Meta-Learning a Solution to the Hanabi Ad-Hoc Challenge

Aron Sarmasi
University of California at Davis
Davis, USA
asarmasi@ucdavis.edu

Timothy Zhang
University of California at Davis
Davis, USA
mtzh@ucdavis.edu

Chu-Hung Cheng
University of California at Davis
Davis, USA
chccheng@ucdavis.edu

Huyen Pham
University of California at Davis
Davis, USA
hdupham@ucdavis.edu

Xuanchen Zhou
University of California at Davis
Davis, USA
xuczhou@ucdavis.edu

Duong Nguyen
University of California at Davis
Davis, USA
mdnnguyen@ucdavis.edu

Soumil Shekdar
University of California at Davis
Davis, USA
sshekdar@ucdavis.edu

Joshua McCoy
University of California at Davis
Davis, USA
jamccoy@ucdavis.edu

ABSTRACT

In this work we demonstrate that the First Order Model Agnostic Meta Learning (FOMAML) algorithm trained on the Hanabi Open Agent Dataset (HOAD) results in a model that is able to outplay both a naive MLP baseline, as well as a randomly selected partner in the Hanabi Ad-Hoc Challenge, in both low-shot and zero-shot setups. We first show that HOAD is well suited for the meta-learning task because its agents are high quality and utilize diverse strategies, thereby confirming that MAML is generalizing, and not memorizing agent strategies. We then detail our application of FOMAML to the cooperative decision making problem Hanabi entails, and we also provide evidence supporting recent results that the task update of MAML gives little to no test time performance boost. The pretrained models and game data are made available online at <https://github.com/aronsar/hoad>.

KEYWORDS

Hanabi, dataset, meta-learning

ACM Reference Format:

Aron Sarmasi, Timothy Zhang, Chu-Hung Cheng, Huyen Pham, Xuanchen Zhou, Duong Nguyen, Soumil Shekdar, and Joshua McCoy. 2021. Meta-Learning a Solution to the Hanabi Ad-Hoc Challenge. In *The 16th International Conference on the Foundations of Digital Games (FDG) 2021 (FDG'21)*, August 3–6, 2021, Montreal, QC, Canada. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3472538.3472546>

This work is licensed under a Creative Commons Attribution International 4.0 License.

FDG'21, August 3–6, 2021, Montreal, QC, Canada
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8422-3/21/08.
<https://doi.org/10.1145/3472538.3472546>

1 INTRODUCTION

Hanabi [5], a tabletop card game for 2-5 players¹, is a challenging application for agent modeling in the multi-agent domain due to its unique combination of partial observability, cooperation, stochasticity, and implicit communication. Recently proposed as a challenge domain [4] to provide a sophisticated yet well-defined set of challenges for artificial intelligence practitioners, Hanabi is of rising interest to the research community. Of particular interest is the Ad-Hoc Challenge, [4] which aims to push research towards agents that are capable of cooperatively interacting with other agents or even humans with whom they are unfamiliar. These problems inherent to Hanabi have also been studied in other contexts, both commercial and industrial, such as self-driving cars, and human-robot cooperation in the manufacturing setting.

In a game of Hanabi, the players see the cards of all their teammates, but not their own. Thus, the defining element of the game is the communication of information about other players' cards either explicitly using hints (a limited resource), or implicitly through playing cards, discarding, or even hinting (e.g. the finesse play [37]). Because of this, a group's success or failure depends on that group's ability to communicate and to understand one another—e.g. if one player makes a finesse play, and another player doesn't pick up on it, the group loses one of their three lives. The result of this mechanic is that unlike other games such as Go [30] and StarCraft[20], Hanabi has a multiplicity of optimal and near-optimal strategies, each corresponding to a different communication schema. While a number of works have focused on creating agents that get high scores in self-play [22] [34] [8], here we instead focus on the more difficult Ad-hoc Challenge, which, to the best of our knowledge, has not yet been studied. In the Ad-hoc Challenge, an agent is asked to play successfully with a held-out (ie. previously unseen) agent after observing only 10 of its self-play games. We break this task down into two subtasks: 1) given 10 of an agent's games, build a model of that agent, and 2) use this model to search over possible future states and pick an action that maximizes the expectation of future

¹in this work we consider only the two-player scenario for simplicity, leaving 3-5 players to future work

reward. In this work we focus on the first of these two subtasks, leaving the second to future work.

Building a model of an agent-based on only 10 of its games is a very difficult task for traditional deep learning algorithms because the 10 games do not provide nearly enough data to cover every aspect of an agent’s strategy, and they also don’t necessarily provide an example for every game mechanic. For example, some agents in Hanabi Open Agent Dataset [29] (HOAD) rarely or never give hints, while other agents may be too weak to reach late game, and so may not provide examples of late-game mechanics in action. As a lower bound, we present a naive multilayer perceptron (an MLP) which does exactly this: it is trained to imitate an agent based on only 10 of its self-play games.

Properly addressing the Ad-Hoc Challenge requires an *a priori* understanding of the game, and of the kind of communication strategies one might employ to play the game. The approach we take in this work is to learn such an understanding using a large number of games played by a variety of different agents (that we call “training agents”). To the best of our knowledge, HOAD is the only dataset of Hanabi playing agents that can facilitate this kind of learning. HOAD is a compilation of different agents from different sources, most using conventions typically seen in human play, and all operating within the same environment and using a binarized game state representation specifically designed for neural learning. The dataset allows pairwise play between any two agents, offers an easy way to add additional agents from most codebases, and also gives access to the Dopamine reinforcement learning framework [7], which we plan to use in future work to further improve and extend the HOAD agents.

To solve the Ad-hoc Challenge, we chose to implement the first-order approximation of Model Agnostic Meta-Learning (FOMAML) [11]. This is a highly generalizable technique that is well-suited to the low-shot learning scenario and has been successfully applied to several problems outside of image recognition, its original application [12] [36]. Despite its popularity however, this is the first time it is applied to the cooperative decision making problem. The goal of MAML is to train a meta-model that is close in parameter space to the optima of many different tasks. During inference, a few gradient updates (hence “low-shot”) are all that is needed to obtain good performance on a previously unseen task. In our case, that means that given just a few games of a held-out agent, the meta-model can be updated to accurately imitate this agent. Recently, however, it has been shown that the meta-model itself is already quite strong on its own, without the task update [25]. Our results provide further evidence of this phenomenon, showing that it also occurs in the context of Hanabi, and using the first-order approximation of MAML. Pseudocode, along with a more detailed explanation of our implementation is provided in Section 5.2.

The main contribution of this work is that we show that MAML can be applied to the cooperative decision making problem to train a meta-model on HOAD, which can then be updated on a small set of a held-out agent’s games, and then played directly with that held-out agent to achieve superior performance to random partner selection and a naive MLP baseline. We also show that given enough training time, the meta-model actually performs better without the task update step, corroborating similar recent findings.

2 RELATED WORKS

Agent imitation is a well-established field of research; some early methods include behavioral cloning [27], which is quite brittle and does not easily adapt to new situations, and inverse reinforcement learning, which is very slow to run [26]. Despite increased attention to this field in the past decade [1] [3], few works study agent imitation under all the same constraints that apply to Hanabi—i.e. partial observability, cooperation, stochasticity, and implicit communication, all in the low-shot regime. As such, we discuss those that adopt a subset of these constraints. He et. al. [15] presents models that jointly learn a policy and the behavior of their opponents using deep reinforcement learning; however, they need access to the agent they model during training, and so are also not well suited to the ad-hoc situation. Le et. al. [18] proposes an approach to learning implicit coordination strategies that jointly learns a latent coordination model as well as the agents’ individual policies, but this approach too is infeasible in the ad-hoc case, as it requires a large number of examples of agent interactions.

Several other works share our low-shot constraint but are overly specialized to their respective tasks. A seminal work in robotic agent imitation, Generative Adversarial Imitation Learning (GAIL), makes use of a generative adversarial technique that achieves good results on low-shot expert imitation; however, GAIL (and its class of solutions) is best suited to the continuous world of robotic movement, and so is difficult to adapt to Hanabi [16] [35] [19]. Duan et. al. [9] proposes one-shot imitation learning to solve the task of robot block-stacking, but their neighborhood attention model is ill-suited to the Hanabi game representation, where there is little spatial information. Similarly, [13] [33] [32] [31] are all designed for the low-shot scenario, but their hand-crafted architectures are too focused on the image classification or robot motion applications to be easily adapted to Hanabi.

The most comparable approach to MAML in terms of generality is the Simple Neural Attentive Learner (SNAIL) [21]; indeed, SNAIL performs comparably or slightly better than MAML when evaluated on Omniglot [17] and Mini-Imagenet [33], and we suspect that it could be a promising baseline to run on HOAD in future work. For this work, however, we preferred MAML’s model agnosticism, which allowed us to use a small and simple fully connected network as the policy predictor. In a similar vein, we chose not to use any of the many MAML derivative works, many of which achieve improvements by introducing the kind of hand-designed architectural and algorithmic tweaks that specialize the approach to a specific domain, and which don’t transfer well to Hanabi [19] [14] [28]. Of additional interest is Yu and Finn’s more recent work using MAML for human imitation in the robot manipulation environment; this work most closely resembles the conceptual underpinnings of our use of MAML [36].

3 RULES OF THE GAME

The base game considered here consists of the following materials:

- 50 cards; 10 each of blue, yellow, red, white, and green; of each color there are three 1 valued cards, two 2’s, two 3’s, two 4’s, and only one 5 valued card
- 8 hint tokens, also referred to as information tokens; these have two faces to indicate whether they’ve been used already

- 3 lives, also referred to as mulligans or danger tokens; these too have two faces to designate their availability

The goal of the game is for the players to collectively form 5 piles of cards, one pile of each color. Each pile must be built in ascending order, starting with the 1 valued card on the bottom, and ending with the 5 valued card on top. At the start of the game, each player is dealt a hand of 5 cards, and the 8 hint tokens and 3 life tokens are all facing up. During play, a player may look at the cards of their teammates, but never their own. Players take legal actions in turns, going around the table in a clockwise fashion. The game ends when any of the below conditions occur:

- If the third life token is flipped face down, the game ends with a score of 0 (this is the most important factor in why it is difficult to play with unfamiliar agents)
- If the 5 piles are completed before the deck runs out, the game ends with a score of 25
- If the deck runs out before the 5 piles are completed, each player gets one last turn, including the one who drew the last card. The final score is equal to the count of the cards in the 5 piles.

According to the official ruleset, no communication may occur during the game other than through the use of hint tokens (explained below). However, the rules also recommend that players interpret this rule in a way that suits them best. Although there is some controversy in the research community over what exactly should be considered acceptable communication [10], in this work we have restricted the agents to only be able to pass information to their teammates through their choice of legal action on their own turn—the same as proposed in [4]. The possible actions in a game are:

- Giving a hint: a player can either give a color hint or a value hint. To give a color hint, pick another player, pick a color, and then point out all cards of that color in their hand. To give a value hint, do the same, but for a card value instead of a color. Giving a hint costs a hint token, and is only a legal action if there are remaining face up hint tokens. No empty hints may be given, i.e. one may not tell a player what is not in their hand.
- Playing a card: a player may choose a card in their hand and play it. To play a card, attempt to place it on one of the five piles. If successful, draw another card from the deck and pass the turn. If the card can’t currently be placed on any of the piles, discard the card, flip a life token face down, and draw another card from the deck. Note: if a 5 is successfully played, flip a hint token face up if possible.
- Discarding a card: a player can choose a card in their hand and place it in the discard pile, flipping a hint token face up, and then drawing a card from the deck (note: players may look through the discard pile at any time). This is only a legal action if there are 7 or fewer hint tokens face up.

4 THE HANABI OPEN AGENT DATASET

The Hanabi Open Agent Dataset consists of Hanabi playing agents that have codebases available online, score reasonably well in self-play (above 10 points), and employ strategies sufficiently different from one another, such that when two different agents play one

Table 1: The agents in HOAD are MLP imitators, which were trained using 500,000 games of each original agent’s game data. The imitators generally score well in self-play, and mimic the original agents with high accuracy.

Agent	Original Self-Play Score	Imitator Self-Play Score	Imitator Accuracy
Simplebot	16.9	16.8	99.7
Valuebot	19.8	18.0	92.0
Holmesbot	20.8	14.7	90.3
Outer	14.5	14.1	66.7
Iggi	17.0	16.2	90.9
Piers	17.3	15.9	85.8
Rainbow	18.5	18.1	77.5
Van-Den-Bergh	14.0	10.5	81.2

another they perform significantly worse than either does in self-play. As a result, the variation of strategy among the agents is ensured. For this reason, some existing agents are excluded from the dataset. These include a number of derivative works that were slight modifications on existing agents, as well as the Actor Critic Hanabi Agent (ACHA) [22], which does not have a public codebase, and Caanan’s agents [6], released concurrently to HOAD.

Canaan’s agents specifically bear further discussion, as they too are motivated by the Hanabi Ad-Hoc Challenge. These rule-based, genetically evolved agents are created using MAP-Elites [23], a Quality Diversity algorithm that parameterizes the agents’ risk aversion and communicativeness, allowing the user to generate agents with a wide array of different strategies. The approach isn’t a panacea, however. Agents close in parameter space overlap significantly in action distribution, and achieve high scores in cross-play. Additionally, the current implementation of the approach doesn’t explore the full space of possible Hanabi strategies (focusing only on risk aversion and communicativeness). Nevertheless, an immediate next step for the HOAD project should be to add support for Canaan’s agents, as the ability to create agents in a structured, parameterized way is very useful in the context of the Ad-hoc Challenge.

One of the advantages of using HOAD is its ease of extensibility. Although some authors of Hanabi playing agents have published multiple agents using the same framework, in general, two arbitrary agents taken from different codebases will not be able to play one another because of differences in implementation in how the agents represent the game state. To circumvent the burden of reimplementing every new agent, the authors of HOAD observe the starting deck order and the actions taken by agents in the native environments, and then use those observations to recreate the games in the Hanabi Learning Environment (HLE) (Section 4.2). Observing the starting deck order and the actions taken is a much simpler task because the set of legal actions is small (≤ 20 in a 2 player game), and the ordering of the deck is typically known by the respective game engine at the start of the game. Once replay data has been gathered for all the agents in the HLE representation, it is possible to train a neural network to imitate each of the original agents. The self-play scores of these imitators, as well as their accuracies (ie.

how well the actions of the imitators match up with the actions of the original agents), are presented in table 1.

4.1 Implementation Details

The workaround technique, described above, was used to generate 500,000 games per original agent, all in the HLE representation. These games were then used to train imitator agents for each original agent. The imitator architecture consisted of: layer sizes of 1024, 1024, 512, 512, 512, and 256, batch normalization at every layer, and a dropout rate of 0.5. Each network was trained for 5 – 10 epochs, until convergence.

4.2 Hanabi Learning Environment

The Hanabi Learning Environment (HLE) [4] is a framework published by DeepMind in 2019 consisting of a fast gameplay engine implemented in C, an easy to use Python API, and integration with the Dopamine reinforcement learning framework. During play, the game state (as seen by the observing agent without revealing hand information) can be accessed as a binarized vector. This is the representation used by all the imitator agents in HOAD, and all the ad-hoc agents described in this work. We use the Dopamine framework only to reproduce the Rainbow agent, although it does present an interesting opportunity for future research, specifically to improve and extend the other HOAD agents.

4.3 HOAD Agent Strategy Summaries

We present a summary of each agent in HOAD, including only necessary detail, and also offering explanation of the imitation accuracies and self-play scores in Table 1. A compilation of commonly employed strategies, used both by human players and HOAD agents, can be found at [37].

Simplebot [24] – Plays only cards which it has enough knowledge about to know they are playable, and gives hints only about playable cards, preferring color hints over value hints. It uses an oldest first discarding strategy. Due to its simplicity and high degree of determinism, Simplebot is easy to imitate.

Valuebot [24] – Same as Simplebot, but before playing, checks to see if the next player is about to discard a valuable card (i.e. the last copy of a card). Also, it is deterministic, and easy to imitate.

Holmesbot [24] – Extends Valuebot by including the use of mulligans and by adding additional inference capabilities. Specifically, card knowledge from hints, the discard pile, and other players' hands are used to make deductions about the agent's hand. Still deterministic, and easy to imitate because all the information it uses is present in the HLE representation.

Iggi [34] – Similar logic to Simplebot but prefers value hints over color hints and prefers discarding unplayable cards over oldest first. Highly deterministic, easy to imitate.

Outer [34] – Similar to Iggi, but prefers discarding over hinting, and uses more randomness in its hinting and discarding logic; this results in significantly reduced imitation accuracy and is also likely the reason for lower published score.

Piers [34] – Extends Iggi by including the use of mulligans (but not as deterministically as Holmesbot), and some additional logic to avoid discarding valuable cards. The lower determinism explains the slight drop in imitation accuracy.

Table 2: Pairwise play scores are produced by playing each MLP imitator agent with every other imitator agent for 500 games. For each combination of agents, we test both possibilities for which agent makes the first move. In the lower left diagonal, the agents in the row labels made the first move, and vice versa.

	Simplebot	Valuebot	Holmesbot	Outer	Iggi	Piers	Rainbow	Van-Den-Bergh
Simplebot	16.8	15.7	12.8	0.0	4.1	1.1	1.7	0.2
Valuebot	15.2	18.0	17.6	0.0	3.8	1.3	2.0	0.0
Holmesbot	11.2	18.3	14.7	0.0	1.4	0.6	0.5	0.0
Outer	0.0	0.0	0.0	14.1	1.0	4.1	6.2	9.0
Iggi	3.9	3.8	1.8	1.8	16.2	11.8	2.7	6.0
Piers	1.8	0.5	0.2	7.2	10.3	15.9	5.5	9.4
Rainbow	0.3	2.0	0.3	6.6	4.0	5.6	18.1	2.6
Van-Den-Bergh	0.0	0.2	0.0	10.6	4.7	8.3	2.7	10.5

Van-der-bergh [34] – Makes some risky plays if they have high likelihood of success and there are remaining mulligans. Prioritizes discarding over hints, gives hints about useless cards, and attempts to maximize transmitted information.

Rainbow [4] – This agent tends to hint for rank instead of color. Conditional action probabilities may be found in the appendix of [4].

4.4 Discussion of Imitator Agent Pairwise Scores

Since the same observation can be passed in to any of the imitator agents, the imitator agents make it possible to play games with agents originating from two different codebases; this is necessary to evaluate how different the players are from each other. Table 2 shows the average score of each imitator agent playing 1000 games with every other imitator agent. As expected, when agents are paired with themselves, they achieve scores much higher than when paired with each other—this corresponds to the high scores on the diagonals and the relatively low scores on the off-diagonals. This confirms the intuition that agents must use a similar strategy when paired with one another, else risk miscommunicating, losing all three lives, and scoring zero points.

An interesting feature of Table 2 warranting discussion is the high scores achieved by certain combinations of agents. Two groups of agents that play relatively successfully with each other are O'Dwyer's Simplebot – Valuebot – Holmesbot, and Walton-Rivers' Iggi – Piers – Outer – Van-der-bergh. All the agents in each of these two groups belong to the same codebase, so the high scores are likely due to the authors of the two codebases reusing logic between their agents. The exception to this is the Rainbow agent, which was produced using reinforcement learning, and so it is surprising that it performs so well with Walton-Rivers' agents. Our best explanation is that both agents reportedly prefer value hints and can presumably also respond well to game states where value hints have been given. Even more interestingly, some agent combinations

(e.g. Holmesbot – Valuebot) result in scores greater than either agent can achieve in self-play. We surmise this is a result of synergy between the two agents' strategies. Although such synergies could cast doubt on the ability of pairwise scores to tell whether agents' strategies are different, we argue that they are rare, and thus have an overall small effect on the trustworthiness of using the pairwise scores as a validation metric.

Finally, we consider the phenomenon that for some agent combinations, it matters significantly who goes first. Specifically, for the combinations Simplebot–Holmesbot, Outer–Piers, and Outer–Vanden-bergh, if the first of the listed agents goes first, the combination gets around 2 points more on average than if the second listed agent goes first. This effect too is likely due to synergies between the agents, and additionally serves to show that in general, synergies have a fairly small effect on the final score.

5 AD-HOC AGENT MODELING

In this work, we propose to apply MAML to the cooperative decision making problem of the ad-hoc challenge. The goal of this algorithm is to train a set of model parameters that are a single gradient update away from good generalization performance on a new learning task (imitating the strategy of a previously unseen agent), using only a small number of training examples—just 10 games in this case. Intuitively, the training process finds a meta-optimum that is close in parameter space to many other task-specific optima. Since MAML is model agnostic—meaning that the algorithm can be applied to any model that trains via backpropagation—we had to make a design decision on what model to choose. We decided not to use a recurrent neural network since none of the original agents consider a game's action history when making moves; rather they take moves considering only the information available in the current game state, as represented by the HLE. Preferring simplicity, we decided to make the MAML model a fully connected network.

We also train a set of "naive" MLPs, so called because they only use the test agent's 10 games to learn its strategy. This isn't meant to be a fair comparison to MAML, but simply a lower bound to demonstrate that it is necessary to use data from other agents to build up an inductive bias before tackling the Ad-hoc Challenge. As a point of clarity, the naive MLPs and the imitator MLPs have different architectures; the naive MLPs have much fewer parameters, because they are trained on so little data. Please refer to the end of section 5.2 for architecture details.

5.1 Summary of Experiments

First, the agent imitations from section 4 were used to generate 500,000 games per agent. Then, for each held-out agent, the MAML algorithm was run for 5 epochs on all the HOAD agents' games except for those of the held-out one, resulting in eight different meta-models. A checkpoint was taken after just 2 epochs for each meta-model.

Evaluation occurred in two scenarios: low-shot and zero-shot. The only difference between the two is that in the low-shot scenario each of the above models was updated with 10 games of their respective held-out agents, whereas this was not done in the zero-shot scenario. The now agent-specific models (or just the meta-models in the zero-shot scenario) were then played with their respective

held-out agent imitations for 50 games each. We then repeated the above evaluation process 10 times, averaging scores across both games and trials.

Each naive MLP, on the other hand, was trained on only 10 games from each agent. The resulting model was then played directly with the corresponding agent imitator for 50 games each. This evaluation process too was repeated 10 times (including retraining each MLP on different games each time), and the scores averaged across both games and trials, for 500 total evaluation games per agent. The results for both of the above experiments may be found in table 3.

Algorithm 1 FOMAML for HOAD

Require: M self-play games from each HOAD agent
Require: a differentiable model to use with FOMAML

Definitions:

k : number of agents to sample
 d : number of ad-hoc games
 α, β : task and meta learning rates
 $\mathcal{L}$: cross entropy loss
 θ : parameters of meta-model
 f_θ : output of forward pass of meta-model

end Definitions:

```

Randomly initialize  $\theta$ 
while not done do
  sample  $k$  agents
  for each agent  $i$  do
    support  $\leftarrow$  sample  $d$  games from agent  $i$ 
    Compute  $\nabla_\theta \mathcal{L}(f_\theta)$  using support
    Compute adapted model with gradient descent:
     $\theta'_i \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}(f_\theta)$ 
    query  $\leftarrow$  sample  $d$  more games from agent  $i$ 
    Compute  $\nabla_{\theta'_i} \mathcal{L}(f_{\theta'_i})$  using query
    Update meta-model using first order approximation:
     $\theta \leftarrow \theta - \beta \nabla_{\theta'_i} \mathcal{L}(f_{\theta'_i})$ 
  end for
end while

```

5.2 Implementation Details

The model used with MAML consists of 12 fully connected layers with hidden sizes 2048, 2048, 1024, 1024, 512, 512, 256, 256, 128, 128, 64, and 64. The model input is the binary vector representation of the game state for 2 players (length 658), as defined by the HLE, and the output is the argmax of the logits layer, which is of length 20 since there are 20 possible actions in a 2 player game. We use leaky ReLU with a negative slope of -0.3, batch normalization at each layer, and skip connections every other layer for increased performance.

In order to model previously unseen agents, it is necessary to modify the experimental setup typically used with MAML [33]. Typically, an N-way K-shot setup is used, where N unseen classes are selected and K examples of each class are provided, and then the model is expected to successfully classify among the N classes. In our implementation, we instead select one unseen agent, update the

Table 3: Ad-hoc play scores for randomly selected partners, the naive MLP, and MAML methods. The naive MLP is trained on the 10 ad-hoc games of each held-out agent. The MAML models are trained on 500,000 games per each of the other seven HOAD agents (3.5M games total), and in the low-shot scenario are also updated with 10 ad-hoc games from the held-out agent. Evaluation is done by playing 500 games with the held-out agent.

Held-out Agent	Random Partner	Naive MLP	MAML 2 epochs		MAML 5 epochs	
			Zero Shot	Low Shot	Zero Shot	Low Shot
Simplebot	5.1	1.1	4.2	5.3	4.6	3.8
Valuebot	5.7	1.5	5.3	7.6	8.4	8.9
Holmesbot	4.6	0.5	6.2	6.8	9.4	7.9
Outer	2.9	2.5	3.0	3.6	3.4	3.9
Iggi	4.5	2.5	5.9	11.1	6.2	8.1
Piers	5.0	2.0	9.3	8.3	11.0	10.2
Rainbow	3.1	1.8	2.9	1.8	3.1	2.6
Van-Den-Bergh	3.8	2.2	4.5	4.7	6.7	5.4
Averages	4.3	1.8	5.2	6.2	6.6	6.3

model with 10 of that agent’s games (which is anywhere from 40 – 800 observation-action pairs depending on the length of the game), and then give the model an observation, and ask it to predict the action that the unseen agent would have taken. The technique and intuition behind MAML (i.e. training to find the meta-optimum) still hold; we have simply changed the objective.

The most significant difference between our implementation and the original MAML is that we’re using the first-order approximation (FOMAML) described in [2] and [11]. In practice, the only difference between it and the original MAML is that the meta-model is updated with $\nabla_{\theta'_i} \mathcal{L}(f_{\theta'_i})$ instead of $\nabla_{\theta} \mathcal{L}(f_{\theta'_i})$ (please refer to our reproduction of FOMAML in Algorithm 1 for clarification on this notation). During each meta training step, 4 training agents are sampled. From each agent, we sample 10 games—this is the support data. Each agent’s support data is used to make a gradient update with the SGD optimizer and a learning rate of 0.0001 (the objective is the imitation task); thus, we get an adapted model for each agent. Next, 10 more games are sampled from each agent—this is the query data. We now find the loss using the adapted parameters and the query data, and then we compute the gradient with respect to this loss and the adapted parameters. This is the point where FOMAML diverges from MAML; in MAML, we would instead have computed the gradient with respect to this loss and the original parameters. Once we have computed this gradient, we use it to update the meta-model. As in the original MAML, this procedure optimizes the meta-model parameters such that a single gradient update on a new agent’s games will produce a maximally accurate imitation of that agent. Finally, every 10,000 episodes the meta-learning rate is halved, until convergence.

The naive MLPs are a fully connected network with layer sizes 512, 512, 512, and 256, batch normalization, dropout = 0.5, and Adam optimizer with learning rate 1.5e-4. These parameters gave the best results after a thorough hyperparameter search.

5.3 Discussion

The first thing we notice in Table 3 is that the naive MLP gets fairly low scores across the board; this backs up our intuition that it is

very difficult to learn the mechanics of the game (not to mention the specific strategy of the held out agent) from just 10 games of data. Comparing these scores to the case where we play the held-out agent with a random other HOAD agent, we can see that at least being familiar with the game mechanics helps to greatly boost the score, by more than a factor of 2. Playing the held-out agent with random HOAD agents also provides a good baseline for how similar the strategies of the various agents are. Whereas in self-play the average HOAD agent achieves around 15.5 points, paired with a randomly selected agent, a held-out agent can only expect to earn about 4.3 points. This baseline is also useful when measuring the performance of an ad-hoc learner—if the learner cannot outperform a randomly selected partner, then it is not a very competitive learner. In this case, we can see that the MAML approach outperformed the randomly selected partner in both the zero and low shot scenarios, by around 2 points.

An important caveat that bears mentioning is that the MAML results are affected to an unknown extent by the fact that some of the agents are drawn from the same codebase and share some conventions. It is possible that some of the improvement seen in the MAML trials is due to agents copying conventions from other agents in the same family. Because diversity is an important aspect of the HOAD dataset, an essential piece of future work will be to develop a more principled approach to measuring similarity between bots—one that accounts for in-family similarities as well.

The most interesting result among the MAML runs is how the effect of adapting the parameters with the ad-hoc games changes as the model converges. Namely, in the early stages of training, the adaptation step seems to help, gaining about 1 point on average. However, once the model is converged, adapting to the held-out agent has no significant effect on performance (and indeed, seems to have a slightly deleterious one). This goes against the results of Raghu, who found that inner loop updates have negligible effect on learned representation at both early and late stages of training [25]. There are, however, some important differences between our work and theirs that may explain this discrepancy—namely, they performed this test using second-order MAML on Omniglot

and MiniImageNet, whereas we use first-order MAML on Hanabi. Raghu does test first-order MAML as well, but only to compare the computational speed of FOMAML to their own algorithm ANIL (Almost No Inner Loop). They show that nearly eliminating the inner loop of MAML can speed up computation by about as much as using first-order MAML. They do not, however, show whether FOMAML benefits from an inner loop in the early stages of training, which in our case, it does.

6 CONCLUSION

In this work we demonstrated that the FOMAML algorithm can be applied to the Hanabi Ad-Hoc Challenge and the cooperative decision making problem it entails. We also argued that HOAD is a good benchmark for meta-learning based approaches, because of the quality and diversity of its agents. When FOMAML is trained on HOAD, the resulting meta-model was able to outperform both a naive MLP baseline, and a randomly selected partner. Furthermore, we showed that after a sufficient number of training steps, the meta-model does not need to be updated on the 10 games of the held-out agent, and will actually perform better without it. This backs up recent evidence that MAML-type algorithms don't benefit from the task update step during test-time evaluation.

ACKNOWLEDGMENTS

To Robert, for the bagels and explaining CMYK and color spaces.

REFERENCES

- [1] Stefano V. Albrecht and Peter Stone. 2018. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence* 258 (may 2018), 66–95. <https://doi.org/10.1016/j.artint.2018.01.002> arXiv:1709.08071
- [2] Antreas Antoniou, Harrison Edwards, and Amos Storkey. 2019. How to train your MAML. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=HJGven05Y7>
- [3] Alexandre Attia and Sharone Dayan. 2018. unpublished. *Global overview of Imitation Learning*. Technical Report. arXiv:1801.06503 <http://arxiv.org/abs/1801.06503>
- [4] Nolan Bard, Jakob N. Foerster, Sarath Chandar, Neil Burch, Marc Lanctot, H. Francis Song, Emilio Parisotto, Vincent Dumoulin, Subhdeep Moitra, Edward Hughes, Iain Dunning, Shibli Mourad, Hugo Larochelle, Marc G. Bellemare, and Michael Bowling. 2019. *The Hanabi Challenge: A New Frontier for AI Research*. Technical Report. arXiv:1902.00506 <http://arxiv.org/abs/1902.00506>
- [5] Antoine Bauza. 2010. Hanabi. <https://www.boardgamegeek.com/boardgame/98778/hanabi>
- [6] Rodrigo Cnaan, Julian Togelius, Andy Nealen, and Stefan Menzel. 2019. Diverse Agents for Ad-Hoc Cooperation in Hanabi. Institute of Electrical and Electronics Engineers (IEEE), 1–8. <https://doi.org/10.1109/cig.2019.8847944> arXiv:1907.03840
- [7] Pablo Samuel Castro, Subhdeep Moitra, Carles Gelada, Saurabh Kumar, Marc G. Bellemare, and Google Brain. 2018. Dopamine: A Research Framework For Deep Reinforcement Learning. In *unpublished*. arXiv:1812.06110v1 <https://github.com/google/dopamine>
- [8] Christopher Cox, Jessica De Silva, Philip Deorsey, Franklin H.J. Kenter, Troy Retter, and Josh Tobin. 2015. How to make the perfect fireworks display: Two strategies for Hanabi. *Mathematics Magazine* 88, 5 (2015), 323–336. <https://doi.org/10.4169/math.mag.88.5.323>
- [9] Yan Duan, Marcin Andrychowicz, Bradly Stadie, Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. 2017. One-Shot Imitation Learning. In *NIPS*.
- [10] Markus Eger and Daniel Gruss. 2019. Wait a Second: Playing Hanabi without Giving Hints. (2019). <https://doi.org/10.1145/3337722>
- [11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *ICML*. arXiv:1703.03400v3
- [12] Chelsea Finn, Tianhe Yu, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. 2017. One-Shot Visual Imitation Learning via Meta-Learning. In *CoRL*. arXiv:1709.04905v1
- [13] Victor Garcia and Joan Bruna. 2018. Few-Shot Learning With Graph Neural Net-Works. In *ICLR*. arXiv:1711.04043v3
- [14] Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. 2018. Recasting Gradient-Based Meta-Learning as Hierarchical Bayes. In *ICLR*. arXiv:1801.08930v1
- [15] He He, Jordan Boyd-Graber, Kevin Kwok, and Hal Daumé. 2016. Opponent Modeling in Deep Reinforcement Learning. In *ICML*. arXiv:1609.05559
- [16] Jonathan Ho and Stefano Ermon. 2016. Generative Adversarial Imitation Learning. In *NIPS*. arXiv:1606.03476
- [17] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. 2015. Human-level concept learning through probabilistic program induction. *Science* 350, 6266 (2015), 1332–1338.
- [18] Hoang M. Le, Yisong Yue, Peter Carr, and Patrick Lucey. 2017. Coordinated Multi-Agent Imitation Learning. In *ICML*. arXiv:1703.03121
- [19] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. unpublished. *Meta-SGD: Learning to Learn Quickly for Few-Shot Learning*. Technical Report. arXiv:1707.09835v2
- [20] Chris Metzzen and James Phinney. 1998. StarCraft: Remastered. <https://starcraft.com/en-us/>
- [21] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. A Simple Neural Attentive Meta-Learner. In *ICLR*. arXiv:1707.03141
- [22] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Tim Harley, Timothy P Lillicrap, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous Methods for Deep Reinforcement Learning. *JMLR* 48 (2016). arXiv:1602.01783v2
- [23] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [24] Arthur O'Dwyer. 2018. Github - quuxplusone/hanabi: Framework for writing bots that play Hanabi. <https://github.com/Quuxplusone/Hanabi>
- [25] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. 2020. Rapid Learning or Feature Reuse? Towards Understanding the Effectiveness of MAML. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rkgMkCetPB>
- [26] Nathan D Ratliff, David Silver, and J Andrew Bagnell. 2009. Learning to search: Functional gradient techniques for imitation learning. *Auton. Robots* 27 (2009), 25–53. <https://doi.org/10.1007/s10514-009-9121-3>
- [27] Stéphane Ross and J Andrew Bagnell. 2010. Efficient Reductions for Imitation Learning. *JMLR* 9 (2010), 661–668.
- [28] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. 2019. Meta-Learning With Latent Embedding Optimization. In *ICLR*. arXiv:1807.05960v3
- [29] Aron Sarmasi, Timothy Zhang, Chu-Hung Cheng, Huyen Pham, Xuanchen Zhou, Duong Nguyen, Soumil Shekdar, and Joshua McCoy. 2021. HOAD: The Hanabi Open Agent Dataset. In *AAMAS*. Montreal, Canada.
- [30] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529 (2016). <https://doi.org/10.1038/nature16961>
- [31] Jake Snell, Kevin Swersky, and Twitter Richard Zemel. 2017. Prototypical Networks for Few-shot Learning. In *NIPS*.
- [32] Flood Sung, Yang Yongxin, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. 2018. Learning to Compare: Relation Network for Few-Shot Learning. In *CVPR*. 1199–1208.
- [33] Oriol Vinyals, Google Deepmind, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching Networks for One Shot Learning. In *NIPS*.
- [34] Joseph Walton-Rivers, Piers R. Williams, Richard Bartle, Diego Perez-Liebana, and Simon M. Lucas. 2017. Evaluating and Modelling Hanabi-Playing Agents. In *CEC*. 1382 – 1389.
- [35] Ziyu Wang, Josh Merel, Scott Reed, Greg Wayne, Nando de Freitas, and Nicolas Heess. 2017. Robust Imitation of Diverse Behaviors. In *Conference on Neural Information Processing Systems*. arXiv:1707.02747
- [36] Tianhe Yu, Chelsea Finn, Annie Xie, Sudeep Dasari, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. 2018. unpublished. *One-Shot Imitation from Observing Humans via Domain-Adaptive Meta-Learning*. Technical Report. arXiv:1802.01557v1
- [37] James Zamiell. [n.d.]. GitHub - Zamiell/hanabi-conventions: A list of Hanabi strategies. <https://github.com/Zamiell/hanabi-conventions>