_{tr} + m_{te}$. The within-task base-learner $P_{W|Z_i^{m_{tr}},u} \in \mathcal{P}(\mathcal{W})$ maps the per-task training subset $Z_i^{m_{tr}}$ to random model parameter $W_i \sim P_{W|Z_i^{m_{tr}},u}$ for a given hyperparameter $U = u$. The test subset is used to evaluate the empirical training loss of a model w for task T_i as

$$L_t(w|Z_i^{m_{te}}) = \frac{1}{m_{te}} \sum_{j=1}^{m_{te}} l(w, Z_{i,j}^{m_{te}}), \tag{13}$$

where $Z_{i,j}^{m_{te}}$ denote the j th example of the test subset $Z_i^{m_{te}}$. Furthermore, the overall empirical meta-training loss for a hyperparameter u is computed by summing up all meta-training tasks as

$$\mathcal{L}_t^{sep}(u|Z_{1:N}^m) = \frac{1}{N} \sum_{i=1}^N L_t^{sep}(u|Z_i^m), \tag{14}$$

where

$$L_t^{sep}(u|Z^m) = \mathbb{E}_{P_{W|Z^{m_{tr}},u}} [L_t(W|Z^{m_{te}})] \tag{15}$$

is the average per-task training loss over the base-learner.

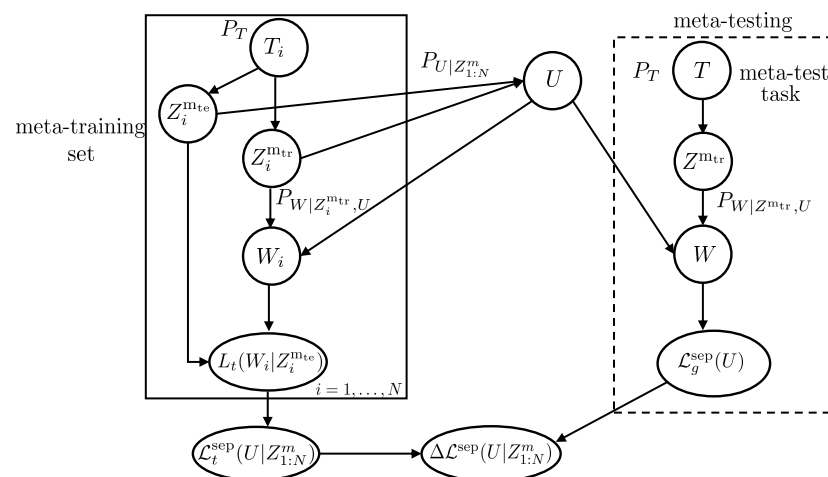


Figure 2. Directed graph representing the variables involved in the definition of meta-generalization gap (18) for separate within-task training and testing sets.

We emphasize that the meta-training loss (14) can be computed by the meta-learner and used as a criterion to select the meta-learning procedure (12), since it is obtained from the meta-training data $Z_{1:N}^m$. We also note that the rationale of splitting training and test sets is that the average training loss $L_t^{\text{sep}}(u|Z_i^m)$ is an unbiased estimate of the corresponding average generalization loss $\mathbb{E}_{P_{W|Z_i^m, \mu}}[L_g(W|T_i)]$.

The true goal of the meta-learner is to minimize the meta-generalization loss,

$$\mathcal{L}_g^{\text{sep}}(u) = \mathbb{E}_{P_{T, Z^{\text{mtr}}}} \mathbb{E}_{P_{W|Z^{\text{mtr}}, \mu}} [L_g(W|T)], \quad (16)$$

where $P_{T, Z^{\text{mtr}}} = P_T P_{Z^{\text{mtr}}|T}$ and $L_g(W|T)$ are as defined in (8). Unlike the meta-training loss (14), the meta-generalization loss is evaluated on a new, meta-test task T and on the corresponding training data Z^{mtr} . We distinguish the meta-generalization loss and meta-training loss by the subscripts g and t , respectively in (16) and (14). The difference between the meta-generalization loss (16) and the meta-training loss (14), known as the meta-generalization gap, is defined as

$$\Delta \mathcal{L}^{\text{sep}}(u|Z_{1:N}^m) = \mathcal{L}_g^{\text{sep}}(u) - \mathcal{L}_t^{\text{sep}}(u|Z_{1:N}^m). \quad (17)$$

The quantity of interest to us is the average meta-generalization gap, defined as

$$\mathbb{E}_{P_{Z_{1:N}^m, U}} [\Delta \mathcal{L}^{\text{sep}}(U|Z_{1:N}^m)], \quad (18)$$

where the expectation is with respect to the joint distribution $P_{Z_{1:N}^m, U} = P_{Z_{1:N}^m} P_{U|Z_{1:N}^m}$, of the meta-training set $Z_{1:N}^m$ and of the hyperparameter U . Note that $P_{Z_{1:N}^m}$ is the marginal of the joint distribution $\prod_{i=1}^N P_{T=T_i} P_{Z^M|T=T_i}$.

Intuitively, if the meta-generalization gap is small, on average or with high probability, the meta learner can take the performance (14) on the meta-training data as a reliable measure of the accuracy of the inferred hyperparameter vector in terms of the meta-generalization loss (16). Furthermore, data-dependant bounds on the meta-generalization gap can be used as regularization terms to avoid meta-overfitting. Meta-overfitting occurs when the meta-trained hyperparameter yields a small meta-training loss but a large meta-test loss, due to an excessive dependence on the meta-training set [13].

2.2.2. Joint Within-Task Training and Test Sets

In the second formulation of meta-learning, as illustrated in Figure 3, the entire data set Z_i^m is used for within-task training and testing. Accordingly, the meta-learner computes the meta-training loss

$$\mathcal{L}_t^{\text{joint}}(u|Z_{1:N}^m) = \frac{1}{N} \sum_{i=1}^N L_t^{\text{joint}}(u|Z_i^m), \quad (19)$$

where

$$L_t^{\text{joint}}(u|Z^m) = \mathbb{E}_{P_{W|Z^m, \mu}} [L_t(W|Z^m)] \quad (20)$$

is the average per-task training loss. Note here that in evaluating the meta-training loss in (19), the data set Z_i^m is used to infer model parameters W and to evaluate the per-task training loss. The expectation in (20) is taken over the output of the base-learner W given the hyperparameter vector u . As discussed, the meta-generalization loss for hyperparameter $u \in \mathcal{U}$ is computed by randomly selecting a novel task $T \sim P_T$ as

$$\mathcal{L}_g^{\text{joint}}(u) = \mathbb{E}_{P_{T, Z^m}} \mathbb{E}_{P_{W|Z^m, \mu}} [L_g(W|T)], \quad (21)$$

where $P_{T,Z^m} = P_T P_{Z^m|T}$ and $L_g(W|T)$ is as defined in (8). In a manner similar to (17), the *meta-generalization gap* for a task distribution P_T , data distribution $P_{Z^m|T}$, meta-learning algorithm $P_{U|Z_{1:N}^m}$, and base-learner $P_{W|Z^m,U}$ is defined as

$$\Delta \mathcal{L}^{\text{joint}}(u|Z_{1:N}^m) = \mathcal{L}_g^{\text{joint}}(u) - \mathcal{L}_t^{\text{joint}}(u|Z_{1:N}^m). \tag{22}$$




Figure 3. Directed graph representing the variables involved in the definition of meta-generalization gap (22) for joint within-task training and testing sets.

The average meta-generalization gap is then given as $\mathbb{E}_{P_{Z_{1:N}^m,U}}[\Delta \mathcal{L}^{\text{joint}}(U|Z_{1:N}^m)]$, where the expectation is taken over all meta-training sets and over the output of the meta-learner.

3. Information-Theoretic Generalization Bounds for Single-Task Learning

In this section, we review two information-theoretic bounds on the generalization gap (11) for conventional learning derived in [25,26]. The material covered in this section provides the necessary background for the analysis of the meta-generalization gap to be studied in the rest of the paper. Throughout this section, we fix a task $\tau \in \mathcal{T}$. Since the generalization and meta-generalization gaps measure the deviation of empirical-mean random variables representing training and meta-training losses from reference values, we will make use of tools and definitions from large-deviation theory (see, e.g., [49]). We discuss the key essential definitions below.

3.1. Preliminaries

To start, the cumulant generating function (CGF) of a random variable $X \sim P_X \in \mathcal{P}(\mathcal{X})$ is defined as $\Lambda_X(\lambda) = \log \mathbb{E}_{P_X}[e^{\lambda(X - \mathbb{E}_{P_X}[X])}]$. If it is well-defined, the CGF $\Lambda_X(\lambda)$ is convex and it satisfies the equalities $\Lambda_X(0) = \Lambda'_X(0) = 0$. A random variable X with finite mean, i.e., with $\mathbb{E}_{P_X}[X] < \infty$, is said to σ^2 -sub-Gaussian if its CGF is bounded as

$$\Lambda_X(\lambda) \leq \frac{\lambda^2 \sigma^2}{2}, \quad \text{for all } \lambda \in \mathbb{R}. \tag{23}$$

As a special case, if X is bounded in the interval $[a, b]$, i.e., if the inequality $0 < a \leq X \leq b < \infty$ holds for some constants a and b , then X is $(b - a)^2/4$ -sub-Gaussian.

3.2. Mutual Information (MI) Bound

We first present the mutual information (MI)-based upper bound obtained in [25]. Key to this result is the following Assumption.

Assumption 1. The loss function $l(w, Z)$ is δ_τ^2 -sub-Gaussian under $Z \sim P_{Z|\tau}$ for all model parameters $w \in \mathcal{W}$.

In particular, if the loss function is bounded, i.e., if the inequalities $-\infty < a \leq l(w, z) \leq b < \infty$ hold for all for $w \in \mathcal{W}$ and $z \in \mathcal{Z}$, Assumption 1 is satisfied with $\delta_\tau^2 = (b - a)^2/4$. The main result is as follows.

Lemma 1 ([25]). Under Assumption 1, the following bound on the generalization gap holds for any base-learner $W \sim P_{W|Z^m, \mu}$

$$|\mathbb{E}_{P_{Z^m, W|\tau, \mu}}[\Delta L(W|Z^m, \tau)]| \leq \sqrt{\frac{2\sigma^2}{m} I(W; Z^m)}. \tag{24}$$

The proof of Lemma 1 is based on a decoupling estimate Lemma, which is reported for completeness in Lemma A1. We also note that the result in Lemma 1 can be extended to account for loss function $l(w, Z)$ with bounded CGF [14].

The bound (24) on the generalization gap is in terms of the mutual information $I(W; Z^m)$, which quantifies the overall dependence between the base-learner output W and the input training data set Z^m . The mutual information in (24) is hence a measure of the sensitivity of the base-learner output to the data set. Using the terminology in [25], if $I(W; Z^m) \leq \epsilon$, the base-learner $P_{W|Z^m, \mu}$ is said to be $(\epsilon, P_{Z|\tau})$ -MI stable, in which case the bound in (24) evaluates to $\sqrt{2\delta_\tau^2 \epsilon / m}$. The relationship between generalization and stability of a training algorithm is well-established [1], and the result (24) amounts to a formulation of this link in information-theoretic terms.

The traditional notion of algorithmic stability measures how much the base-learner output changes with the replacement of an individual training sample [30,50]. In the next section, we review the bound in [26] that translates this per-sample stability concept within an information-theoretic framework.

3.3. Individual Sample MI (ISMI) Bound

The MI-based bound in Lemma 1 has the disadvantage of being vacuous, i.e., $I(W; Z^m) = \infty$, for deterministic base-learning algorithms $P_{W|Z^m, \mu}$ defined on continuous parameter space \mathcal{W} . An individual sample MI (ISMI)-based bound that address this shortcoming was introduced in [26]. The ISMI bound borrows the standard algorithmic stability notion of sensitivity of the base-learner output to the replacement of any individual training sample [17,18]. Accordingly, the resulting bound is in terms of the MI between the trained parameter W and each data point Z_i of the training data set Z^m . The bound, summarized in Lemma 2, applies under the following assumption.

Assumption 2. The loss function $l(w, z)$ satisfies either of the following two conditions:

- (a) Assumption 1, or
- (b) $l(W, Z)$ is a δ_τ^2 -sub-Gaussian random variable when $(W, Z) \sim P_{W|u, \tau} P_{Z|\tau}$, where $P_{W|u, \tau} \in \mathcal{P}(\mathcal{W})$ is the marginal of the joint distribution $P_{W|Z^m, \mu} P_{Z^m|\tau}$.

We note that, in general, Assumption 1 does not imply Assumption 2(b) (see ([40], Appendix C)), and vice versa (see [26]). There are, however, loss functions $l(w, z)$ and relevant distributions for which both the assumptions hold, including the case of loss functions $l(\cdot, \cdot)$ which takes values in a bounded interval $[a, b]$.

Lemma 2 ([26]). Under Assumption 2, the following bound on the average generalization gap holds for any base-learner $P_{W|Z^m, \mu}$

$$|\mathbb{E}_{P_{Z^m, W|\tau, \mu}}[\Delta L(W|Z^m, \tau)]| \leq \frac{1}{m} \sum_{i=1}^m \sqrt{2\sigma^2 I(W; Z_i)}. \tag{25}$$

For a loss function satisfying Assumption 1, the ISMI bound (25) is tighter than (24), i.e.,

$$\frac{1}{m} \sum_{i=1}^m \sqrt{2\delta_\tau^2 I(W; Z_i)} \leq \sqrt{\frac{2\delta_\tau^2}{m} I(W; Z^m)}. \tag{26}$$

The inequality in (26) follows from the chain rule of mutual information and Jensen’s inequality [26].

4. Information-Theoretic Generalization Bounds for Meta-Learning

In this section, we first derive novel MI-based bounds on the meta-generalization gap with separate within-task training and test sets, as introduced in Section 4.1, and then we consider joint within-task training and test sets, as described in Section 4.2.

4.1. Bounds on Meta-Generalization Gap with Separate Within-Task Training and Test Sets

In this section, we present two novel MI-based bounds on the meta-generalization gap (18) for the setup with separate within-task training and testing sets. The first is an MI-based bound, which is akin to Lemma 1, and the second is an individual task MI (ITMI) bound, which resembles Lemma 2 for conventional learning.

4.1.1. MI-Based Bound

In order to derive the MI-based bound, we make the following assumption on $L_t^{\text{sep}}(u|Z^m)$ in (15). Throughout, we use P_{Z^m} to denote the marginal of the joint distribution $P_{T,Z^m} = P_T P_{Z^m|T}$.

Assumption 3. For all $u \in \mathcal{U}$, the average per-task training loss $L_t^{\text{sep}}(u|Z^m)$ is σ^2 -sub-Gaussian under $Z^m \sim P_{Z^m}$.

Distinct from the assumptions in Section 3 on loss function $l(w, z)$, we note that Assumption 3 is on the average per-task training loss $L_t^{\text{sep}}(u|Z^m)$. This is because the loss function $l(w, z)$ satisfying Assumption 1 do not in general guarantee the sub-Gaussianity of $L_t^{\text{sep}}(u|Z^m)$ with respect to $Z^m \sim P_{Z^m}$. However, if the loss function is bounded, Assumption 3 can be easily verified to hold, as given in the following lemma.

Lemma 3. If the loss function $l(\cdot, \cdot)$ is $[a, b]$ -bounded, then $L_t^{\text{sep}}(\cdot|Z^m)$ is also $[a, b]$ bounded for all $Z^m \in \mathcal{Z}^m$. Consequently, $L_t^{\text{sep}}(u|Z^m)$ is $(b - a)^2/4$ -sub-Gaussian under $Z^m \sim P_{Z^m}$ for all $u \in \mathcal{U}$.

Under Assumption 3, the following theorem presents an upper bound on the meta-generalization gap (18).

Theorem 1. Let Assumption 3 hold for the base-learner $P_{W|Z^m_{\text{tr}}, \mathcal{U}}$. Then, for any meta learner $P_{U|Z^m_{1:N}}$ such that the inequality $I(U; Z^m_{1:N}) < \infty$ holds, we have the following bound on the average meta-generalization gap

$$\left| \mathbb{E}_{P_{Z^m_{1:N}, \mathcal{U}}} [\Delta \mathcal{L}^{\text{sep}}(U|Z^m_{1:N})] \right| \leq \sqrt{\frac{2\sigma^2}{N} I(U; Z^m_{1:N})}. \tag{27}$$

Proof. See Appendix B. □

The technical lemmas required for the proof of Theorem 1 and the theorems that follow are included in Appendix A.

In order to prove Theorem 1, one needs to overcome an additional challenge as compared to the derivation of bounds for learning reviewed in Section 3. In fact, the meta-generalization gap is caused by two distinct sources of uncertainty: (a) environment-level

uncertainty due to a finite number N of observed tasks, and (b) within-task uncertainty resulting from the finite number m of per-task data samples. Our proof approach involves applying the single-task MI-based bound in Lemma 1 to bound the effect of both sources of uncertainties.

Towards this, we start by introducing the average training loss for the randomly selected meta-test task as

$$\mathcal{L}_{g,t}^{\text{sep}}(u) = \mathbb{E}_{P_{T,Z^m}} [L_t^{\text{sep}}(u|Z^m)]. \tag{28}$$

The subscript g, t denotes that the loss is generalization (g) with expectation over P_{T,Z^m} at the environment level, and training (t) at the task level with $L_t^{\text{sep}}(u|Z^m)$. Note that this differs from the meta-test loss $\mathcal{L}_g^{\text{sep}}(u)$ in (16) in that the per-task loss is evaluated in (28) on the training set. With this definition, the meta-generalization gap can be decomposed as

$$\begin{aligned} & \mathbb{E}_{P_{Z_{1:N}^m, U}} [\Delta \mathcal{L}^{\text{sep}}(U|Z_{1:N}^m)] \\ &= \mathbb{E}_{P_{Z_{1:N}^m, U}} \left[(\mathcal{L}_g^{\text{sep}}(U) - \mathcal{L}_{g,t}^{\text{sep}}(U)) + (\mathcal{L}_{g,t}^{\text{sep}}(U) - \mathcal{L}_t^{\text{sep}}(U|Z_{1:N}^m)) \right]. \end{aligned} \tag{29}$$

In (29), the second difference $\mathcal{L}_{g,t}^{\text{sep}}(U) - \mathcal{L}_t^{\text{sep}}(U|Z_{1:N}^m)$ corresponds to the environment-level uncertainty and arises from the observation of a finite number N of tasks. In fact, as N increases, the meta-training loss $\mathcal{L}_t^{\text{sep}}(u|Z_{1:N}^m)$ almost surely tends to $\mathcal{L}_{g,t}^{\text{sep}}(u)$ by the law of large numbers. However, the average $\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{\text{sep}}(U) - \mathcal{L}_t^{\text{sep}}(U|Z_{1:N}^m)]$ is not equal to zero in general for finite values of N . The within-task generalization gap is instead measured by the difference $\mathcal{L}_g^{\text{sep}}(u) - \mathcal{L}_{g,t}^{\text{sep}}(u)$. In the setup under study with separate within-task training and test sets, this term equals zero, since, as we discussed, the average empirical loss $L_t^{\text{sep}}(u|Z_i^m)$ is an unbiased estimate of the corresponding average test loss $\mathbb{E}_{P_{W|Z_i^{\text{mtr}}, u}} [L_g(W|T_i)]$ (cf. (28)). This is no longer true for joint within-task training and test sets, as we discuss in Section 4.2.

The decomposition approach adopted here follows the main steps of the bounding techniques introduced in ([16], Equation (6)). In contrast, the PAC-Bayesian bounds in [20,21] rely on a different decomposition of the meta-generalization gap. The environment and within-task generalization gaps are then separately bounded in high probability, and are combined via union bound to obtain the required PAC-Bayesian bounds.

The bound (27) relates the meta-generalization gap to the information-theoretic stability of the meta-training procedure. As first introduced here, this stability is measured by the MI $I(U; Z_{1:N}^m)$ between the hyperparameter U and the meta-training data set $Z_{1:N}^m$, in a manner similar to the MI-based bounds in Lemma 1 for conventional learning. Importantly, as we will discuss in Section 4.2, this direct parallel between learning and meta-learning no longer applies with joint within-task training and test data sets.

4.1.2. ITMI Bound

We now present the ITMI bound, which holds under the following assumption.

Assumption 4. *Either of the following assumptions on the average per-task training loss, $L_t^{\text{sep}}(u|Z^m)$ holds:*

- (a) $L_t^{\text{sep}}(u|Z^m)$ satisfies Assumption 3, or
- (b) $L_t^{\text{sep}}(U|Z^m)$ is σ^2 -sub-Gaussian under $(U, Z^m) \sim P_U P_{Z^m}$, where P_U is the marginal of the joint distribution $P_{Z_{1:N}^m, U}$ and P_{Z^m} is the marginal of the joint distribution P_{T, Z^m} .

Assumption 4 can be seen to be implied by the sufficient conditions in Lemma 3.

Theorem 2. Let Assumption 4 hold for the base-learner $P_{W|Z^{m_{tr}},U}$. Then, for any meta learner $P_{U|Z^{m_{1:N}}}$, the following bound on the meta-generalization gap (18) holds

$$\left| \mathbb{E}_{P_{Z^{m_{1:N}},U}} [\Delta \mathcal{L}^{\text{sep}}(U|Z^{m_{1:N}})] \right| \leq \frac{1}{N} \sum_{i=1}^N \sqrt{2\sigma^2 I(U; Z_i^m)}. \tag{30}$$

where the MI $I(U; Z_i^m)$ is computed with respect to the joint distribution $P_{Z_i^m,U}$ obtained by marginalizing the probability distribution $P_{Z^{m_{1:N}},U}$.

Proof. See Appendix B. \square

As can be seen from (30), the ITMI bound on the meta-generalization gap is in terms of the MI $I(U; Z_i^m)$ between the output U of the meta learner and each per-task data set Z_i^m . This, in turn, quantifies the sensitivity of the meta learner output to the replacement of a single per-task data set. Moreover, under Assumption 3, the ITMI bound (30) yields a tighter bound than the MI-based bound (27). This can be seen from the following sequence of relations

$$\sqrt{\frac{1}{N} I(U; Z^{m_{1:N}})} = \sqrt{\frac{1}{N} \sum_{i=1}^N I(U; Z_i^m | Z_{(i-1)}^m)} \tag{31a}$$

$$\stackrel{(a)}{\geq} \sqrt{\frac{1}{N} \sum_{i=1}^N I(U; Z_i^m)} \tag{31b}$$

$$\stackrel{(b)}{\geq} \frac{1}{N} \sum_{i=1}^N \sqrt{I(U; Z_i^m)}, \tag{31c}$$

where $Z_{(i-1)}^m = (Z_1^m, \dots, Z_{i-1}^m)$; (a) follows, since Z_i^m is independent of $Z_{(i-1)}^m$; and (b) follows from Jensen’s inequality.

4.2. Bounds on Generalization Gap with Joint Within-Task Training and Test Sets

We now derive MI and ITMI-based bounds on the meta-generalization gap in (22) for the case with joint within-task training and test sets. As we will see, the key difference with respect to the case with separate within-task training and test sets is that the uncertainty due to finite number of per-task samples, measured by the second term in the decomposition (29), contributes in a non-negligible way to the meta-generalization gap. Since there is no split into separate within-task training and test sets, the average training loss with respect to the learning algorithm is given by $L_t^{\text{joint}}(u|Z^m)$ in (20).

4.2.1. MI-Based Bound

In order to derive the MI-based bound, we make the following assumptions.

Assumption 5. We consider the following assumptions.

- (a) For each task $\tau \in \mathcal{T}$, the loss function $l(w, Z)$ satisfies Assumption 1, and
- (b) The average per-task training loss $L_t^{\text{joint}}(u|Z^m)$ in (20) is σ^2 -sub-Gaussian for all $u \in \mathcal{U}$ when $Z^m \sim P_{Z^m}$.

An easily verifiable sufficient condition for the above assumption to hold is the boundedness of loss function $l(w, z)$, which follows in a manner similar to Lemma 3.

Theorem 3. Let Assumption 5 hold for a base-learner $W \sim P_{W|Z^m, U}$. Then, for any meta learner $P_{U|Z_{1:N}^m}$, we have the following bound on the meta-generalization gap (22)

$$\left| \mathbb{E}_{P_{Z_{1:N}^m, U}} [\Delta \mathcal{L}^{\text{joint}}(U|Z_{1:N}^m)] \right| \leq \sqrt{\frac{2\sigma^2}{N} I(U; Z_{1:N}^m)} + \mathbb{E}_{P_T} \left[\sqrt{\frac{2\delta_T^2}{m} I(W; Z^m|T = \tau)} \right]. \quad (32)$$

where the MI $I(W; Z^m|T = \tau)$ is evaluated with respect to the distribution $P_{Z^m, W|T=\tau}$ obtained by marginalizing the joint distribution $P_{W|Z^m, U} P_{Z_{1:N}^m, U} P_{Z^m|T=\tau}$.

Proof. See Appendix C. \square

With joint within-task training and test sets, the bound (32) on the meta-generalization gap contains the contributions of two mutual informations. The first, $I(U; Z_{1:N}^m)$, quantifies the sensitivity of the meta learner output U to the meta-training data set $Z_{1:N}^m$. This term also appeared in the bound (27) with separate within-task training and test sets. Decomposing the meta-generalization gap in a manner analogous to (29), it corresponds to a bound on the average of the second difference. The second contribution, $I(W; Z^m|T = \tau)$, quantifies the sensitivity of the output of the base-learner $P_{W|Z^m, U}$ to the data set Z^m of the meta-test task T , when the hyperparameter is randomly selected by the meta-learner $P_{U|Z_{1:N}^m}$ using the meta-training set $Z_{1:N}^m$. This second term is in line with the single-task generalization gap bounds (24), and it bounds the corresponding first difference in the decomposition (29).

We finally note that the dependence of the bound in (32) on the number of tasks N and per-task samples m is of the order $1/\sqrt{N} + 1/\sqrt{m}$. Meta-generalization bounds with similar dependence have been derived in [20] using PAC-Bayesian arguments. The bounds on excess risk for representation learning also follow a similar order of dependence on N and m (c.f [51], [Thm. 2]).

4.2.2. ITMI Bound on (22)

For deriving the ITMI bound on the meta-generalization gap (22), we assume the following.

Assumption 6. Either of the following assumptions hold:

- (a) Assumption 5 holds, or
- (b) For each task $\tau \in \mathcal{T}$, the loss function $l(W, Z)$ is δ_τ^2 -sub-Gaussian when $(W, Z) \sim P_{W|T} P_{Z|T}$, where $P_{W|T}$ is the marginal of the joint distribution $P_{W|Z^m, U} P_{Z_{1:N}^m, U} P_{Z^m|T}$. The average per-task training loss $L_t^{\text{joint}}(U|Z^m)$ is σ^2 -sub-Gaussian when $(U, Z^m) \sim P_U P_{Z^m}$.

As in Section 4.1.2, Assumption 6 can be seen to be implied by the sufficient conditions in Lemma 3.

Theorem 4. Under Assumption 6, for any meta learner $P_{U|Z_{1:N}^m}$, the following bound holds on the average meta-generalization gap

$$\left| \mathbb{E}_{P_{Z_{1:N}^m, U}} [\Delta \mathcal{L}^{\text{joint}}(U|Z_{1:N}^m)] \right| \leq \frac{1}{N} \sum_{i=1}^N \sqrt{2\sigma^2 I(U; Z_i^m)} + \mathbb{E}_{P_T} \left[\frac{1}{m} \sum_{j=1}^m \sqrt{2\delta_T^2 I(W; Z_j|T = \tau)} \right], \quad (33)$$

where the MI $I(U; Z_i^m)$ is evaluated with respect to $P_{Z_i^m, U}$ obtained by marginalizing $P_{Z_{1:N}^m, U}$, and the MI $I(W; Z_j|T = \tau)$ is with respect to $P_{Z_j, W|T=\tau}$ obtained by marginalizing $P_{Z^m, W|T=\tau}$.

Proof. See Appendix C. \square

Similar to the bound in (32), the bounds on meta-generalization gap in (33) are in terms of two types of mutual information, the first describing the sensitivity of the meta-learner and the second the sensitivity of the base-learner. Specifically, the MI $I(U; Z_i^m)$

quantifies the sensitivity of the output of the meta learner to per-task data set Z_i^m , and the MI $I(W; Z_j | T = \tau)$ measures the sensitivity of the output of the base-learner, $P_{W|Z^m, U}$ to each data sample Z_i within the training set Z^m of the meta-test task T . Moreover, it can be shown, in a manner similar to (31c), that, under Assumption 5, the ITMI bound in (33) is tighter than the MI bound in (32).

4.3. Discussion on Bounds

The bounds on the average meta-generalization gap obtained in this section generalize the bounds for conventional single-task learning in Section 3. To see this, consider the task distribution $P_T = \delta(T - \tau)$ to be centered at some task $\tau \in \mathcal{T}$. Recall that in conventional learning, the hyperparameter u is fixed a priori. As such, the mutual information $I(U; Z_{1:N}^m)$ (for MI-based bounds) and $I(U; Z_i^m)$ (for ITMI-based bounds) vanishes. For the separate within-task training and test sets, this implies that the average generalization gap is zero, which follows since the per-task test loss $L_t(W|Z_i^{\text{me}})$ is an unbiased estimate of per-task generalization loss $L_g(W|T_i)$. The MI- and ITMI-based bounds for the joint within-task training and test sets then reduce to

$$\left| \mathbb{E}_{P_{Z^m, W|\tau, u}} [\Delta L(W|Z^m, \tau)] \right| \leq \sqrt{\frac{2\delta_\tau^2}{m} I(W; Z^m)}, \quad (34)$$

and

$$\left| \mathbb{E}_{P_{Z^m, W|\tau, u}} [\Delta L(W|Z^m, \tau)] \right| \leq \frac{1}{m} \sum_{j=1}^m \sqrt{2\delta_\tau^2 I(W; Z_j)} \quad (35)$$

respectively, where $I(W; Z^m)$ is evaluated with respect to the joint distribution $P_{W, Z^m|\tau, u}$ and $I(W; Z_j)$ with respect to $P_{W, Z_j|\tau, u}$.

The MI- and ITMI-based bounds derived in this section point that a smaller correlation between hyperparameters and meta-training set and thus small mutual information $I(U; Z_{1:N}^m)$ improves the meta-generalization gap, although this seems deleterious to performance. To clarify this contradiction, we would like to emphasize that these bounds quantify the difference between meta-generalization loss and empirical training loss, which in turn depends on the sensitivity of the meta-learner and base-learner to their input meta-training set and per-task training set, respectively. The mutual information terms in our bounds capture these sensitivities. Consequently, our bounds suggest that a meta-learner that is highly correlated to the input meta-training set (i.e., when $I(U; Z_{1:N}^m)$ is large) does not generalize well (i.e., yields large meta-generalization gap). This property aligns with a previous information-theoretic analysis for generalization in conventional learning [25].

To the best of our knowledge, the MI- and ITMI-based bounds studied here are the first bounds on the average meta-generalization gap. As discussed in the introduction, these bounds are distinct from the high-probability PAC and PAC-Bayesian bounds on the meta-generalization gap studied previously on meta-learning. Consequently, the bounds studied in this work are not directly comparable with the existing high-probability bounds.

Finally, we note that similarity between tasks is crucial to meta-learning. If the per-task data distributions $P_{Z|T=\tau}$ in the task environment are ‘closer’ to each other, a meta-learner can efficiently learn the shared characteristics of tasks, and can generalize well to new tasks from the task environment. In our setting, the statistical properties of the task environment $(P_T, \{P_{Z|T=\tau}\}_{\tau \in \mathcal{T}})$ dictate this similarity. Although our MI- and ITMI-based bounds do not explicitly capture this, we note that the properties of task environment are implicitly accounted for by the mutual information terms $I(U; Z_{1:N}^m)$ and $I(U; Z_i^m)$, where the meta-training data set $Z_{1:N}^m$ is generated from the task environment, and also by the sub-Gaussianity considerations in Assumptions 3–6. From preliminary studies, we believe that information-theoretic bounds that explicitly capture the impact of task

similarity require a different performance metric than the average meta-generalization gap considered here, and is left to future work.

5. Applications

In this section, we consider two applications of the information-theoretic bounds proposed in Section 4.1. The first, simpler, example concerns a parameter estimation problem for which an optimized meta-learner can be obtained in closed form. In contrast, the second application covers a broad class of iterative meta-training schemes.

5.1. Parameter Estimation

To illustrate the bounds on the meta-generalization gap derived in Section 4.1, we first consider the problem of prediction for a Bernoulli process with a ‘soft’ predictor that uses only a few samples from the process, as well as meta-training data. Towards this, we consider an arbitrary discrete finite set of tasks $\mathcal{T} = \{\tau_1, \dots, \tau_M\}$. The data distribution $P_{Z|T=\tau_k}$ for each task $\tau_k \in \mathcal{T}, k \in \{1, \dots, M\}$, is given as Bernoulli(μ_{τ_k}) with mean parameter μ_{τ_k} . The task distribution P_T is then defined over the finite set of mean parameters $\{\mu_{\tau_1}, \dots, \mu_{\tau_M}\}$. The base-learner uses training data, distributed i.i.d. from Bernoulli(μ_{τ_k}) to determine the parameter W_k , which is used as a predictor of new observation $Z \sim \text{Bernoulli}(\mu_{\tau_k})$ at test time. The loss function is defined as $l(w, z) = (w - z)^2$, measuring the quadratic error between prediction and realized test input z . Note that the optimal (Bayes) predictor, computable in the ideal case of known distribution $P_{Z|T=\tau_k}$, is given as $W = \mu_{\tau_k}$. We now distinguish the two cases with separate and joint within-task training and test sets.

5.1.1. Separate Within-Task Training and Test Sets

The base-learner $P_{W|Z_k^{m_{tr}}, u}$ for task $\tau_k \in \mathcal{T}$, deterministically selects the prediction

$$W_k = \alpha D_k^{m_{tr}} + (1 - \alpha)u, \tag{36}$$

where $D_k^{m_{tr}} = \frac{1}{m_{tr}} \sum_{j=1}^{m_{tr}} Z_{k,j}^{m_{tr}}$ is an empirical average over the training set $Z_{k,j}^{m_{tr}}$, u is a hyperparameter defining a bias that can be meta-trained, and $\alpha \in [0, 1)$ is a fixed scalar. Here, $Z_{k,j}^{m_{tr}}$ denote the j th data sample in the training set $Z_k^{m_{tr}}$ of task τ_k . The bias term in (36) may help approximate the ideal Bayes predictor in the presence of limited data $Z_k^{m_{tr}}$.

The objective of the meta-learner is to infer the hyperparameter u . For a given meta-training data set $Z_{1:N}^m$, comprising of data sets from N tasks sampled according to P_T , the meta-learner can compute the empirical meta-training loss as

$$\mathcal{L}_t^{\text{sep}}(u|Z_{1:N}^m) = \frac{1}{N} \sum_{k=1}^N \frac{1}{m_{te}} \sum_{j=1}^{m_{te}} (W_k - Z_{k,j}^{m_{te}})^2, \tag{37}$$

where $Z_{k,j}^{m_{te}}$ denote the j th example in the test set of Z_k^m , the k th sub-data set of $Z_{1:N}^m$. The meta-learner $P_{U|Z_{1:N}^m}$ then deterministically selects the minimizing hyperparameter u of the meta-training empirical loss function in (37). This optimization yields

$$U = \frac{(1 - \alpha)^{-1}}{N} \left(\sum_{k=1}^N D_k^{m_{te}} - \alpha D_k^{m_{tr}} \right), \tag{38}$$

where $D_k^{m_{te}} = \sum_{j=1}^{m_{te}} Z_{k,j}^{m_{te}} / m_{te}$. Note that $D_k^{m_{te}}$ and $D_k^{m_{tr}}$ are binomial random variables and by (38), U takes finitely many discrete values and is bounded as $-\alpha(1 - \alpha)^{-1} \leq U \leq (1 - \alpha)^{-1}$. The meta-test loss can be explicitly computed as

$$\mathcal{L}_g^{\text{sep}}(u) = (1 - \alpha)^2 (u^2 - 2u\mathbb{E}_{P_T}[\mu_T]) + \mathbb{E}_{P_T} \left[\alpha^2 \left(\mu_T^2 + \frac{\mu_T \bar{\mu}_T}{m_{tr}} \right) + \mu_T - 2\alpha\mu_T^2 \right], \tag{39}$$

where $\bar{\mu}_T = 1 - \mu_T$, and the average meta-generalization gap evaluates to

$$\mathbb{E}_{P_{Z_{1:N}^m, U}}[\Delta\mathcal{L}^{\text{sep}}(U|Z_{1:N}^m)] = \frac{2(1 + \alpha^2)}{N} \text{Var}_T + \frac{2\mathbb{E}_{P_T}[\mu_T\bar{\mu}_T]}{N} \left(\frac{1}{m_{\text{te}}} + \frac{\alpha^2}{m_{\text{tr}}} \right), \quad (40)$$

where $\text{Var}_T = (\mathbb{E}_{P_T}[\mu_T^2] - (\mathbb{E}_{P_T}[\mu_T])^2)$ is the variance of μ_T .

To compute the MI- and ITMI-based bounds on the meta-generalization gap (40), it is easy to verify that the average training loss $L_t^{\text{sep}}(\cdot|Z^m)$ is bounded, i.e., $0 \leq L_t^{\text{sep}}(\cdot|Z^m) \leq (1 + \alpha)^2$ for all $u \in \mathcal{U}$ and $Z^m \in \mathcal{Z}^m$. Thus, Assumption 3 for the MI bound and also Assumption 4 for the ITMI bound hold with $\sigma^2 = (1 + \alpha)^4/4$. For the MI bound, we note that, since the meta-learner is deterministic, we have that $I(U; Z_{1:N}^m) = H(U)$. The ITMI bound (30) is given as

$$|\mathbb{E}_{P_{Z_{1:N}^m, U}}[\Delta\mathcal{L}^{\text{sep}}(U|Z_{1:N}^m)]| \leq \frac{1}{N} \sum_{i=1}^N \sqrt{\frac{(1 + \alpha)^4}{2} I(U; Z_i^m)}. \quad (41)$$

The information-theoretic measures in (41) can be evaluated numerically as discussed in Appendix D.

For a numerical illustration, Figure 4 plots the average of the meta-generalization loss (39) and average meta-training loss (A16) along with the ITMI bound in (41) and MI bound in (27). It can be seen that the ITMI bound is tighter than MI bound and correctly predicts the decrease in the meta-generalization gap as the number N of tasks increases.




Figure 4. Comparison of the MI bound in (27) and ITMI-based bound obtained in (41) with the meta-generalization gap for meta-learning with separate within-task training and test sets. The task environment is defined by $M = 12$ tasks. Other parameters are set as $\alpha = 0.15$, $m_{\text{tr}} = 15$, $m_{\text{te}} = 5$.

5.1.2. Joint Within-Task Training and Testing sets

We now consider the case with joint within-task training and test sets. The base-learner $P_{W|Z_k^m, U}$ for task $\tau_k \in \mathcal{T}$ still uses the predictor (36), but now the empirical average over the training set is given as $D_k = \sum_{j=1}^m Z_{k,j}^m/m$. As before, the meta-learner $P_{U|Z_{1:N}^m}$ deterministically selects the minimizing hyperparameter u of the meta-training empirical loss function, $L_{Z_{1:N}^m}(u) = (1/N) \sum_{k=1}^N (1/m) \sum_{j=1}^m (W_i - Z_{k,j}^m)^2$, yielding $U = \frac{1}{N} \sum_{k=1}^N D_k$. As discussed in Appendix D, the meta-generalization loss for this example can also be explicitly computed and the meta-generalization gap bounds in (32) and (33) can be evaluated numerically. Figure 5 plots the average meta-generalization loss and average meta-training loss along with the MI bound in (32) and ITMI bound in (A18), as a function of per-task

data samples m . The ITMI bound is seen to better reflect the decrease of the meta-training loss as a function of m .

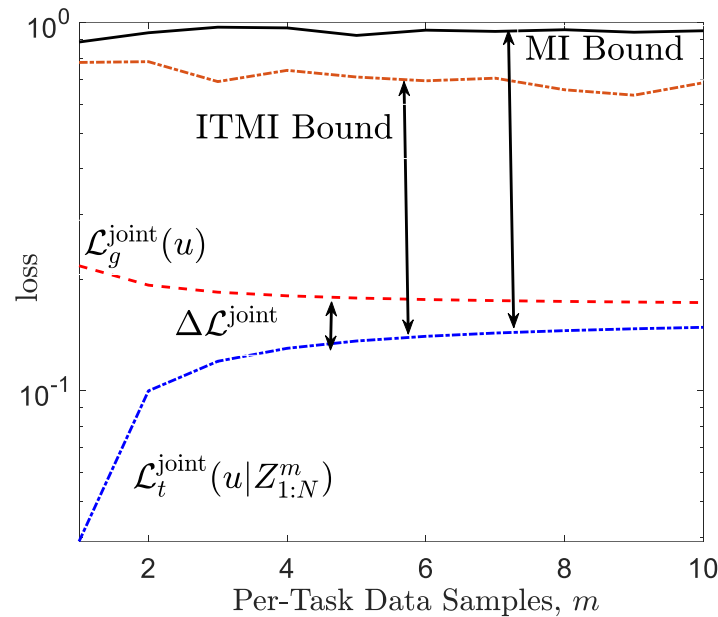


Figure 5. Comparison of the MI- and ITMI-based bound obtained in (A18) with the meta-generalization gap for meta-learning with joint within-task training and test sets, as a function of the per-task data samples m for $N = 5$ and $\alpha = 0.55$. The task environment is defined by $M = 9$ tasks.

5.2. Noisy Iterative Meta-Learning Algorithms

Most meta-learning algorithms are built around a nested loop structure, with the inner loop applying the base-learner on the meta-training set and the outer loop updating the hyperparameters U . In this section, we focus on a vast class of such meta-learning algorithms in which the inner loop applies training procedures dependent on the current iterate of the hyperparameter, while the outer loop updates the hyperparameter using a stochastic rule. This class includes stochastic variants of state-of-the-art algorithms such as MAML [23] and reptile [24]. We apply the derived information-theoretic bounds to study the meta-generalization performance of the mentioned class of meta-training iterative stochastic rules by focusing on the case of separate within-task training and test sets here, which is assumed e.g., by MAML. The analysis for the setup with joint within-task training and test sets can also be carried out at the cost of a more cumbersome notation.

To start, let $U^j \in \mathbb{R}^d$ denote the hyperparameter vector at outer iteration j , with $U^0 \in \mathbb{R}^d$ being an arbitrary initialization. For example, in MAML, the hyperparameter U defines the initial iterate used by each base-learner in the inner loop to update the model parameter W_τ corresponding to task τ . At each iteration $j \geq 1$, we randomly select a mini-batch of task indices $K_j \subseteq [1, \dots, N]$ from the meta-training data $Z_{1:N}^m$, obtaining the corresponding data set $Z_{K_j}^m = (Z_{K_j}^{\text{mtr}}, Z_{K_j}^{\text{mte}}) \subseteq Z_{1:N}^m$, where $Z_{K_j}^{\text{mtr}} = \{Z_k^{\text{mtr}}\}_{k \in K_j}$ and $Z_{K_j}^{\text{mte}} = \{Z_k^{\text{mte}}\}_{k \in K_j}$ are the separate training and test sets for the selected tasks. For each index $k \in K_j$, in the inner loop, the base-learner selects the model parameter W_k^j as a possibly stochastic function

$$W_k^j = g(U^{j-1}, Z_k^{\text{mtr}}). \tag{42}$$

For instance, in MAML, the function $g(U^{j-1}, Z_k^{\text{mtr}}) \in \mathbb{R}^d$ in (42) represents the output of an SGD procedure that starts from initialization U^{j-1} and uses the task training data Z_k^{mtr} to iteratively update the model parameters, producing the final iterate W_k^j . We denote

as $W_{K_j} = \{W_k^j\}_{k \in K_j}$ the collection of the base-learners' outputs for all task indices $k \in K_j$ at outer iteration j .

In the outer loop, the meta-learner uses the task-specific adapted parameters W_{K_j} from the inner loop and the meta-test set $Z_{K_j}^{mte}$ to update the past iterate U^{j-1} according to the general update rule

$$U^j = F(U^{j-1}) + \beta_j G(U^{j-1}, W_{K_j}, Z_{K_j}^{mte}) + \xi_j, \tag{43}$$

where $F(\cdot)$ and $G(\cdot, \cdot, \cdot)$ are arbitrary deterministic functions; β_j is the step-size; and $\xi_j \sim \mathcal{N}(0, \gamma_j^2 I_d)$ is an isotropic Gaussian noise, independently drawn for $j = 1, 2, \dots$. As an example, in MAML, the function $F(\cdot)$ is the identity function and function $G(\cdot, \cdot, \cdot)$ equals the gradient of the empirical loss $1/|K_j| \sum_{k \in K_j} L_t^{sep}(W_k^j | Z_k^{mte})$ in (14) with respect to U^{j-1} . Note, however, that MAML does not add noise, i.e., $\gamma_j^2 = 0$ for all j .

The final output of the meta-learning algorithm is then defined as an arbitrary function $U = f(U^1, \dots, U^J)$, of all iterates. Examples of function f include the last update $f(U^1, \dots, U^J) = f(U^J)$ and average of the updates $f(U^1, \dots, U^J) = 1/J \sum_{j=1}^J U^j$. A graphical model representation of the variables involved is shown in Figure 6.

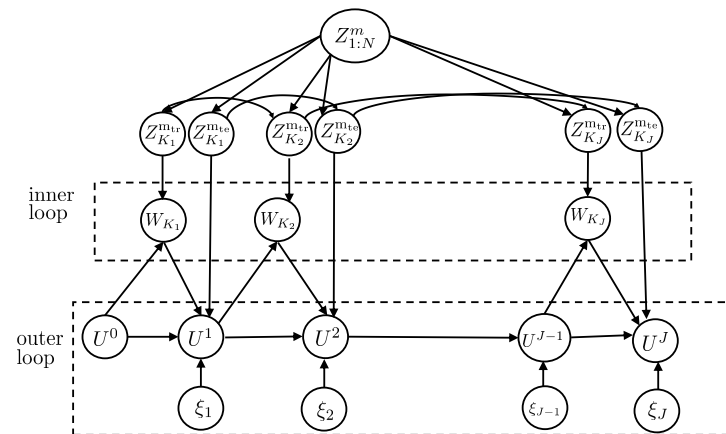


Figure 6. A graphical model representation of the variables involved in the Definition of noisy iterative algorithms.

We now derive an upper bound on the meta-generalization gap for the general class of iterative meta-learning algorithm satisfying (42) and (43) under the following assumptions.

Assumption 7.

- (1) For the base-learner given in (42), the average training loss $L_t^{sep}(u | Z^m)$ in (15) is σ^2 -sub-Gaussian for all $u \in \mathcal{U}$ when $Z^m \sim P_{Z^m}$;
- (2) The meta-training data set $Z_{K_j}^m$ sampled at each iteration j is conditionally independent of the history of model-parameter vectors $\{W_{K_i}\}_{i=1}^{j-1}$ and hyperparameter $U^{(j-1)} = (U^1, U^2, \dots, U^{j-1})$, i.e.,

$$P_{Z_{K_j}^m | \{Z_{K_i}^m\}_{i=1}^{j-1}, Z_{1:N}^m, U^{(j-1)}, \{W_{K_i}\}_{i=1}^{j-1}} = P_{Z_{K_j}^m | \{Z_{K_i}^m\}_{i=1}^{j-1}, Z_{1:N}^m}; \tag{44}$$

- (3) The meta-parameter update function $G(\cdot, \cdot, \cdot)$ is uniformly bounded, i.e., $\|G(\cdot, \cdot, \cdot)\|_2 \leq L$ for some $L > 0$.

Lemma 4. Under Assumption 7, the following upper bound on the meta-generalization gap (18) holds for the class of noisy iterative meta-training algorithms (42) and (43)

$$\mathbb{E}_{P_{Z_{1:N}^m, U}}[\Delta\mathcal{L}^{\text{sep}}(U|Z_{1:N}^m)] \leq \sqrt{\frac{2\sigma^2}{N} \sum_{j=1}^J \frac{d}{2} \log\left(1 + \frac{\beta_j^2 L^2}{d\gamma_j^2}\right)}. \tag{45}$$

Proof. See Appendix E. \square

The bound in (45) has the same form as the generalization gap derived in [27] for conventional learning. From (45), the generalization gap can be reduced by increasing the variance γ_j^2 of the injected Gaussian noise. In particular, the meta-generalization gap depends on the ratios β_j^2/γ_j^2 between squared step size β_j^2 and variance γ_j^2 . For example, SGLD sets $\gamma_j = \sqrt{\beta_j}$, and a step size β_j decaying over time according to the standard Robbins–Monro conditions, in order to ensure convergence of the output samples to the generalized posterior distribution of the hyperparameters [52].

Example: To illustrate bound (45), we now consider a simple logistic regression problem that generalizes the example studied in Section 5.1. Accordingly, each data point Z corresponds to labelled data $Z = (X, Y)$, where $X \in \{0, 1\}^d$ represents the input vector and $Y \in \{0, 1\}$ represents the corresponding binary label. The data distribution $P_{Z|\tau_k} = P_{X|\tau_k}P_{Y|X, \tau_k}$ for each task $\tau_k \in \mathcal{T} = \{\tau_1, \dots, \tau_M\}$ is such that $X \sim P_{X|\tau_k}$ is a d -dimensional Bernoulli vector obtained via d independent draws from $\text{Bernoulli}(v)$ and Y is distributed as $Y \sim \text{Bernoulli}(\phi(\mu_{\tau_k}^T X))$, where $\phi(a) = 1/(1 + \exp(-a))$ is the sigmoid function and $\mu_{\tau_k} \in \mathbb{R}^d$, with $\|\mu_{\tau_k}\|_2 \leq 1$. The task distribution P_T then defines a distribution over the parameter vectors $\{\mu_{\tau_1}, \dots, \mu_{\tau_M}\}$. The base-learner uses training data generated i.i.d. from $P_{Z|\tau_k}$ to obtain a prediction w of the parameter vector μ_{τ_k} for task $\tau_k \in \mathcal{T}$. The loss function is taken as the quadratic error $l(w, z) = (\phi(w^T x) - y)^2$.

At each iteration j , starting from initialization point U^{j-1} , the base-learner in (42) uses a one-step projected gradient descent algorithm on the training data set Z_k^{mtr} to obtain the prediction W_k^j as

$$W_k^j = \text{proj}_{\mathcal{W}}\left(U^{j-1} - \alpha \nabla_w L_t^{\text{sep}}(w|Z_k^{\text{mtr}})|_{w=U^{j-1}}\right), \tag{46}$$

where $\alpha > 0$ is the step-size $\mathcal{W} = \{w \in \mathbb{R}^d \mid \|w\|_2 \leq 1\}$ is the set of feasible model parameters and $\text{proj}_{\mathcal{A}}(b) = \frac{1}{2} \min_{a \in \mathcal{A}} \|a - b\|_2^2$ is the projection operator. The meta-learner (43) updates the initialization vector according to the noisy gradient descent rule

$$U^j = U^{j-1} - \beta_j \left(\frac{1}{|K_j|} \sum_{k=1}^{|K_j|} \nabla_w L_t^{\text{sep}}(w|Z_k^{\text{mte}})|_{w=W_k^j} \right) + \xi_j, \tag{47}$$

where β_j is the step-size; and $\xi_j \sim \mathcal{N}(0, \gamma_j^2 I_d)$ is isotropic Gaussian noise. This update rule corresponds to performing a first order MAML (FOMAML) [23] with the addition of noise.

For this problem, it is easy to verify that Assumption 7 is satisfied, since the loss function $l(\cdot, \cdot)$ is bounded in the interval $[0, 1]$, whereby $L_t^{\text{sep}}(u|Z^m)$ is also $[a, b]$ -bounded. We also have the inequality

$$\left\| \frac{1}{|K_t|} \sum_{i=1}^{|K_t|} \nabla_w L_{Z^{\text{mte}}}(w)|_{w=W_i^t} \right\|_2 \leq 2\sqrt{d}e^{\sqrt{d}} \triangleq L. \tag{48}$$

The MI bound in (45) then evaluates to

$$\mathbb{E}_{P_{Z_{1:N}^m, U}}[\Delta\mathcal{L}(U|Z_{1:N}^m)] \leq \sqrt{\frac{1}{2N} \sum_{j=1}^J \frac{d}{2} \log\left(1 + \frac{4\beta_j^2 e^{2\sqrt{d}}}{\gamma_j^2}\right)}. \quad (49)$$

We now evaluate the meta-training and meta-test loss, along with the bound (49) as a function of the ratio γ_j^2/β_j^2 in Figure 7. For the experiment, we considered a task environment of $M = 20$ tasks with $\nu = 0.4$, $d = 3$, $N = 4$ meta-training tasks with $m_{tr} = 10$ training data samples and $m_{te} = 5$ test data samples. For the inner-loop (46), we fixed step-size $\alpha = 10^{-4}$ and for the outer-loop (47), we set $|K_t| = N$, $\beta_j = 0.25$ and $T = 200$ iterations.

As suggested by Lemma 4, the meta-generalization gap decreases with addition of noise. While the MI bound (45) is generally loose, it correctly quantifies the dependence of the meta-generalization loss and the ratio γ_j^2/β_j^2 , and it can hence serve as a useful meta-training criterion [20,48].




Figure 7. Comparison of the meta-generalization gap with the MI-based bound in (49) as function of the ratio γ_t^2/β_t^2 .

6. Conclusions

This work has presented novel information-theoretic upper bounds on the average generalization gap of meta-learning algorithms, thereby extending the well-studied information-theoretic approaches in conventional learning to meta-learning. The proposed bounds capture two sources of uncertainty—environment-level uncertainty and within-task uncertainty—and bound them via separate mutual information terms. Applications were also discussed, with the aim of elucidating the use of the bounds to quantify meta-overfitting and guide the choice of the meta-inductive bias, i.e., the class of inductive biases. The derived bounds are amenable to further refinements, such as those along the lines of [39,40,45]. It would also be interesting to study the meta-generalization bounds on noisy iterative meta-learning algorithms using the tighter information-theoretic bounds such as [26,40].

Author Contributions: Conceptualization, formal analysis, writing—original draft, S.T.J.; conceptualization, funding acquisition and supervision, O.S. Both authors have read and agreed to the published version of the manuscript.

Funding: The authors have received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 Research and Innovation Programme (Grant Agreement No. 725731).

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Decoupling Estimate Lemmas

The proofs of the main results rely on the following decoupling estimate lemmas, which bound the difference in expectations under a change of measure from the joint $P_{X,Y}$ to the product of the marginals $P_X P_Y$. In order to state the most general form of decoupling estimate lemmas, we first define a generalized sub-Gaussian random variable.

Definition A1. A random variable X is said to be $(\Psi_+, \Psi_-, b_+, b_-)$ -generalized sub-Gaussian if there exist convex functions $\Psi_+ : \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\Psi_- : \mathbb{R}_+ \rightarrow \mathbb{R}$ that satisfy the equalities $\Psi_+(0) = \Psi_-(0) = \Psi'_+(0) = \Psi'_-(0) = 0$ and bound the CGF of X as

$$\Lambda_X(\lambda) \leq \Psi_+(\lambda), \quad \text{for } \lambda \in [0, b_+) \tag{A1a}$$

$$\Lambda_X(\lambda) \leq \Psi_-(-\lambda), \quad \text{for } \lambda \in (b_-, 0], \tag{A1b}$$

for some constants $0 < b_+ \leq \infty$ and $-\infty \leq b_- < 0$.

For a $(\Psi_+, \Psi_-, b_+, b_-)$ -generalized sub-Gaussian random variable, we also introduce the following standard definitions. First, the Legendre dual of function $\Psi_+(\lambda)$ is defined as

$$\Psi_+^*(x) = \sup_{\lambda \in [0, b_+)} (\lambda x - \Psi_+(\lambda)). \tag{A2}$$

It can be easily seen that $\Psi_+^*(\cdot)$ is a non-negative, convex, and non-decreasing function on $[0, \infty)$ with $\Psi_+^*(0) = 0$. Second, the inverse Legendre dual of function $\Psi_+(\lambda)$ is defined as $\Psi_+^{*-1}(y) = \inf\{x \geq 0 : \Psi_+^*(x) \geq y\}$. This function is concave, and it can be equivalently written as [26]

$$\Psi_+^{*-1}(y) = \inf_{\lambda \in [0, b_+)} \frac{y + \Psi_+(\lambda)}{\lambda}. \tag{A3}$$

Similar definitions and results apply for $\Psi_-(\cdot)$.

A σ^2 -sub-Gaussian random variable X is a generalized sub-Gaussian variable with $\Psi_+(\lambda) = \Psi_-(\lambda) = \lambda^2 \sigma^2 / 2$, $b_+ = \infty$ and $b_- = -\infty$. Furthermore, the Legendre dual functions are given as $\Psi_+^*(x) = \Psi_-^*(x) = x^2 / (2\sigma^2)$, and the inverse Legendre dual functions evaluate to

$$\Psi_+^{*-1}(y) = \Psi_-^{*-1}(y) = \sqrt{2\sigma^2 y}. \tag{A4}$$

We are now ready to state the decoupling estimate lemmas.

Lemma A1 (Decoupling Estimate [14]). Let $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ be two jointly distributed random variables with joint distribution $P_{X,Y}$, and let $f(X, Y)$ be a real valued function such that $f(x, Y)$ is $(\Psi_+, \Psi_-, \infty, -\infty)$ -generalized sub-Gaussian for all $x \in \mathcal{X}$ when $Y \sim P_Y$. Then we have the following inequalities

$$\pm \left(\mathbb{E}_{P_X P_Y} [f(\tilde{X}, \tilde{Y})] - \mathbb{E}_{P_{X,Y}} [f(X, Y)] \right) \leq \Psi_{\mp}^{*-1}(I(X; Y)), \tag{A5}$$

where $(\tilde{X}, \tilde{Y}) \sim P_X P_Y$.

Lemma A2 (General Decoupling Estimate [26]). *Let $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ be two jointly distributed random variables with joint distribution $P_{X,Y}$, and let $f(X, Y)$ be a real valued function such that $f(X, Y)$ is a $(\Psi_+, \Psi_-, b_+, b_-)$ -generalized sub-Gaussian when $(X, Y) \sim P_X P_Y$. Then, we have the inequality (A5).*

Note that in Lemma A2, the random variables X, Y are jointly distributed according to $P_{X,Y}$. Assuming that the function $f(X, Y)$ is generalized sub-Gaussian under $X \sim P_X$ and $Y \sim P_Y$ with P_X and P_Y being the marginals of $P_{X,Y}$, the lemma provides an upper bound on the difference between average of $f(X, Y)$ when (X, Y) is jointly distributed according to $P_{X,Y}$ and the average of $f(X, Y)$ when (X, Y) is independent with $X \sim P_X$ and $Y \sim P_Y$. The resultant bound thus provides an estimate of the effect of decoupling of the joint distribution to its marginals with respect to function $f(X, Y)$.

Appendix B. Proofs of Theorems 1 and 2

For the proof of Theorem 1, we use the decomposition (29) of the meta-generalization gap into average environment-level and within-task generalization gaps as

$$\begin{aligned} & \mathbb{E}_{P_{Z_{1:N}^m, U}} [\Delta \mathcal{L}^{\text{sep}}(U | Z_{1:N}^m)] \\ &= \mathbb{E}_{P_{Z_{1:N}^m, U}} \left[(\mathcal{L}_g^{\text{sep}}(U) - \mathcal{L}_{g,t}^{\text{sep}}(U)) + (\mathcal{L}_{g,t}^{\text{sep}}(U) - \mathcal{L}_t^{\text{sep}}(U | Z_{1:N}^m)) \right] \tag{A6} \\ &= \mathbb{E}_{P_T} \left[\mathbb{E}_{P_{Z^m | T} P_{W, U, Z_{1:N}^m | Z^{\text{mtr}}}} [\Delta L(W | Z^{\text{mte}}, T)] \right] + \mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{\text{sep}}(U) - \mathcal{L}_t^{\text{sep}}(U | Z_{1:N}^m)] \end{aligned}$$

where (A6) follows since the average within-task generalization gap for a random meta-test task $\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_g^{\text{sep}}(U) - \mathcal{L}_{g,t}^{\text{sep}}(U)]$ can be equivalently written as $\mathbb{E}_{P_T} [\mathbb{E}_{P_{Z^m | T} P_{W, U, Z_{1:N}^m | Z^{\text{mtr}}}} [\Delta L(W | T)]]$, with $\Delta L(W | Z^{\text{mte}}, T) = L_{P_{Z|T}}(W) - L_t(W | Z^{\text{mte}})$ denoting the generalization gap of the meta-test task T , and the joint distribution $P_{W, U, Z_{1:N}^m | Z^{\text{mtr}}}$ factorizes as $P_{W | Z^{\text{mtr}}, U} P_{Z_{1:N}^m | U}$. To obtain an upper bound on the average meta-generalization gap $\mathbb{E}_{P_{Z_{1:N}^m, U}} [\Delta \mathcal{L}^{\text{sep}}(U | Z_{1:N}^m)]$, we bound each of the two differences in (A6) separately.

We first bound the second difference in (A6) $\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{\text{sep}}(U) - \mathcal{L}_t^{\text{sep}}(U | Z_{1:N}^m)]$, which represents the expected environment-level uncertainty measured using the average per-task training loss $L_t^{\text{sep}}(u | Z^m)$ defined in (15). To this end, we extend the single-task learning generalization bound of Lemma 1 by resorting to the decoupling estimate in Lemma A1 with $X = U, Y = Z_{1:N}^m$ and $f(X, Y) = \mathcal{L}_t^{\text{sep}}(U | Z_{1:N}^m)$, so that $\mathbb{E}_{P_{X,Y}} [f(X, Y)] = \mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_t^{\text{sep}}(U | Z_{1:N}^m)]$ and $\mathbb{E}_{P_X P_Y} [f(\tilde{X}, \tilde{Y})] = \mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_g^{\text{sep}}(U)]$.

Recall from Assumption 3 that for a given $u \in \mathcal{U}, \mathcal{L}_t^{\text{sep}}(u | Z_{1:N}^m) = \frac{1}{N} \sum_{i=1}^N L_t^{\text{sep}}(u | Z_i^m)$ is the average of σ^2 -sub-Gaussian i.i.d. random variables $L_t^{\text{sep}}(u | Z_i^m)$ under $Z_i^m \sim P_{Z^m}$ for all $i \in \{1, \dots, N\}$. It then follows that $\mathcal{L}_t^{\text{sep}}(u | Z_{1:N}^m)$ is σ^2/N -sub-Gaussian under $Z_{1:N}^m \sim P_{Z_{1:N}^m}$ for all $u \in \mathcal{U}$ [53]. Applying Lemma A1 with Ψ_{\mp}^{*-1} specialized to the σ^2/N -sub-Gaussian loss in (A4) gives that

$$\left| \mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{\text{sep}}(U) - \mathcal{L}_t^{\text{sep}}(U | Z_{1:N}^m)] \right| \leq \sqrt{\frac{2\sigma^2 I(U; Z_{1:N}^m)}{N}}. \tag{A7}$$

We now evaluate the first difference in (A6). It can be seen that for a fixed task $\tau \in \mathcal{T}$, the average within-task uncertainty evaluates to

$$\begin{aligned} \mathbb{E}_{P_{Z^m | T=\tau} P_{W, U, Z_{1:N}^m | Z^{\text{mtr}}}} [\Delta L(W | Z^{\text{mte}}, T = \tau)] &= \mathbb{E}_{P_{Z^m | T=\tau} P_{W | Z^{\text{mtr}}}} [\Delta L(W | Z^{\text{mte}}, T = \tau)] \\ &= \mathbb{E}_{P_{Z^{\text{mtr}} | T=\tau}} \mathbb{E}_{P_{W | Z^{\text{mtr}}}} [L_g(W | T = \tau) - \mathbb{E}_{P_{Z^{\text{mte}} | T=\tau}} L_t(W | Z^{\text{mte}})] \tag{A8} \\ &\stackrel{(a)}{=} 0, \end{aligned}$$

where (a) follows, since W and $Z^{m_{te}}$ are conditionally independent given $Z^{m_{tr}}$, whereby $\mathbb{E}_{P_{Z^{m_{tr}}|T=\tau}} \mathbb{E}_{P_{W|Z^{m_{tr}}}} [\mathbb{E}_{P_{Z^{m_{te}}|T=\tau}} L_t(W|Z^{m_{te}})] = \mathbb{E}_{P_{Z^{m_{tr}}|T=\tau}} \mathbb{E}_{P_{W|Z^{m_{tr}}}} [L_g(W|T = \tau)]$. Substituting (A7) and (A8) in (A6), then concludes the proof.

For Theorem 2, the proof follows along the same line, bounding the average environment-level uncertainty $\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{sep}(U) - \mathcal{L}_t^{sep}(U|Z_{1:N}^m)]$. Towards this, we note that the environment-level uncertainty can be equivalently written as

$$\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{sep}(U) - \mathcal{L}_t^{sep}(U|Z_{1:N}^m)] = \frac{1}{N} \sum_{i=1}^N \left(\mathbb{E}_{P_{Z_i^m, P_U}} [L_t^{sep}(U|Z_i^m)] - \mathbb{E}_{P_{Z_i^m, U}} [L_t^{sep}(U|Z_i^m)] \right), \tag{A9}$$

where Z^m and U in the first term are conditionally independent random variables distributed as $(Z_i^m, U) \sim P_{Z_i^m} P_U$, while, in the second term, they are jointly distributed according to $P_{Z_i^m, U}$, which is obtained by marginalizing the joint distribution $P_{Z_{1:N}^m, U}$. Under Assumption 4(a), we can bound the difference $\mathbb{E}_{P_{Z_i^m, P_U}} [L_t^{sep}(U|Z_i^m)] - \mathbb{E}_{P_{Z_i^m, U}} [L_t^{sep}(U|Z_i^m)]$ by resorting to the decoupling estimate in Lemma A1 with $X = U, Y = Z_i^m, f(X, Y) = L_t^{sep}(U|Z_i^m)$ such that $\mathbb{E}_{P_{X,Y}} [f(X, Y)] = \mathbb{E}_{P_{U, Z_i^m}} [L_t^{sep}(U|Z_i^m)]$ and $\mathbb{E}_{P_X P_Y} [f(\tilde{X}, \tilde{Y})] = \mathbb{E}_{P_U P_{Z_i^m}} [L_t^{sep}(U|Z_i^m)]$. Since $L_t^{sep}(U|Z_i^m)$ is σ^2 -sub-Gaussian under Assumption 4(a) for all $u \in \mathcal{U}$, Lemma A1 yields the following bound

$$\left| \mathbb{E}_{P_{Z_i^m, P_U}} [L_t^{sep}(U|Z_i^m)] - \mathbb{E}_{P_{Z_i^m, U}} [L_t^{sep}(U|Z_i^m)] \right| \leq \sqrt{2\sigma^2 I(U; Z_i^m)}. \tag{A10}$$

The bound in (A10) can also be obtained using Assumption 4(b) by resorting to the general decoupling estimate in Lemma A2 by fixing $X = U, Y = Z_i^m, f(X, Y) = L_t^{sep}(U|Z_i^m)$ such that $\mathbb{E}_{P_{X,Y}} [f(X, Y)] = \mathbb{E}_{P_{U, Z_i^m}} [L_t^{sep}(U|Z_i^m)]$ and $\mathbb{E}_{P_X P_Y} [f(\tilde{X}, \tilde{Y})] = \mathbb{E}_{P_U P_{Z_i^m}} [L_t^{sep}(U|Z_i^m)]$. Substituting the bound in (A10) in (A9) then yields the required bound in (27).

Appendix C. Proofs of Theorems 3 and 4

For Theorem 3, we start from the following decomposition of the average meta-generalization gap analogous to (A6)

$$\mathbb{E}_{P_{Z_{1:N}^m, U}} [\Delta \mathcal{L}^{joint}(U|Z_{1:N}^m)] = \mathbb{E}_{P_T} \left[\mathbb{E}_{P_{Z^m|T} P_{W|Z^m}} [\Delta L(W|Z^m, T)] \right] + \mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{joint}(U) - \mathcal{L}_t^{joint}(U|Z_{1:N}^m)] \tag{A11}$$

where $\mathcal{L}_{g,t}^{joint}(u) = \mathbb{E}_{P_{T, Z^m}} [L_t^{joint}(u|Z^m)]$ is the average training loss for the randomly selected meta-test task as a function of the hyperparameter u , and $\Delta L(W|Z^m, T) = L_{P_{Z|T}}(W) - L_t(W|Z^m)$ is the generalization gap for the meta-test task T . The MI bound on the expected environment-level uncertainty, $\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_{g,t}^{joint}(U) - \mathcal{L}_t^{joint}(U|Z_{1:N}^m)]$, can be obtained by using Lemma A1 and the Assumption 5(b) as in (A7).

The main difference between the separate and joint within-task training and test sets scenarios is that while the average within-task uncertainty vanishes in the former scenario, this is not the case for joint within-task training and training sets. Consequently, we now bound the average within-task generalization gap denoted by the first difference in (A11). For given task $\tau \in \mathcal{T}$, to bound the within-task generalization gap $\mathbb{E}_{P_{Z^m|T=\tau} P_{W|Z^m}} [\Delta L(W|Z^m, T = \tau)]$, we resort to Lemma A1 with $X = W, Y = Z^m$ and $f(X, Y) = L_t(W|Z^m)$, so that $\mathbb{E}_{P_{X,Y}} [f(X, Y)] = \mathbb{E}_{P_{W, Z^m|T=\tau}} [L_t(W|Z^m)]$. It can be then verified that $\mathbb{E}_{P_X P_Y} [f(\tilde{X}, \tilde{Y})] = \mathbb{E}_{P_{W|T=\tau}} \mathbb{E}_{P_{Z^m|T=\tau}} [L_t(W|Z^m)] = \mathbb{E}_{P_{W|T=\tau}} [L_g(W|T = \tau)] = \mathbb{E}_{P_{W, Z^m|T=\tau}} [L_g(W|T = \tau)]$, where $P_{W, Z^m|T=\tau} = P_{W|Z^m} P_{Z^m|T=\tau}$. Since $L_t(w|Z^m)$ is the sum of i.i.d δ_τ^2 -sub-Gaussian random variables $l(w, Z_i)$ (from Assumption 5(a)), we have that

$L_t(w|Z^m)$ is δ_τ^2/m -sub-Gaussian under $Z^m \sim P_{Z^m|T=\tau}$ for all $w \in \mathcal{W}$ [53]. Consequently, Lemma A1 yields the following bound

$$\left| \mathbb{E}_{P_{Z^m|T=\tau}P_{W|Z^m}} [\Delta L(W|Z^m, T = \tau)] \right| \leq \sqrt{\frac{2\delta_\tau^2}{m} I(W; Z^m|T = \tau)}. \tag{A12}$$

Averaging with respect to P_T on both sides of (A12), and combining with the bound on average environment-level uncertainty yields the required bound in (32) via Jensen’s inequality.

For Theorem 4, the proof follows along the same line. The ITMI bound on the expected environment-level uncertainty can be obtained along the lines of (A10), using the assumption on $L_t^{\text{joint}}(u|Z^m)$ in either Assumption 6(a) or Assumption 6(b). We now show that we can similarly bound the within-task uncertainty using the assumption on loss function $l(w, z)$ in either Assumption 6(a) or Assumption 6(b). Towards this, for fixed task $\tau \in \mathcal{T}$, we write the average within-task uncertainty equivalently as

$$\mathbb{E}_{P_{Z^m|T=\tau}P_{W|Z^m}} [\Delta L(W|Z^m, T = \tau)] = \frac{1}{m} \sum_{j=1}^m \left(\mathbb{E}_{P_{W|T=\tau}P_{Z_j|T=\tau}} [l(W, Z_j)] - \mathbb{E}_{P_{W, Z_j|T=\tau}} [l(W, Z_j)] \right), \tag{A13}$$

where W and Z_j in the second term are jointly distributed according to $P_{W, Z_j|T=\tau}$, which is the marginal of the joint distribution $P_{W, Z^m|T=\tau}$. In contrast, W and Z_j in the first term are conditionally independent random variables distributed as $(W, Z_j) \sim P_{W|T=\tau}P_{Z_j|T=\tau}$ where $P_{W|T=\tau}$ is the marginal distribution of $P_{W, Z_j|T=\tau}$. Now, fixing $X = W, Y = Z_j$ and $f(X, Y) = l(W, Z_j)$ so that $\mathbb{E}_{P_X P_Y} [f(\tilde{X}, \tilde{Y})] = \mathbb{E}_{P_{W|T=\tau}P_{Z_j|T=\tau}} [l(W, Z_j)]$ and $\mathbb{E}_{P_{X, Y}} [f(X, Y)] = \mathbb{E}_{P_{W, Z_j|T=\tau}} [l(W, Z_j)]$ in Lemma A1 under the assumption on $l(w, z)$ in Assumption 6(a), or in Lemma A2 under the assumption on $l(w, z)$ in Assumption 6(b) yields the following bound,

$$\left| \mathbb{E}_{P_{Z^m|T=\tau}P_{W|Z^m}} [\Delta L(W|Z^m, T = \tau)] \right| \leq \frac{1}{m} \sum_{j=1}^m \sqrt{2\delta_\tau^2 I(W; Z_j|T = \tau)}. \tag{A14}$$

Averaging with respect to P_T on both sides of (A14), and combining with the bound on average environment-level uncertainty yields the required bound in (33).

Appendix D. Details of Example

We first give details of the derivation of meta-generalization gap for the case with separate within-task training and test sets. The average meta-generalization loss can be computed as $\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_g^{\text{sep}}(U)] =$

$$\begin{aligned} & \mathbb{E}_{P_{Z_{1:N}^m, U}} \left[(1 - \alpha)^2 U^2 + \mathbb{E}_{P_{T, Z^{\text{mtr}}}} [\alpha^2 (D_T^{\text{mtr}})^2 + \mathbb{E}_{P_{Z|T}} [Z^2] - 2\alpha D_T^{\text{mtr}} \mu_T + 2(1 - \alpha)U(\alpha D_T^{\text{mtr}} - \mu_T)] \right] \\ & \stackrel{(a)}{=} \mathbb{E}_{P_{Z_{1:N}^m, U}} \left[(1 - \alpha)^2 (U^2 - 2U\mathbb{E}_{P_T} [\mu_T]) \right] + \mathbb{E}_{P_T} \left[\alpha^2 \left(\mu_T^2 + \frac{\mu_T \bar{\mu}_T}{m_{\text{tr}}} \right) + \mu_T - 2\alpha \mu_T^2 \right], \end{aligned} \tag{A15}$$

where the equality in (a) follows since $\mathbb{E}_{P_{Z|T}} [Z^2] = \mu_T, \mathbb{E}_{P_{Z^{\text{mtr}}|T}} [D_T^{\text{mtr}}] = \mu_T$ and $\mathbb{E}_{P_{Z^{\text{mtr}}|T}} [(D_T^{\text{mtr}})^2] = \mu_T^2 + \mu_T \bar{\mu}_T / m_{\text{tr}}$. In a similar manner, the average meta-training loss can be computed as

$$\begin{aligned} \mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_t^{\text{sep}}(U|Z_{1:N}^m)] &= \mathbb{E}_{P_{Z_{1:N}^m}} \left[-(1 - \alpha)^2 U^2 + \frac{1}{N} \sum_{i=1}^N \alpha^2 (D_i^{\text{mtr}})^2 \right. \\ & \quad \left. + \frac{1}{N} \sum_{i=1}^N \frac{1}{m_{\text{te}}} \sum_{j=1}^{m_{\text{te}}} (Z_{i,j}^{\text{mte}})^2 - 2\alpha \frac{1}{N} \sum_{i=1}^N D_i^{\text{mtr}} D_i^{\text{mte}} \right], \end{aligned} \tag{A16}$$

with U defined as in (38). The meta-generalization gap in (40) then results by taking the difference of (A15) and (A16), and using that $\mathbb{E}_{P_{Z_{1:N}^m}} [(1 - \alpha)^2 U^2] = \mathbb{E}_{P_T} [\mu_T \bar{\mu}_T] (\frac{1}{Nm_{te}} + \frac{\alpha^2}{Nm_{tr}}) + \frac{1}{N} (1 + \alpha^2) \text{Var}_T + (1 - \alpha)^2 (\mathbb{E}_{P_T} [\mu_T])^2$ and $\mathbb{E}_{P_{Z_{1:N}^m}} [U] = \mathbb{E}_{P_T} [\mu_T]$ with $\text{Var}_T = \mathbb{E}_{P_T} [\mu_T^2] - (\mathbb{E}_{P_T} [\mu_T])^2$.

We now evaluate the mutual pieces of information $I(U; Z_{1:N}^m)$ and $I(U; Z_i^m)$. For the first MI, note that, since the meta-learner is deterministic (see (38)), $H(U|Z_{1:N}^m) = 0$ and thus $I(U; Z_{1:N}^m) = H(U)$. For the second MI, we can write $I(U; Z_i^m) = H(U) - \mathbb{E}_{Z_i^m} [H(U|Z_i^m = z^m)]$. It can be seen that random variables U and $U|Z_i^m = z^m$ are mixtures of probability distributions, whose entropies can be evaluated following standard methods [54].

For the case with joint within-task training and test sets, the meta-generalization gap can be obtained in a similar way as

$$\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_t^{\text{joint}}(U|Z_{1:N}^m)] = \frac{2}{N} \left[\mathbb{E}_{P_T} \left[\frac{\mu_T \bar{\mu}_T}{m} \right] + \text{Var}_T \right] + \frac{2\alpha}{m} \mathbb{E}_{P_T} [\mu_T \bar{\mu}_T]. \tag{A17}$$

For the MI- and ITMI-based bounds, note that with $\mathcal{W} = [0, 1]$, the loss function $l(\cdot, \cdot)$ is $[0, 1]$ -bounded, and for the deterministic base-learner in (36) with $\mathcal{U} = [0, 1]$, the average training loss, $L_t^{\text{joint}}(u|Z^m)$ is also $[0, 1]$ -bounded for all $Z^m \in \mathcal{Z}^m$. Thus, Assumptions 5 and 6 hold with $\sigma^2 = \delta_t^2 = 1/4$. For the MI bound in (32), we have $I(U; Z_{1:N}^m) = H(U)$ and $I(W; Z^m|T = \tau) = H(W|T = \tau) - \mathbb{E}_{Z^m} [H(W|Z^m, T = \tau)]$. For the ITMI bound (33), we have

$$|\mathbb{E}_{P_{Z_{1:N}^m, U}} [\mathcal{L}_t^{\text{joint}}(U|Z_{1:N}^m)]| \leq \frac{1}{N} \sum_{i=1}^N \sqrt{\frac{1}{2} I(U; Z_i^m)} + \mathbb{E}_{P_T} \left[\frac{1}{m} \sum_{j=1}^m \sqrt{\frac{1}{2} I(W; Z_j|T = \tau)} \right]. \tag{A18}$$

All information measures can be easily evaluated numerically [54].

Appendix E. Proof of Lemma 4

From the update rule of the meta-learner in (43), we get the Markov dependency

$$P_{U^j|U^{(j-1)}, \{W_{K_i}\}_{i=1}^j, \{Z_{K_i}^m\}_{i=1}^j, Z_{1:N}^m} = P_{U^j|U^{j-1}, W_{K_j}, Z_{K_j}^m}, \tag{A19}$$

where $U^{(j-1)} = \{U^1, \dots, U^{j-1}\}$ is the history vector of hyperparameters. The sampling strategy in (44) together with (A19) then implies the following relation

$$P_{U^j|U^{(j-1)}, \{W_{K_i}\}_{i=1}^j, \{Z_{K_i}^m\}_{i=1}^j, Z_{1:N}^m} = P_{U^j|U^{j-1}, W_{K_j}, Z_{K_j}^m}. \tag{A20}$$

Using $U^{(j)} = \{U^1, \dots, U^j\}$ to denote the set of all updates, we have the following relations

$$I(U; Z_{1:N}^m) \stackrel{(a)}{\leq} I(U^{(j)}; Z_{1:N}^m) \stackrel{(b)}{\leq} I(U^{(j)}; \{Z_{K_i}^m\}_{i=1}^j) = \sum_{j=1}^J I(U^j; \{Z_{K_i}^m\}_{i=1}^j | U^{(j-1)}) \tag{A21}$$

$$\leq \sum_{j=1}^J I(U^j; \{Z_{K_i}^m\}_{i=1}^j, \{W_{K_i}\}_{i=1}^j | U^{(j-1)}) \tag{A22}$$

$$= \sum_{j=1}^J h(U^j|U^{(j-1)}) - h\left(U^j|U^{(j-1)}, \{Z_{K_i}^m\}_{i=1}^j, \{W_{K_i}\}_{i=1}^j\right) \tag{A23}$$

$$\stackrel{(c)}{=} \sum_{j=1}^J \left[h(U^j|U^{j-1}) - h(U^j|U^{j-1}, W_{K_j}, Z_{K_j}^m) \right], \tag{A24}$$

where, the inequality in (a) follows from data processing inequality on Markov chain $Z_{1:N}^m \rightarrow U^{(j)} \rightarrow U$; (b) follows from the Markov chain $Z_{1:N}^m \rightarrow \{Z_{K_i}^m\}_{i=1}^J \rightarrow U^{(j)}$; and the equality in (c) follows from $U^{(j-2)} \rightarrow U^{j-1} \rightarrow U^j$ and (A20). Finally, the computation of bound in (A24) follows similar to Lemma 5 in [27].

References

- Shalev-Shwartz, S.; Ben-David, S. *Understanding Machine Learning: From Theory to Algorithms*; Cambridge University Press: Cambridge, UK, 2014.
- Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2006.
- Simeone, O. A Brief Introduction to Machine Learning for Engineers. *Found. Trends Signal Process.* **2018**, *12*, 200–431. [\[CrossRef\]](#)
- Schmidhuber, J. Evolutionary Principles in Self-Referential Learning, or On Learning How to Learn: The Meta-meta-... Hook. Ph.D. Thesis, Technische Universität München, München, Germany, 1987.
- Thrun, S.; Pratt, L. Learning to Learn: Introduction and Overview. In *Learning to Learn*; Springer: Berlin/Heidelberg, Germany, 1998; pp. 3–17.
- Thrun, S. Is Learning the N-th Thing Any Easier Than Learning the First? In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*; MIT Press: Cambridge, MA, USA, 1996; pp. 640–646.
- Vilalta, R.; Drissi, Y. A Perspective View and Survey of Meta-Learning. *Artif. Intell. Rev.* **2002**, *18*, 77–95. [\[CrossRef\]](#)
- Simeone, O.; Park, S.; Kang, J. From Learning to Meta-Learning: Reduced Training Overhead and Complexity for Communication Systems. *arXiv* **2020**, arXiv:2001.01227.
- Kuzborskij, I.; Orabona, F. Fast rates by transferring from auxiliary hypotheses. *Mach. Learn.* **2017**, *106*, 171–195. [\[CrossRef\]](#)
- Kienzle, W.; Chellapilla, K. Personalized handwriting recognition via biased regularization. In *Proceedings of the 23rd International Conference on Machine Learning*, Pittsburgh, PA, USA, 25–29 June 2006; pp. 457–464.
- Denevi, G.; Ciliberto, C.; Grazi, R.; Pontil, M. Learning-to-learn stochastic gradient descent with biased regularization. *arXiv* **2019**, arXiv:1903.10399.
- Denevi, G.; Pontil, M.; Ciliberto, C. The advantage of conditional meta-learning for biased regularization and fine-tuning. *arXiv* **2020**, arXiv:2008.10857.
- Baxter, J. A Model of Inductive Bias Learning. *J. Artif. Intell. Res.* **2000**, *12*, 149–198. [\[CrossRef\]](#)
- Russo, D.; Zou, J. Controlling Bias in Adaptive Data Analysis Using Information Theory. In *Proceedings of the Artificial Intelligence and Statistics (AISTATS)*, Cadiz, Spain, 9–11 May 2016; pp. 1232–1240.
- Xu, H.; Farajtabar, M.; Zha, H. Learning Granger causality for Hawkes processes. In *Proceedings of the International Conference on Machine Learning*, New York City, NY, USA, 20–22 June 2016; pp. 1717–1726.
- Maurer, A. Algorithmic Stability and Meta-Learning. *J. Mach. Learn. Res.* **2005**, *6*, 967–994.
- Devroye, L.; Wagner, T. Distribution-Free Performance Bounds for Potential Function Rules. *IEEE Trans. Inf. Theory* **1979**, *25*, 601–604. [\[CrossRef\]](#)
- Rogers, W.H.; Wagner, T.J. A Finite Sample Distribution-Free Performance Bound for Local Discrimination Rules. *Ann. Stat.* **1978**, pp. 506–514. [\[CrossRef\]](#)
- Pentina, A.; Lampert, C. A PAC-Bayesian Bound for Lifelong Learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, Beijing, China, 21–26 June 2014; pp. 991–999.
- Amit, R.; Meir, R. Meta-Learning by Adjusting Priors Based on Extended PAC-Bayes Theory. In *Proceedings of the International Conference on Machine Learning (ICML)*, Beijing, China, 21–26 June 2018; pp. 205–214.
- Rothfuss, J.; Fortuin, V.; Krause, A. PACOH: Bayes-Optimal Meta-Learning with PAC-Guarantees. *arXiv* **2020**, arXiv:2002.05551.
- Denevi, G.; Ciliberto, C.; Stamos, D.; Pontil, M. Incremental learning-to-learn with statistical guarantees. *arXiv* **2018**, arXiv:1803.08089.
- Finn, C.; Abbeel, P.; Levine, S. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proceedings of the International Conference on Machine Learning—Volume 70*, Sydney, NSW, Australia, 6–11 August 2017; pp. 1126–1135.
- Nichol, A.; Achiam, J.; Schulman, J. On First-Order Meta-Learning Algorithms. *arXiv* **2018**, arXiv:1803.02999.
- Xu, A.; Raginsky, M. Information-Theoretic Analysis of Generalization Capability of Learning Algorithms. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, Long Beach, CA, USA, 4–9 December 2017; pp. 2524–2533.
- Bu, Y.; Zou, S.; Veeravalli, V.V. Tightening Mutual Information Based Bounds on Generalization Error. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Paris, France, 7–12 July 2019; pp. 587–591.
- Pensia, A.; Jog, V.; Loh, P.L. Generalization Error Bounds for Noisy, Iterative Algorithms. In *Proceedings of the 2018 IEEE International Symposium on Information Theory (ISIT)*, Veil, CO, USA, 17–22 June 2018; pp. 546–550.
- Vapnik, V.N.; Chervonenkis, A.Y. On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. In *Theory of Probability and Its Applications*; SIAM: Philadelphia, PA, USA, 1971; Volume 16, pp. 264–280.
- Koltchinskii, V.; Panchenko, D. Rademacher Processes and Bounding the Risk of Function Learning. In *High Dimensional Probability II*; Springer: Berlin/Heidelberg, Germany, 2000; Volume 47, pp. 443–457.
- Bousquet, O.; Elisseeff, A. Stability and Generalization. *J. Mach. Learn. Res.* **2002**, *2*, 499–526.

31. Kearns, M.; Ron, D. Algorithmic Stability and Sanity-Check Bounds for Leave-One-Out Cross-Validation. *Neural Comput.* **1999**, *11*, 1427–1453. [[CrossRef](#)]
32. Poggio, T.; Rifkin, R.; Mukherjee, S.; Niyogi, P. General Conditions for Predictivity In Learning Theory. *Nature* **2004**, *428*, 419–422. [[CrossRef](#)]
33. Kutin, S.; Niyogi, P. Almost-Everywhere Algorithmic Stability and Generalization Error. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI'02)*; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2002; pp. 275–282.
34. Dwork, C.; Feldman, V.; Hardt, M.; Pitassi, T.; Reingold, O.; Roth, A.L. Preserving Statistical Validity in Adaptive Data Analysis. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing (STOC)*, Portland, OR, USA, 15–17 June 2015; pp. 117–126.
35. Bassily, R.; Nissim, K.; Smith, A.; Steinke, T.; Stemmer, U.; Ullman, J. Algorithmic Stability for Adaptive Data Analysis. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing (STOC)*, Cambridge, MA, USA, 19–21 June 2016; pp. 1046–1059.
36. McAllester, D.A. PAC-Bayesian Model Averaging. In *Proceedings of the Twelfth Annual Conference on Computational Learning Theory (COLT)*, Santa Cruz, CA, USA, 7–9 July 1999; pp. 164–170.
37. Seeger, M. PAC-Bayesian Generalization Error Bounds for Gaussian Process Classification. *J. Mach. Learn. Res.* **2002**, *3*, 233–269.
38. Alquier, P.; Ridgway, J.; Chopin, N. On the Properties of Variational Approximations of Gibbs Posteriors. *J. Mach. Learn. Res.* **2016**, *17*, 8374–8414.
39. Asadi, A.; Abbe, E.; Verdú, S. Chaining Mutual Information and Tightening Generalization Bounds. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, Montreal, QC, Canada, 2–8 December 2018; pp. 7234–7243.
40. Negrea, J.; Haghifam, M.; Dziugaite, G.K.; Khisti, A.; Roy, D.M. Information-Theoretic Generalization Bounds for SGLD via Data-Dependent Estimates. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, Vancouver, BC, Canada, 8–14 December 2019; pp. 11013–11023.
41. Alabdulmohsin, I. Towards a Unified Theory of Learning and Information. *Entropy* **2020**, *22*, 438. [[CrossRef](#)]
42. Jiao, J.; Han, Y.; Weissman, T. Dependence Measures Bounding the Exploration Bias for General Measurements. In *Proceedings of the 2017 IEEE International Symposium on Information Theory (ISIT)*, Aachen, Germany, 25–30 June 2017; pp. 1475–1479.
43. Issa, I.; Gastpar, M. Computable Bounds On the Exploration Bias. In *Proceedings of the 2018 IEEE International Symposium on Information Theory (ISIT)*, Vail, CO, USA, 17–22 June 2018; pp. 576–580.
44. Issa, I.; Esposito, A.R.; Gastpar, M. Strengthened Information-Theoretic Bounds on the Generalization Error. In *Proceedings of the 2019 IEEE International Symposium on Information Theory (ISIT)*, Paris, France, 7–12 July 2019; pp. 582–586.
45. Steinke, T.; Zakyntinou, L. Reasoning About Generalization via Conditional Mutual Information. *arXiv* **2020**, arXiv:2001.09122.
46. Knoblauch, J.; Jewson, J.; Damoulas, T. Generalized Variational Inference. *arXiv* **2019**, arXiv:1904.02063.
47. Bissiri, P.G.; Holmes, C.C.; Walker, S.G. A General Framework for Updating Belief Distributions. *J. R. Stat. Soc. Ser. Stat. Methodol.* **2016**, *78*, 1103–1130. [[CrossRef](#)] [[PubMed](#)]
48. Yin, M.; Tucker, G.; Zhou, M.; Levine, S.; Finn, C. Meta-Learning Without Memorization. *arXiv* **2019**, arXiv:1912.03820.
49. Wainwright, M.J. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*; Cambridge University Press: Cambridge, UK, 2019; Volume 48.
50. Shalev-Shwartz, S.; Shamir, O.; Srebro, N.; Sridharan, K. Learnability, Stability and Uniform Convergence. *J. Mach. Learn. Res.* **2010**, *11*, 2635–2670.
51. Maurer, A.; Pontil, M.; Romera-Paredes, B. The Benefit of Multitask Representation Learning. *J. Mach. Learn. Res.* **2016**, *17*, 2853–2884.
52. Welling, M.; Teh, Y.W. Bayesian Learning via Stochastic Gradient Langevin Dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, Bellevue, WA, USA, 28 June–2 July 2011; pp. 681–688.
53. Boucheron, S.; Lugosi, G.; Massart, P. *Concentration Inequalities: A Nonasymptotic Theory of Independence*; Oxford University Press: Oxford, UK, 2013.
54. Michalowicz, J.V.; Nichols, J.M.; Bucholtz, F. *Handbook of Differential Entropy*; CRC Press: New York, NY, USA, 2013.