

Search in Cooperative Partially-Observable Games

Adam Lerer

Facebook AI Research
alerer@fb.com

Hengyuan Hu

Facebook AI Research
hengyuan@fb.com

Jakob Foerster

Facebook AI Research
jnf@fb.com

Noam Brown

Facebook AI Research
noambrown@fb.com

Abstract

Recent superhuman results in games have largely been achieved in a variety of zero-sum settings, such as Go and Poker, in which agents need to compete against others. However, just like humans, real-world AI systems have to coordinate and communicate with other agents in cooperative partially observable environments as well. These settings commonly require participants to both interpret the actions of others and to act in a way that is informative when being interpreted. Those abilities are typically summarized as *theory of mind* and are seen as crucial for social interactions. In this paper we propose two different search techniques that can be applied to improve an arbitrary agreed-upon policy in a cooperative partially observable game. The first one, *single-agent search*, effectively converts the problem into a single agent setting by making all but one of the agents play according to the agreed-upon policy. In contrast, in *multi-agent search* all agents carry out the same common-knowledge search procedure whenever doing so is computationally feasible, and fall back to playing according to the agreed-upon policy otherwise. We prove that these search procedures are theoretically guaranteed to at least maintain the original performance of the agreed-upon policy (up to a bounded approximation error). In the benchmark challenge problem of Hanabi, our search algorithm greatly improves the performance of every agent we tested and when applied to a policy trained using RL achieves a new state-of-the-art score of 24.61 / 25 in the game, compared to a previous-best of 24.07 / 25.

Introduction

Real-world situations such as driving require humans to coordinate with others in a partially-observable environment with limited communication. In such environments, humans have a mental model of how other agents will behave in different situations (theory of mind). This model allows them to change their beliefs about the world based on why they think an agent acted as they did, as well as predict how their own actions will affect others' future behavior. Together, these capabilities allow humans to search for a good action to take while accounting for the behavior of others.

Despite the importance of these cooperative settings to real-world applications, most recent progress on AI in large-

scale games has been restricted to zero-sum settings where agents compete against each other, typically rendering communication useless. Search has been a key component in reaching professional-level performance in zero-sum games, including backgammon (Tesauro 1994), chess (Campbell, Hoane Jr, and Hsu 2002), Go (Silver et al. 2016; 2017; 2018), and poker (Moravčík et al. 2017; Brown and Sandholm 2017; 2019).

Inspired by the success of search techniques in these zero-sum settings, in this paper we propose methods for agents to conduct search given an agreed-upon 'convention' policy (which we call the *blueprint policy*) in cooperative partially observable games. In the first method, a single agent performs search assuming all other agents play according to the blueprint policy. This allows the search agent to treat the known policy of other agents as part of the environment and maintain beliefs about the hidden information based on others' actions.

In the second method, multiple agents can perform search simultaneously but must simulate the search procedure of other agents in order to understand why they took the actions they did. We propose a modification to the multi-agent search procedure - *retrospective belief updates* - that allows agents to fall back to the blueprint policy when it is too expensive to compute their beliefs, which can drastically reduce the amount of computation while allowing for multi-agent search in most situations.

Going from just the blueprint policy, to single-agent search, to multi-agent search each empirically improves performance at the expense of increased computation. Additionally, we prove that both search procedures cannot result in a lower expected value than the blueprint policy except for an error term that decays in the number of Monte Carlo (MC) rollouts.

We test these techniques in Hanabi, which has been proposed as a new benchmark challenge problem for AI research (Bard et al. 2019). Hanabi is a popular fully cooperative, partially observable card game with limited communication. Most prior agents for Hanabi have been developed using handcrafted algorithms or deep reinforcement learning (RL). However, the Hanabi challenge paper itself remarks that "humans approach Hanabi differently than cur-

rent learning algorithms” because while RL algorithms perform exploration to find a good joint policy (convention), humans instead typically start with a convention and then individually search for the best action assuming that their partners will play the convention (Bard et al. 2019).

Applying our search methods to an RL blueprint (that we train in self-play) establishes a new state-of-the-art score of 24.61 / 25 on 2-player Hanabi, compared to a previous-best of 24.07 / 25. Our search methods also achieve state-of-the-art scores in 3, 4, and 5-player Hanabi (Table 1). To our knowledge, this is the first application of theoretically sound search in a large partially observable cooperative game.

We provide code for single- and multi-agent search in Hanabi at https://github.com/anonymous/Hanabi_Search

Related Work

Search has been necessary to achieve superhuman performance in almost every benchmark game. For fully observable games, examples include two-ply search in backgammon (Tesauro 1994), alpha-beta pruning in chess (Campbell, Hoane Jr, and Hsu 2002), and Monte Carlo tree search in Go (Silver et al. 2016; 2017; 2018). The most prominent partially observable benchmark game is poker, where search based on beliefs was key to achieving professional-level performance (Moravčík et al. 2017; Brown and Sandholm 2017; 2019). Our search technique most closely resembles the one used in the superhuman multi-player poker bot *Pluribus*, which conducts search given each agents’ presumed beliefs and conducts MC rollouts beyond the depth limit of the search space assuming all agents play one of a small number of blueprint policies.

Cooperative multi-agent settings have been studied extensively under the DEC-POMDP formalism (Oliehoek, Spaan, and Vlassis 2008). Finding optimal policies in DEC-POMDPs is known to be NEXP-hard (Bernstein et al. 2002), so much prior work studies environments with structure such as factorized local interactions (Oliehoek et al. 2008), hierarchical policies (Amato et al. 2015), or settings with explicit costs of communication (Goldman and Zilberstein 2003).

There has been a good deal of prior work developing agents in Hanabi. At a high level this can be split into two categories: The first consists of a number of hand-crafted bots that incorporate human conventions, heuristics and information theory in order to achieve high scores. Notable examples of human convention bots include Smart-Bot (O’Dwyer 2019) and Fireflower (Wu 2018a). So called ‘hat-coding’ strategies (Wu 2018b), which are particularly successful for $N > 2$ players, instead use information theory by communicating instructions to all players via hints using modulo-coding, as is commonly employed to solve ‘hat-puzzles’. The second category contains more recent work that has focused on tackling Hanabi as a learning problem. In particular, the Hanabi Learning Environment (Bard et al. 2019) contains a set of baselines multi-agent RL methods. In this domain, the Bayesian Action Decoder (BAD) learning method uses a public belief over private features and explores in the space of deterministic partial policies using deep RL (Foerster et al. 2019), which can be regarded as a

scalable instantiation of the general ideas presented in (Nayyar, Mahajan, and Teneketzis 2013). This established the most recent state of the art in Hanabi.

Lastly, there is recent work on ad-hoc team play in Hanabi, in which agents get evaluated against a pool of different team mates (Canaan et al. 2019; Walton-Rivers et al. 2017).

Background

We consider a Dec-POMDP with N agents. The Markov state of the environment is $s \in \mathcal{S}$ and we use i to denote the agent index. At each time step, t , each agent obtains an observation $o^i = Z(s, i)$, where Z is a deterministic observation function, and takes an action $a^i \in \mathcal{A}$, upon which the environment carries out a state transition based on the transition function, $s_{t+1} \sim \mathcal{P}(s_{t+1}|s_t, \mathbf{a})$, and agents receive a team-reward of $r_t = \mathcal{R}(s_t, \mathbf{a})$, where \mathbf{a} indicates the joint action of all agents. We use $\tau_t = \{s_0, \mathbf{a}_0, r_0, \dots, s_t\}$ to denote the game history (or ‘trajectory’) at time t . Each agent’s policy, π^i , conditions on the agent’s *action-observation history* (AOH), $\tau_t^i = \{o_0^i, a_0^i, r_0, \dots, o_t^i\}$, and the goal of the agents is to maximise the total expected return, $J_\pi = \mathbb{E}_{\tau \sim P(\tau|\pi)} R_0(\tau)$, where $R_0(\tau)$ is the forward looking return of the trajectory, $R_t(\tau) = \sum_{t' \geq t} \gamma^{t'-t} r_{t'}$, and γ is an (optional) discount factor. In this paper we consider both deterministic policies, $a^i = \pi(\tau_t^i)$, typically based on heuristics or search, and stochastic policies, $a^i \sim \pi^\theta(a^i|\tau_t^i)$, where θ are the parameters of a function approximator, e.g., a deep neural network.

In order to represent variable sequence lengths trajectories, τ_t^i , Deep RL in partially observable settings typically uses recurrent networks (RNNs) (Hausknecht and Stone 2015). These RNNs can learn implicit representations of the sufficient statistics over the Markov state given τ_t^i . In contrast, in our work we will use explicit *beliefs* to represent the probability distribution over possible trajectories¹. The private belief of agent i that they are in trajectory τ_t at time step t is $B^i(\tau_t) = P(\tau_t|\tau_t^i)$. We also define $B(\tau_t) = P(\tau_t|\tau_t^G)$ and $B(\tau_t^i) = P(\tau_t^i|\tau_t^G) = \sum_{\tau_t \in \tau_t^i} B(\tau_t)$, which is the probability that agent i is in AOH τ_t^i conditional only on the *common knowledge* (CK) of the entire group of agents τ_t^G . Here CK are things that all agents know that all agents know ad infinitum. Please see (Osborne and Rubinstein 1994) for a formal definition of CK and (Foerster et al. 2018; 2019) for examples of how CK can arise and be used in multi-agent learning. Practically it can be computationally challenging to exactly compute common-knowledge beliefs due a large number of possible states and trajectories. However, in many settings, e.g., poker (Brown and Sandholm 2019), the CK-belief can be factorized across public features and a set of private features associated with the different agents. In these settings the CK trajectory is simply

¹Maintaining exact beliefs in large POMDPs is typically considered intractable. However we find that even in an environment with large state spaces such as card games, the number of states that have non-zero probability conditional on a player’s observations is often much smaller; in the case of Hanabi, this set of possible states can be stored and updated explicitly.

the history of public features.

In our methods we will commonly need to carry out computations over each of the possible trajectories that have non-zero probability given an agent’s AOH or all agents’ common knowledge. We refer to these as the *trajectory range* $\beta_t^i = \{\tau_t | B^i(\tau_t) > 0\}$ and $\beta_t = \{\tau_t | B(\tau_t) > 0\}$. We refer to the vector of probabilities in the trajectory range on timestep t by \mathbf{B}_t^i and \mathbf{B}_t . We also commonly need to carry out computations over each of an agent’s AOHs that have non-zero probability given the common knowledge of the entire group of agents, which we refer to as the *AOH range* $\chi_t^i = \{\tau_t^i | B(\tau_t^i) > 0\}$. We refer to the entire vector of probabilities in the AOH range of agent i on timestep t by \mathbf{C}_t^i .

We further define the typical conditional expectations (‘value functions’),

$$V_\pi(\tau_t) = \mathbb{E}_{\tau'_T \sim P(\tau'_T | \pi, \tau_t)} R_t(\tau'_T) \quad (1)$$

$$Q_\pi(\tau_t, a^i) = \mathbb{E}_{\tau'_T \sim P(\tau'_T | \pi, \tau_t, a^i)} R_t(\tau'_T) \quad (2)$$

Given the range defined above we also introduce expectations conditioned on the AOHs, τ_t^i :

$$V_\pi(\tau_t^i) = \sum_{\tau_t \in \beta_t^i} B^i(\tau_t) V_\pi(\tau_t) \quad (3)$$

$$Q_\pi(\tau_t^i, a^i) = \sum_{\tau_t \in \beta_t^i} B^i(\tau_t) Q_\pi(\tau_t, a^i) \quad (4)$$

Even though the optimal policies in the fully cooperative setting are deterministic, we consider stochastic policies for the purpose of reinforcement learning.

We further assume that a deterministic *blueprint* policy π_b , defining what each player should do for all possible trajectories, is common knowledge amongst all players. Player i ’s portion of the blueprint policy is denoted π_b^i and the portion of all players other than i as π_b^{-i} . In some settings, when actually playing, players may choose to not play according to the blueprint and instead choose a different action determined via online search. In that case π_b differs from π , which denotes the policy that is actually played.

In our setting all of the past actions taken by all agents and the observation functions of all players are common knowledge to all agents. As such, if an agent is known to be playing according to a given policy, each action taken by this agent introduces a belief update across all other agents and the public belief. Suppose agent i has a current belief B_{t-1}^i and next observes (a_t^j, o_t^i) , where we have broken up the observation to separate out the observed partner action. Then,

$$B^i(\tau_t) = P(\tau_t | \tau_t^i) = P(\tau_t | \tau_{t-1}^i, o_t^i, a_t^j) \quad (5)$$

$$= \frac{B^i(\tau_{t-1}) \pi^j(a_t^j | \tau_{t-1}) P(o_t^i | \tau_{t-1}, a_t^j)}{\sum_{\tau'_{t-1}} B^i(\tau'_{t-1}) \pi^j(a_t^j | \tau'_{t-1}) P(o_t^i | \tau'_{t-1}, a_t^j)} \quad (6)$$

In other words, the belief update given (a_t^j, o_t^i) consists of two updates: one based on the partner’s known policy π^j , and the other based on the dynamics of the environment. The common knowledge belief B is updated in the same way, using the common-knowledge observation τ^G rather than τ^i .

Method

In this section we describe our online search algorithm for cooperative partially-observable games.

Single-Agent Search

We first consider the case in which only one agent conducts online search. We denote the searching agent as agent i . Every other agent simply plays according to the blueprint policy (and we assume agent i knows all other agents play according to the blueprint).

Since agent i is the only agent determining her policy online while all other agents play a fixed common-knowledge policy, this is effectively a single-agent POMDP for agent i . Specifically, agent i maintains a belief distribution \mathbf{B}_t^i over trajectories she might be in based on her AOH τ_t^i and the known blueprint policy of the other agents. Each time she receives an observation or another agent acts, agent i updates her belief distribution according to (6) (see Figure 1, left). Each time agent i must act, she estimates via Monte Carlo rollouts the expected value $Q_{\pi_b}(\tau_t^i, a^i)$ of each action assuming *all* agents (including agent i) play according to the joint blueprint policy π_b for the remainder of the game following the action (Figure 1, right). A precise description of the single-agent search algorithm is provided in the appendix.

As described, agent i calculates the expected value of only the next action. This is referred to as 1-ply search. One could achieve even better performance by searching further ahead, or by having the agent choose between multiple blueprint policies for the remainder of the game. However, the computational cost of the search would also increase, especially in a game like Hanabi that has a large branching factor due to chance. In this paper all the experiments use 1-ply search.

Since this search procedure uses exact knowledge of all other agents’ policies, it cannot be conducted correctly by multiple agents independently. That is, if agent j conducts search on a turn after agent i conducted search on a previous turn, then agent j ’s beliefs are incorrect because they assume agent i played π_b^i while agent i actually played the modified policy π^i . If more than one agent independently performs search assuming that others follow the blueprint, policy improvement cannot be guaranteed, and empirical performance is poor (Table 4, Appendix).

Multi-Agent Search

In order for an agent to conduct search effectively, her belief distribution must be accurate. In the case of single-agent search, this was achieved by all agents agreeing beforehand on a blueprint policy, and then also agreeing that only one agent would ever conduct search and deviate from the blueprint. In this section we instead assume that all agents agree beforehand on both a blueprint policy and on what search procedure will be used. When agent i acts and conducts search, the other agents exactly replicate the search procedure conducted by agent i (including the random seed) and compute agent i ’s resulting policy accordingly. In this way, the policy played so far is always common knowledge.

Since the other agents do not know agent i 's private observations, we have all agents (including agent i) conduct search and compute agent i 's policy for *every possible* AOH that agent i might be in based on the common-knowledge observations. Specifically, all agents conduct search for every AOH $\tau_t^i \in \chi_t^i$. When conducting search for a particular τ_t^i as part of this loop, the agents also compute what B_t^i would be assuming agent i 's AOH is τ_t^i and compute $Q_{\pi_b}(\tau_t^i, a^i)$ for every action a^i based on this B_t^i . We refer to this loop of search over all $\tau_t^i \in \chi_t^i$ as *range-search*.

Agent i must also compute her policy via search for every $\tau_t^i \in \chi_t^i$ (that is, conduct range-search) even though she knows τ_t^i , because the search procedure of other agents on future timesteps may be based on agent i 's policy for $\tau_t^i \neq \tau_t^i$ where $\tau_t^i \in \chi_t^i$ and it is necessary for all agents to be consistent on what that policy is to ensure that future search procedures are replicated identically by all agents.

In a game like two-player Hanabi, $|\chi_t^i|$ could be nearly 10 million, which means the range-search operation of multi-agent search could be 10 million times more expensive than single-agent search in some situations and therefore infeasible. Fortunately, the actual number of positive-probability AOHs will usually not be this large. We therefore have all agents agree beforehand on a budget for range-search, which we refer to as a *max range* (abbreviated *MR*). If $|\chi_t^i| > MR$ on timestep t where agent i is the acting agent, then agent i does not conduct search and instead simply plays according to the blueprint policy π_b . Since χ_t^i and the max range are common knowledge, it is also common knowledge when an agent does not search on a timestep and instead plays according to the blueprint.

Using a max range makes multi-agent search feasible on certain timesteps, but the fraction of timesteps in which search can be conducted may be less than single-agent search (which in a balanced two-player game is 50%). The next section describes a way to use a max range while guaranteeing that there's always at least one agent who can perform search.

Retrospective Policy Updates

As discussed in the previous section, χ_t^i on some timesteps may be too large to conduct search on. Using a max range mitigates this problem, but may result in search only rarely being conducted. Fortunately, in many domains more common knowledge information is revealed as the game progresses, which reduces the number of positive-probability AOHs on previous timesteps.

For example, at the start of a two-player game of Hanabi in which agent i acts first (and then agent j), $|\chi_0^i|$ might be nearly 10 million. If agent i were to use search to choose an action at this point, it would be too expensive for the agents to run range-search given the magnitude of $|\chi_0^i|$, so the other agents would not be able to conduct search on future timesteps.

However, suppose the action chosen from agent i 's search results in agent i giving a hint to agent j . Given this new common-knowledge information, it might now be known that only 100,000 of the 10 million seemingly possible AOHs at the first timestep were actually possible. It may

now be feasible for the agents to run range-search on this reduced set of 100,000 possible AOHs. In this way, agent i is able to conduct search on a timestep t where the max range is exceeded, and the agents can execute range-search at some later timestep t' once further observations have reduced the size of χ_t^i below the max range.

We now introduce additional notation to generalize this idea. Agent i 's belief at timestep t' that the trajectory at some earlier timestep t was (or is) τ_t is $B_{t'}^i(\tau_t) = P(\tau_t | \tau_{t'}^i)$. The public belief at timestep t' , which conditions only on the common knowledge of the entire group of agents at timestep t' , that the trajectory at timestep t was (or is) τ_t is $B_{t'}(\tau_t) = P(\tau_t | \tau_{t'}^G)$ and $B_{t'}(\tau_t^i) = P(\tau_t^i | \tau_{t'}^G) = \sum_{\tau_t \in \tau_t^i} B_{t'}(\tau_t)$. The trajectory range at timestep t' of the trajectories at timestep t is $\beta_{t,t'} = \{\tau_t | B_{t'}(\tau_t) > 0\}$ and $\beta_{t,t'}^i = \{\tau_t | B_{t'}^i(\tau_t) > 0\}$. The AOH range at timestep t' of the agent i AOHs at timestep t is $\chi_{t,t'}^i = \{\tau_t^i | B_{t'}^i(\tau_t^i) > 0\}$.

Again, the key idea behind retrospective updates is that an agent can delay running range-search on a timestep until that range shrinks based on subsequent observations. When it is an agent's turn to act, she conducts search if and only if she has run range-search for each previous timestep t where one of the *other* agents played search (because this means she knows her belief distribution). Otherwise, she plays according to the blueprint policy.

Specifically, all agents track the oldest timestep on which search was conducted but range-search was not conducted. This is denoted t^* . Assume the agent acting at t^* was agent i . If at any point $|\chi_{t^*,t}^i| \leq MR$ then all agents conduct range-search for timestep t^* and t^* is incremented up to the next timestep on which search was conducted but range-search was not conducted (but obviously not incremented past the current timestep t). If $|\chi_{t^*,t}^i| \leq MR$ for this new t^* then all agents again conduct range-search and the process repeats. Since $\chi_{t^*,t}^i$ depends only on common knowledge, all agents conduct range-search at the same time and therefore t^* is always consistent across all agents. When it is an agent's turn to act on timestep t , she conducts search if and only if she was the agent to act on timestep t^* or if $t = t^*$. An important upshot of this method is that it's always possible for at least one agent to use search when acting.

If *MR* is set to zero then multi-agent search with retrospective updates is identical to single-agent search, because the first agent to act in the game will conduct search and she will continue to be the only agent able to conduct search on future timesteps. As *MR* is increased, agents are able to conduct search on an increasing fraction of timesteps. In two-player Hanabi, setting *MR* = 10,000 makes it possible to conduct search on 86.0% of timesteps even though in the worst case $|\chi_t^i| \approx 10,000,000$.

Soundness and Convergence Bound for Search

We now prove a theorem that applies to all three search variants described in this section. Loosely, it states that applying search on top of a blueprint policy cannot reduce the expected reward relative to the blueprint, except for an error term that decays as $O(1/\sqrt{N})$, where N is the number of Monte Carlo rollouts.

Theorem 1. Consider a Dec-POMDP with N agents, actions \mathcal{A} , reward bounded by $r_{\min} \leq R(\tau) < r_{\max}$ where $r_{\max} - r_{\min} \leq \Delta$, game length bounded by T , and a set of blueprint policies $\pi_b \equiv \{\pi_b^0, \dots, \pi_b^N\}$. If a Monte Carlo search policy π_s is applied using N rollouts per step, then

$$V_{\pi_s} - V_{\pi_b} \geq -2T\Delta|\mathcal{A}|N^{-1/2} \quad (7)$$

The proof is provided in the Appendix.

Experimental Setup

We evaluate our methods in the partially observable, fully cooperative game Hanabi, which at a high level resembles a cooperative extension of solitaire. Hanabi has recently been proposed as a new frontier for AI research (Bard et al. 2019) with a unique focus on theory of mind and communication.

The main goal of the team of agents is to complete 5 stacks of cards, one for each color, in a legal sequence, starting with a 1 and finishing with a 5. The defining twist in Hanabi is that while players can observe the cards held by their teammates, they cannot observe their own cards. As such, players need to exchange information with their teammates in order to decide which cards to play. Hanabi offers two different means for doing so. First, players can take costly hint actions that reveal part of the state to their teammates. Second, since all actions are observed by all players, each action (such as playing a card or discarding a card) can itself be used to convey information, in particular if players agree on a set of conventions before the game. For further details on the state and action space in Hanabi please see (Bard et al. 2019).

In this work we are focused on the *self-play* part of the challenge, in which the goal is to find a set of policies that achieve a high score when playing together as a team.

For the blueprint strategy used in the experiments, we experimented with open-sourced handcrafted bots WT-FWThat (Wu 2018b) and SmartBot (O’Dwyer 2019). We also created two of our own blueprint strategies. One was generated from scratch using deep reinforcement learning, which we call RLBot. The other, which we call CloneBot, was generated by conducting imitation learning using deep neural networks on the policy produced by single-agent search on top of SmartBot. The details for the generation of both bots are given in the appendix.

All experiments except the imitation learning of CloneBot and the reinforcement learning of RLBot were conducted on CPU using machines with Intel® Xeon® E5-2698 CPUs containing 40 cores each. A game of Hanabi requires about 2 core-hours for single-agent search and 90 core-hours for retrospective multi-agent search using the SmartBot blueprint policy with a max range of 10,000. We parallelize the search procedure over multiple cores on a single machine.

Search Implementation in Hanabi

The public and private observations in Hanabi are factorizable into public and private features.² Specifically, each

²The class of games of this form - which includes most card games - has recently been formalized as ‘Factorized Observation Games’ in (Kovařík et al. 2019).

player’s hidden information consists of the cards held by other agents, so the beliefs for each player can be represented as a distribution over possible hands that player may be holding. The initial private beliefs can be constructed based on the card counts. In addition to updates based on the actions of other agents, updates based on observations amount to (a) adjusting the probabilities of each hand based on the modified card count as cards are revealed, and (b) setting the probability of hands inconsistent with hints to 0.

The public beliefs for 2-player Hanabi can be factored into independent probability distributions over each player’s hand (conditional on the common knowledge observations). These public beliefs are identical to the private beliefs except that the card counts are not adjusted for the cards in the partner’s hand. Given one player’s hand, the private beliefs over the other player’s hand (which is of course no longer independent of the other player’s hand) can be computed from the public beliefs by adjusting the card counts for the privately-observed cards.³

Estimating Action Expected Values via UCB

In order to reduce the number of MC rollouts that must be performed during search, we use a UCB-like procedure that skips MC rollouts for actions that are presumed not to be the highest-value action with high confidence. After a minimum of 100 rollouts per action is performed, the reward sample mean and its standard deviation is computed for each action. If the expected value for an action is not within 2 standard deviations of the expected value of the best action, its future MC rollouts are skipped.

Furthermore, we use a configurable threshold for deviating from the blueprint action. If the expected value of the action chosen by search does not exceed the value of the blueprint action by more than this threshold, the agent plays the blueprint action. We use a threshold of 0.05 in our experiments.

The combination of UCB and the blueprint deviation threshold reduces the number of rollouts required per timestep by 10×, as shown in Figure 3 in the appendix.

Bootstrapping Search-Based Policies via Imitation Learning

Search can be thought of as a policy improvement operator, i.e. an algorithm that takes in a joint policy and outputs samples from an improved joint policy. These samples can be used as training data to learn a new policy via imitation learning. This improved policy approximates the effect of search on the blueprint policy while being cheaper to execute, and search can be run on this learned policy, in effect bootstrapping the search procedure. In principle, repeated application of search and learning could allow the effect of single-agent search to be applied on multiple agents, and could allow the benefits of search to extend beyond the depth limit of the search procedure. However, there is no guarantee that this process would eventually converge to an optimal policy, even in the case of perfect function approximation.

³This conversion from public to private beliefs is what we call ConditionOnAOH() in the algorithm listing in the Appendix.

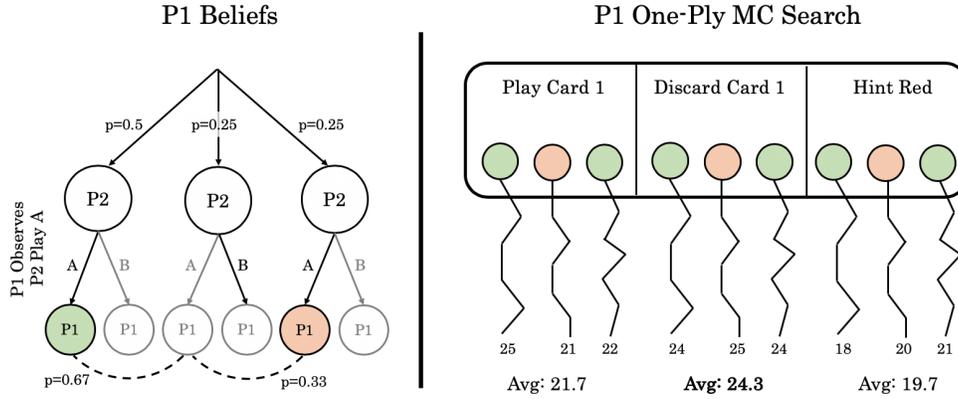


Figure 1: An illustration of the proposed Monte Carlo search procedure for an agent P1. **Left:** The MDP starts in one of three states and P2 plays A. P1 observes that P2 played A, ruling out the middle state in which P2 would have played B. This leaves P1 with a probability distribution (beliefs) over two possible states. **Right:** For each legal action (‘Play Card 1’, ...) P1 performs rollouts from states drawn from the belief distribution, and picks the action with the highest average score: Discard Card 2.

We refer to the policy learned in this manner as CloneBot. We provide details of the training procedure in the Appendix and evaluate its performance in Table 1.

Results

Table 1 shows that adding search leads to a large improvement in performance for all blueprint policies tested in two-player Hanabi (the most challenging variant of Hanabi for computers) and achieves a new state-of-the-art score of 24.61 / 25 compared to the previous-best of 24.07 / 25.⁴ Much of this improvement comes from adding single-agent search, though adding multi-player search leads to an additional substantial improvement.

Unfortunately there are no reliable statistics on top human performance in Hanabi. Discussions with highly experienced human players has suggested that top players might achieve perfect scores in 2-player Hanabi somewhere in the range of 60% to 70% of the time when optimizing for perfect scores. Our strongest agent optimizes for expected value rather than perfect scores and still achieves perfect scores 75.5% of the time in 2-player Hanabi.

Table 2 shows the benefits of single-agent search also extend to the 3, 4, and 5-player variants of Hanabi as well. For these variants, the *WTFWThat* agent (Wu 2018b) was state-of-the-art, achieving close to perfect scores for 4 and 5 players. Applying single-agent search to these policies achieves new state-of-the-art scores of 24.83, 24.96, and 24.94 / 25 for the 3, 4, and 5-player variants. Applying multi-agent search in Hanabi with more than 2 players is currently intractable because the range of AOHs that must be tracked for P players and H possible hands is roughly H^{P-1} .

Table 3 examines the performance and cost in number of rollouts for single-agent search and multi-agent search with different values of the max range (MR) parameter, using SmartBot as the blueprint agent. When MR is set to zero, we are performing single-agent search, so search is conducted

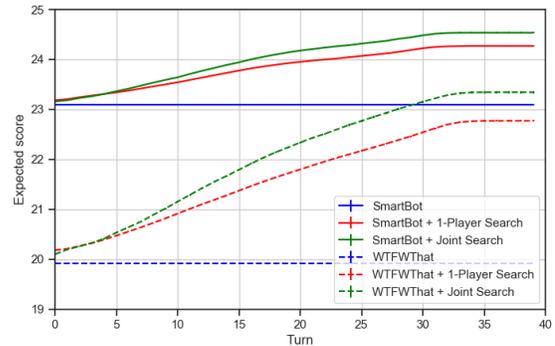


Figure 2: Average score predicted by MC search rollouts at different points in the game, for two different blueprints. Expected score by turn is averaged over replicates for each condition (error bars are included but not too small to be visible). If a game ends before 40 turns, the final score is propagated for all subsequent turns.

on 50% of timesteps, which requires about 10^5 rollouts per game. As the max range is increased, search is conducted more often. At a max range of 10,000, search is conducted on 86% of timesteps even though the maximum possible range for a timestep in two-player Hanabi is nearly 10 million. However, this still requires about $1,000\times$ as many rollouts as single-agent search.

Figure 2 plots the average MC search prediction of the expected payoff of the best move ($Q_{\pi_b}(\tau^i, a^*)$) at different points in the game, for two blueprint policies. As predicted by the theory, the expected score starts at the blueprint expected score and increases monotonically as search is applied at each move of the game.

Conclusions

In this paper we described approaches to search in partially observable cooperative games that improves upon an arbi-

⁴It is impossible to achieve a perfect score for some shuffles of the deck. However, the highest possible average score is unknown.

Blueprint Strategy	No Search	Single-Agent Search	Multi-Agent Search
RL (ACHA) (Bard et al. 2019)	22.73 ± 0.12 15.1%	- -	- -
BAD (Foerster et al. 2019)	23.92 ± 0.01 58.6%	- -	- -
SmartBot (O’Dwyer 2019)	22.99 ± 0.001 29.6%	24.21 ± 0.002 59.1%	24.48 ± 0.01 67.1%
CloneBot (ours)	23.32 ± 0.001 40.6%	24.37 ± 0.02 64.6%	- -
DQN (ours)	23.45 ± 0.01 46.8%	24.30 ± 0.02 63.5%	24.49 ± 0.02 67.7%
SAD (Anonymous 2020)	24.07 ± 0.01 56.4%	24.53 ± 0.01 71.1%	24.61 ± 0.01 75.5%

Table 1: Results for various policies in 2-player Hanabi coupled with single-agent and retrospective multi-agent search. The top row is the average score (25 is perfect), the bottom row is the perfect-score percentage. Single-agent and multi-agent search monotonically improve the performance of different blueprint policies. For multi-agent search, a max range of 400 is used for the RLBot blueprint and 10,000 for the other blueprints. Results in bold beat the prior state of the art of 23.92.

# Players	No Search	Single-Agent Search
2	19.45 ± 0.001 0.28%	22.78 ± 0.02 10.0%
3	24.20 ± 0.01 49.1%	24.83 ± 0.006 85.9%
4	24.83 ± 0.01 87.2%	24.96 ± 0.003 96.4%
5	24.89 ± 0.00 91.5%	24.94 ± 0.004 95.5%

Table 2: Hanabi performance for different numbers of players, using the ‘WTFWThat’ hat-counting blueprint agent which achieved state-of-art performance in 3, 4, and 5-player Hanabi. Results shown in bold are state-of-the-art for that number of players.

trary blueprint policy. Our algorithms ensure that any agent conducting search always has an accurate belief distribution over the possible trajectories they may be in, and provide an alternative in case the search procedure is intractable at certain timesteps. We showed that in the benchmark domain of Hanabi, both search techniques lead to large improvements in performance to all the policies we tested on, and achieves a new state of the art score of 24.61 / 25 compared to a previous-best of 24.07 / 25. We also proved that applying our search procedure cannot hurt the expected reward relative to the blueprint policy, except for an error term that shrinks with more Monte Carlo rollouts. This result fits a theme, also shown in other games such as Chess, Go, and Poker, that search leads to dramatically improved performance compared to learning or heuristics alone.

The performance improvements from search come at a computational cost. Our search procedure involves tracking

Max Range	% search	# Rollouts	Average Score
0	50.0%	1.5×10^5	24.21 ± 0.002
80	56.2%	3.0×10^6	24.30 ± 0.02
400	65.1%	1.3×10^7	24.40 ± 0.02
2000	77.2%	5.2×10^7	24.46 ± 0.01
10000	86.0%	1.8×10^8	24.48 ± 0.01

Table 3: Multi-agent search with a SmartBot blueprint for different values of max range. As max range increases, search is performed more often leading to a higher score, but the rollout cost increases dramatically.

the belief probability of each AOH that a player may be in given the public information, which for Hanabi is no more than 10 million probabilities. Search also requires computing a large number of MC rollouts of the policy, especially for multi-agent search where the number of rollouts scales with the size of the belief space. Conducting search in partially observable games, whether cooperative or competitive, with many more AOHs per instance of public information remains an interesting challenge for future work and may be applicable to settings like Bridge and environments with visual inputs like driving.

This work currently assumes perfect knowledge of other agents’ policies. This is a reasonable assumption in settings involving centralized planning but decentralized execution, such as self-driving cars that are created by a single company or robots working together as a team in a factory. In general however, an agent’s model of others may not be perfect. Investigating how search performs in the presence of an imperfect model of the partner, and how to make search more robust to errors in that model, are important directions for future work as well.

References

- Amato, C.; Konidaris, G.; Cruz, G.; Maynor, C. A.; How, J. P.; and Kaelbling, L. P. 2015. Planning for decentralized control of multiple robots under uncertainty. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 1241–1248.
- Anonymous. 2020. Simplified action decoder for deep multi-agent reinforcement learning. In *Submitted to International Conference on Learning Representations*. under review.
- Bard, N.; Foerster, J. N.; Chandar, S.; Burch, N.; Lanctot, M.; Song, H. F.; Parisotto, E.; Dumoulin, V.; Moitra, S.; Hughes, E.; et al. 2019. The Hanabi challenge: A new frontier for AI research. *arXiv preprint arXiv:1902.00506*.
- Bernstein, D. S.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The complexity of decentralized control of markov decision processes. *Mathematics of operations research* 27(4):819–840.
- Brown, N., and Sandholm, T. 2017. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science* eaao1733.
- Brown, N., and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science* eaay2400.
- Campbell, M.; Hoane Jr, A. J.; and Hsu, F.-h. 2002. Deep Blue. *Artificial intelligence* 134(1-2):57–83.
- Canaan, R.; Togelius, J.; Nealen, A.; and Menzel, S. 2019. Diverse agents for ad-hoc cooperation in hanabi. *arXiv preprint arXiv:1907.03840*.
- Foerster, J. N.; de Witt, C. A. S.; Farquhar, G.; Torr, P. H.; Boehmer, W.; and Whiteson, S. 2018. Multi-agent common knowledge reinforcement learning. *arXiv preprint arXiv:1810.11702*.
- Foerster, J.; Song, F.; Hughes, E.; Burch, N.; Dunning, I.; Whiteson, S.; Botvinick, M.; and Bowling, M. 2019. Bayesian action decoder for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, 1942–1951.
- Goldman, C. V., and Zilberstein, S. 2003. Optimizing information exchange in cooperative multi-agent systems. In *Autonomous Agents and Multiagent Systems*, 137–144. ACM.
- Hausknecht, M., and Stone, P. 2015. Deep recurrent q-learning for partially observable mdps. In *2015 AAAI Fall Symposium Series*.
- Horgan, D.; Quan, J.; Budden, D.; Barth-Maron, G.; Hessel, M.; van Hasselt, H.; and Silver, D. 2018. Distributed prioritized experience replay. *CoRR* abs/1803.00933.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kovařík, V.; Schmid, M.; Burch, N.; Bowling, M.; and Lisý, V. 2019. Rethinking formal models of partially observable multiagent decision making. *arXiv preprint arXiv:1906.11110*.
- Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. 2017. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science* 356(6337):508–513.
- Nayyar, A.; Mahajan, A.; and Teneketzis, D. 2013. Decentralized stochastic control with partial history sharing: A common information approach. *IEEE Transactions on Automatic Control* 58(7):1644–1658.
- O’Dwyer, A. 2019. Hanabi. <https://github.com/Quuxplusone/Hanabi>.
- Oliehoek, F. A.; Spaan, M. T.; Whiteson, S.; and Vlassis, N. 2008. Exploiting locality of interaction in factored dec-pomdps. In *Autonomous Agents and Multiagent Systems-Volume 1*, 517–524.
- Oliehoek, F. A.; Spaan, M. T.; and Vlassis, N. 2008. Optimal and approximate q-value functions for decentralized pomdps. *Journal of Artificial Intelligence Research* 32:289–353.
- Osborne, M. J., and Rubinstein, A. 1994. *A course in game theory*. MIT press.
- Rigollet, P. 2015. High-dimensional statistics.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized Experience Replay. *arXiv e-prints* arXiv:1511.05952.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587):484.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *Nature* 550(7676):354.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science* 362(6419):1140–1144.
- Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Mach. Learn.* 3(1):9–44.
- Tesauro, G. 1994. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural computation* 6(2):215–219.
- van Hasselt, H.; Guez, A.; and Silver, D. 2015. Deep reinforcement learning with double q-learning. *CoRR* abs/1509.06461.
- Walton-Rivers, J.; Williams, P. R.; Bartle, R.; Perez-Liebana, D.; and Lucas, S. M. 2017. Evaluating and modelling hanabi-playing agents. In *IEEE Congress on Evolutionary Computation (CEC)*, 1382–1389.
- Wang, Z.; de Freitas, N.; and Lanctot, M. 2015. Dueling network architectures for deep reinforcement learning. *CoRR* abs/1511.06581.
- Wu, D. 2018a. A rewrite of hanabi-bot in scala. <https://github.com/lightvector/fireflower>.
- Wu, J. 2018b. Hanabi simulation in rust. <https://github.com/WuTheFWasThat/hanabi.rs>.