

## **Abstract**

Representation learning has been widely studied in the context of meta-learning, enabling rapid learning of new tasks through shared representations. Recent works such as MAML have explored using fine-tuning-based metrics, which measure the ease by which adaptation can achieve good performance, as proxies for obtaining representations. We present a theoretical framework for analyzing representations derived from a MAML-like algorithm, assuming the available tasks use *approximately* the same underlying representation. We then provide risk bounds on the best predictor found by fine-tuning via gradient descent, demonstrating that the algorithm can provably leverage the shared structure. The upper bound applies to general function classes, which we demonstrate by instantiating the guarantees of our framework in the logistic regression and neural network settings. In contrast, we establish the existence of settings where any algorithm, using a representation trained with no consideration for task-specific fine-tuning, performs as well as a learner with no access to source tasks. This separation result underscores the benefit of adaptation-based methods, such as MAML, over methods with “frozen representation” objectives in few-shot learning.

## **Reading List**

### **Textbooks**

- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint* (Vol. 48). Cambridge University Press.

### **Papers**

- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., & Darrell, T. (2014, January). Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning* (pp. 647-655). PMLR.
- Finn, C., Abbeel, P., & Levine, S. (2017, July). Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning* (pp. 1126-1135). PMLR.
- Baxter, J. (2000). A model of inductive bias learning. *Journal of artificial intelligence research*, 12, 149-198.
- Maurer, A., & Jaakkola, T. (2005). Algorithmic Stability and Meta-Learning. *Journal of Machine Learning Research*, 6(6).
- Maurer, A., Pontil, M., & Romera-Paredes, B. (2016). The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81), 1-32.
- Du, S. S., Hu, W., Kakade, S. M., Lee, J. D., & Lei, Q. (2020). Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*.
- Tripuraneni, N., Jordan, M., & Jin, C. (2020). On the Theory of Transfer Learning: The Importance of Task Diversity. *Advances in Neural Information Processing Systems*, 33.
- Tripuraneni, N., Jin, C., & Jordan, M. I. (2020). Provable meta-learning of linear representations. *arXiv preprint arXiv:2002.11684*.
- Hu, J., Chen, X., Jin, C., Li, L., & Wang, L. (2021). Near-optimal Representation Learning for Linear Bandits and Linear RL. *arXiv preprint arXiv:2102.04132*.

- Denevi, G., Pontil, M., & Ciliberto, C. (2020). The Advantage of Conditional Meta-Learning for Biased Regularization and Fine Tuning. *Advances in Neural Information Processing Systems*, 33.
- Denevi, G., Pontil, M., & Ciliberto, C. (2021). Conditional Meta-Learning of Linear Representations. *arXiv preprint arXiv:2103.16277*.