

Diversifying AI: Towards Creative Chess with AlphaZero

Tom Zahavy^a, Vivek Veeriah^a, Shaobo Hou^a, Kevin Waugh^{b,*}, Matthew Lai^a, Edouard Leurent^a, Nenad Tomašev^a, Lisa Schut^{c,*}, Demis Hassabis^a and Satinder Singh^a

^aGoogle DeepMind, London, United Kingdom, ^bSony AI, New York, NY, 10010, ^cUniversity of Oxford, Oxford, United Kingdom, ^{*}Work done at Google DeepMind.

In recent years, Artificial Intelligence (AI) systems have surpassed human intelligence in a variety of computational tasks. However, AI systems, like humans, make mistakes, have blind spots, hallucinate, and struggle to generalize to new situations. This work explores whether AI can benefit from creative decision-making mechanisms when pushed to the limits of its computational rationality. In particular, we investigate whether a team of diverse AI systems can outperform a single AI in challenging tasks by generating more ideas as a group and then selecting the best ones. We study this question in the game of chess, the so-called “drosophila of AI”. We build on AlphaZero (AZ) and extend it to represent a league of agents via a latent-conditioned architecture, which we call AZ_{db} . We train AZ_{db} to generate a wider range of ideas using behavioral diversity techniques and select the most promising ones with sub-additive planning. Our experiments suggest that AZ_{db} plays chess in diverse ways, solves more puzzles as a group and outperforms a more homogeneous team. Notably, AZ_{db} solves twice as many challenging puzzles as AZ, including the challenging Penrose positions. When playing chess from different openings, we notice that players in AZ_{db} specialize in different openings, and that selecting a player for each opening using sub-additive planning results in a 50 Elo improvement over AZ. Our findings suggest that diversity bonuses emerge in teams of AI agents, just as they do in teams of humans and that diversity is a valuable asset in solving computationally hard problems.




Figure 1 | Left: an illustration of diverse AI systems playing chess. Right: AZ_{db} architecture. For each player i , a matchmaker selects an opponent j . A latent variable l^i associated with the player conditions the network outputs p, v, v_d . A score for each position is computed using v, v_d and the diversity intrinsic reward r^i and used with the prior p for MCTS.

Contents

1	Introduction	2
2	Background	5
3	Methods	5
4	Experiments	8
5	Related work	17
6	Summary	19
7	Acknowledgements	20
A	Additional related work on QD	21
B	Response diversity	22
	B.1 Technical details	22
	B.2 Matchmaker comparisons	25
C	Diversity bonuses in AZ	28
	C.1 Evaluation of external chess engines in puzzle	29
	C.2 Visualizing the piece occupancies	29
	C.3 Details on self-play from puzzle positions.	32
	C.4 Additional results on AZ’s understanding of puzzles	33
	C.5 Sub-additive planning in chess matchups	34
D	Puzzle sets	38

1. Introduction

Artificial Intelligence (AI) systems have achieved superhuman performance in a variety of fields, including playing games like chess, Go, poker, and Atari (Schrittwieser et al., 2020), and more recently, in retrieving knowledge, reasoning over it and answering questions in [medical](#), [law](#), and [business](#) domains (Chowdhery et al., 2022; Katz et al., 2023; OpenAI, 2023; Singhal et al., 2023). AI can outperform humans in many ways; it has greater computational power and ability to experience and learn from more data. Nevertheless, AI is limited by the amount of data and compute available to it, as well as the quality of human feedback. Reinforcement learning (RL) can further improve AI systems by enabling them to learn from trial and error, explore new possibilities, and discover new knowledge.

AlphaZero (AZ; [Silver et al., 2018a](#)), an RL algorithm, learned to play Go, chess, and shogi from scratch, without any prior knowledge of the games. It not only outperformed human experts, but also played moves that were considered to be creative ([Metz, 2016](#)). Millions of people watched AZ play against [StockFish](#) (SF), ([Agadmator](#), [GothamChess](#)), and books have been written about it ([Sadler and Regan, 2019](#)). AZ has influenced the style of professional chess players ([González-Díaz and Palacios-Huerta, 2022](#)) and is associated with increased novelty in the decision-making of Go players ([Shin et al., 2023](#)). Indeed, AZ has had a profound impact on the world of chess, inspiring most contemporary chess engines, such as [Leela Chess Zero](#) (Lc0) and SF, to incorporate neural networks into their evaluation functions. More recently, AZ has been extended beyond games to discover better algorithms to compress videos ([Mandhane et al., 2022](#)), for matrix multiplication ([Fawzi et al., 2022](#)), and similar techniques combining search and a neural network have been applied to other domains ([Kemmerling et al., 2023](#)).

Even though chess engines are very powerful, we have most likely not reached optimal play yet. For instance, in the related game of Go, superhuman AI has been shown to be vulnerable to adversarial attacks ([Timbers et al., 2022](#); [Wang et al., 2022](#)). Other evidence includes the results of computer chess competitions like the [Top Chess Engine Championship](#) (TCEC), where the strength of the top engines continues to increase with each season. It is well known that the strength of modern chess engines depends on the amount of compute they have for search; for example, AZ’s policy network, i.e., AZ without search, is rated at least 1000 Elo points lower than AZ. More evidence can be found in [anti-computer tactics](#), which exploit blind spots in how computers reason about chess. We can find blind spots in chess puzzles, which computers struggle to solve for different reasons: they can be counterintuitive, may require long-term planning, were designed to trick specific chess engines, or involve querying a chess engine in positions that are different from what it was trained on (out-of-distribution (OOD)). Finding positions that computers don’t understand attracted the interest of many chess players ([Copeland, 2021a](#); [Doggers, 2017](#); [Penrose, 2020](#)). While these positions become rarer over time, computers still can’t solve all of them. In addition, despite being considered weaker players overall, humans can solve most of these puzzles, demonstrating a gap between human and machine thinking. Some scientists even ask if understanding these positions can “*help to understand human consciousness*” ([Doggers, 2017](#)).

This work aims to bridge the gap between human and artificial intelligence by investigating whether creative problem-solving mechanisms identified in human intelligence can be used to make artificial intelligence more creative. We acknowledge that there are different definitions of creativity, which vary across domains and disciplines, but broadly speaking there is “a general agreement that creativity involves the production of novel, useful products” ([Mumford, 2003](#)). In this work, we focus on creative problem-solving rather than creativity in other domains such as art or AI-generated art. Even within this narrower context, there is no agreed-upon definition of creativity. Chess players, for example, have difficulty describing their creative thinking process in words: “In my games I have sometimes found a combination intuitively, simply feeling that it must be there. Yet, I was not able to translate my thoughts into normal human language” (Michail Tal). Similarly, creativity in mathematical invention is considered spontaneous and unconscious ([Hadamard, 2020](#)). For these reasons, we have chosen to limit our discussion to a specific definition for creative-problem solving by Osborn: “*searching for an original and previously unknown solution to a problem*” ([Osborn, 2012](#)).

In particular, we ask if a team of diverse chess agents can become more creative than a single super human AI. We propose to measure this by looking at the number of challenging chess puzzles that the team solves as a group. These ideas are inspired by studies that have shown that diverse human thinkers outperform homogenous groups on complex tasks, producing “diversity bonuses” ([Page, 2019](#)): improved problem solving, increased innovation, and more accurate predictions—all of which lead to better results. One example is the story of the Netflix prize ([Wired](#)) – an open

competition to predict customers’ movie ratings – where the top two teams worked together to beat the challenge by cleverly combining each other’s algorithms. Other examples of the correlation between heterogeneous teams and diversity bonuses include: teams of three or more songwriters writing the majority of the Billboard 100 hits (Priceonomics), larger and diverse teams outperform individuals in scientific writing (Singh and Fleming, 2010), companies with diverse managements outperforming homogeneous ones (Dezso and Ross, 2012; Hunt et al., 2015) and many more (Page, 2019, Chapter 5).

In the past, chess clubs used to collaborate and play each other via [correspondence chess](#), which led to creative and impressive games. One notable mention is the “Kasparov versus the World” game (Kasparov and King, 2000), where a team of thousands played against Kasparov and decided on each move by plurality vote. The team demonstrated strong chess on which Kasparov wrote that he had “never expended as much effort on any other game in his life”. In a similar vein, chess grandmasters (GMs) often prepare for important games or matches with a team, composed of strong players with diverse styles and qualities (e.g., [GM Magnus Carlsen’s team](#)).

It might surprise some readers that RL can benefit from creative-problem solving mechanisms, since in principle, RL can solve any problem via trial and error. However, RL methods cannot find truly optimal policies in most problems of practical interest and hence find a particular (hopefully) near-optimal solution whose choice is determined by the nature of the agent’s bounds (memory, compute, training experience). This singular sub-optimal choice can potentially generalise poorly when faced with situations it did not encounter during its training. Humans, on the other hand, seem capable of considering multiple good solutions when attempting to solve a new problem instance (Page, 2019), ruling out solutions that had worked before (Stokes, 2005), and accepting the notion of failure (Kasparov and King, 2000). This leads to a kind of creativity through diversity (of solutions) that is missing in current RL agents (see Appendix C for more details).

To test our hypothesis we designed an algorithm to train a league of high-quality, diverse AZ agents, which we call AZ_{db} . The league is represented as a single latent-conditioned architecture (Fig. 1) such that each player i in the league is represented via a latent variable l^i . We encourage diversity via behavioral and response diversity techniques, which complemented each other in other domains (Liu et al., 2021). Behavioral diversity is implemented via intrinsic motivation: each player i receives an intrinsic reward r^i that motivates the agent to visit different positions from the others. Response diversity is implemented via a matchmaker that samples opponent j for each player i .

We find that the players in the AZ_{db} league play chess differently: there’s more variance in whether they castle long or short, develop different pawn structures, prefer different moves in common openings, and more. We collected puzzles from various sources and evaluated the performance of AZ_{db} in these positions. We observed that different players solve different puzzles, that the AZ_{db} team can solve twice as many challenging puzzles as an individual player, and that diverse teams solve more puzzles than more homogeneous teams. We note that getting better at solving chess puzzles may not necessarily improve the overall strength of a chess program, but we believe that improving the creativity and generalization of machines to these human-made “questions” is of great importance to the AI and chess communities.

Lastly, we also see diversity bonuses in chess matchups: The best player in the AZ_{db} league is 30 ELO stronger than AZ indicating that AZ_{db} benefits from diversity bonuses via its training pipeline. But even more importantly, we observe that different players in the league specialize to different openings: With sub-additive-planning AZ_{db} is 50 ELO stronger than AZ.

2. Background

AZ uses a deep neural network $(p, v) = f_\theta(s)$ with parameters θ . This neural network $f_\theta(s)$ takes the board position s as an input and outputs a vector of move probabilities p with components $p_a = \Pr(a|s)$ for each action a , and a scalar value v estimating the expected outcome z of the game from position s ($\mathbb{E}[z|s]$). AZ learns these move probabilities and value estimates entirely from self-play (playing vs. itself); these are then used to guide its search in future games. The board s is represented as an image stack composed of a concatenation of planes of size 8×8 , and also includes metadata about the position (e.g., history/castling rights, etc.). Further details can be found in [Silver et al. \(2018b\)](#), Representation section in the Appendix).

AZ uses a Monte Carlo tree search (MCTS) algorithm as described in [Silver et al. \(2017\)](#). Each state-action pair (s, a) stores a set of statistics, $\{N(s, a), W(s, a), Q(s, a), P(s, a)\}$, where $N(s, a)$ is the visit count, $W(s, a)$ is the total action-value, $Q(s, a)$ is the mean action-value, and $P(s, a)$ is the prior probability of selecting a in s . Each simulation begins at the root node of the search tree, s_0 , and finishes when the simulation reaches a leaf node s_L at time-step L . At each timestep, $t < L$, an action a_t is selected using the PUCT algorithm ([Rosin, 2011](#)):

$$a_t = \arg \max_a Q(s_t, a) + U(s_t, a), \quad (1)$$

where $U(s, a) = C(s)P(s, a)\sqrt{N(s)}/(1 + N(s, a))$, where $N(s)$ is the parent visit count and $C(s)$ is the exploration rate, which grows slowly with search time, $C(s) = \log((1 + N(s) + c_{base})/c_{base}) + c_{init}$. The leaf node s_L is added to a queue for neural network evaluation, $(p, v) = f_\theta(s_L)$. The leaf node is expanded and each state-action pair (s_L, a) is initialized to $\{N(s_L, a) = 0, W(s_L, a) = 0, Q(s_L, a) = 0, P(s_L, a) = p_a\}$. The visit counts and values are then updated in a backward pass through each step $t \leq L$, $N(s_t, a_t) = N(s_t, a_t) + 1$, $Q(s_t, a_t) = W(s_t, a_t)/N(s_t, a_t)$ and

$$W(s_t, a_t) = W(s_t, a_t) + v. \quad (2)$$

The search returns a vector π representing a probability distribution over moves, $\pi_a = \Pr(a|s_{root})$. Further details can be found in [Silver et al. \(2018b\)](#), see the pseudocode in the Appendix). Note that Eq. (2) is highlighted since the only change we make to MCTS is to the value estimate v . AZ is trained by RL from self-play games, starting from randomly initialized parameters θ . Each game is played by running an MCTS from the current position $s_{root} = s_t$ at turn t , and then selecting a move, $a_t \sim \pi_t$, either proportionally (for exploration) or greedily (for exploitation) with respect to the visit counts at the root state. At the end of the game, the terminal position s_T is scored according to the rules of the game to compute the game outcome $z \in \{-1, 0, 1\}$ (for loss, draw or win). The neural network parameters θ are adjusted by gradient descent to minimize the error between the predicted outcome v and the game outcome z , and to maximize the similarity of policy vector p to the probabilities π :

$$(p, v) = f_\theta(s), \ell = (z - v)^2 - \pi^T \log p. \quad (3)$$

3. Methods

On a high level, AZ_{db} is composed of a latent-conditioned architecture (Fig. 1), which is commonly used for modeling diverse policies, see for example, ([Eysenbach et al., 2019](#); [Gregor et al., 2017](#)) for behavioral diversity and ([Liu et al., 2022b](#)) for response diversity. The architecture has a shared torso and three heads: policy, value, and intrinsic value. The policy and the value heads have identical architecture to the ones used by AZ and are trained in the same manner. The agent uses a single set of weights θ to represent the league by conditioning the network on a latent variable l^i that is associated

with each player $i \in \Pi = \{0, \dots, N_{\text{players}} - 1\}$. We represent these latent variables as one-hot vectors and concatenate them to the input as additional constant planes.

Our objective is to discover a set of Quality-Diverse (QD; [Mouret and Clune, 2015](#); [Pugh et al., 2016](#)) policies – policies that play chess well but do so differently from each other. We use response and behavioral diversity techniques, which have been demonstrated to complement each other ([Liu et al., 2021](#)). The next section explains how we implemented behavioural diversity in AZ_{db} and how we combine the policies via sub additive planning. In Appendix B, we further provide details regarding match-making techniques for response diversity. Our results on different match makers were inconclusive and we discuss them in the Appendix for completeness.

Behavioral diversity. To optimize the diversity of the policies in AZ_{db} we focus on behavioural diversity. We refer to the behavior of a policy π as the *state-action occupancy measure* d_π , which measures how often a policy visits each state-action pair when following π . We define a diversity objective $\text{Diversity}(d_{\pi^1}, \dots, d_{\pi^n})$ over the state occupancies of the policies in AZ_{db}. To optimize this objective, we rely on the Reward is Enough principle ([Silver et al., 2021](#)), which states that any objective can be optimized by maximizing a local reward signal. Concretely, we rely on the *intrinsic reward equation*:

$$r_d(s, a) = \nabla_{d_\pi} f(d_\pi)(s, a), \quad (4)$$

which states that one can optimise the function f by locally maximizing its gradient as an intrinsic reward ([Zahavy et al., 2021b](#)). This approach allows us to use the AZ algorithm with only a few changes to maximize diversity. We will now provide more formal details.

Let $\mathbb{P}_\pi(s_t = \cdot)$ be the probability measure over states at time t under policy π , then $d_\pi^{\text{avg}}(s, a) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \sum_{t=1}^T \mathbb{P}_\pi(s_t = s) \pi(s, a)$, $d_\pi^{\gamma}(s, a) = (1 - \gamma) \mathbb{E} \sum_{t=1}^{\infty} \gamma^t \mathbb{P}_\pi(s_t = s) \pi(s, a)$, for the average reward case and the discounted case respectively. It is worth noting that d_π can be calculated for any stationary policy π , and that the RL problem of reward maximization can be written as finding a policy whose state occupancy has the largest inner product with the extrinsic reward (r_e), i.e. $d_\pi \cdot r_e$, which is known as the dual formulation of RL.

We use $\phi(s, a)$ to denote the feature representation of a state-action pair, which represents the chess board after taking action a from state s . It includes the locations of the pieces on the board as well as castling rights (see examples in Fig. 17). Note that the board representation s used as input for AZ and AZ_{db} is slightly different from ϕ , and includes history, side to play and a 50 moves counter (see more details in the background section). We removed these features from ϕ because the first is redundant and the latter two occur in a predictable and relatively independent manner.

Given $\phi(s, a)$, we define $\psi = \mathbb{E}_{s, a \sim d_\pi} [\phi(s, a)]$ to be the expected features under the state occupancy d_π . With this notation, we define our objective as maximizing a non-linear quality-diversity objective over occupancies:

$$\max_{\Pi^n} \text{Diversity}(d_{\pi^1}, \dots, d_{\pi^n}) + \sum_i \lambda_i d_{\pi^i} \cdot r_e, \quad (5)$$

where λ_i is a policy-specific scalar, r_e is the extrinsic reward vector and $\text{Diversity}(\Pi^n)$ is a non-linear utility function that we will soon define. To solve this problem, our agent maximizes a reward that combines the extrinsic reward r_e (win/draw/loss) and a policy-specific diversity intrinsic reward r_d^i :

$$r(s, a) = (1 - \lambda_i) r_d^i(s, a) + \lambda_i r_e(s, a). \quad (6)$$

For simplicity, $\forall i > 0$ we set $\lambda_i = \lambda$, but for policy $i = 0$ we set $\lambda_0 = 1$ so it only maximizes extrinsic reward.

The non-linear diversity utility function is defined over the expected features ψ of the policies in the set. Diverse policies, under this definition, are policies that have different piece occupancies:

$$\text{Diversity}(\Pi^n) = \sum_{i=1}^n (1 - \lambda_i) \left(\underbrace{0.5 \|\psi^i - \psi^{j_i^*}\|_2^2}_{\text{Repulsive}} - \underbrace{0.2 (\|\psi^i - \psi^{j_i^*}\|_2^5 / \ell_0^3)}_{\text{Attractive}} \right), \quad (7)$$

where index j_i^* refers to the player with the closest expected features to player i , *i.e.*, $j_i^* = \arg \min_{j \neq i} \|\psi^i - \psi^j\|_2^2$. The utility function, $\text{Diversity}(\Pi^n)$, is designed to balance the QD trade-off, *i.e.*, to discover policies that are diverse and optimal. Simply maximizing a linear combination of a diversity objective and the extrinsic reward, often results in policies that are either optimal or diverse. Instead, Eq. (7) introduces a hyperparameter ℓ_0 to control the QD trade-off. ℓ_0 defines an equilibrium between two distance-dependent forces: one force that attracts the policies to behave similarly to each other and another that repulses them (Zahavy et al., 2023). According to the *intrinsic reward equation* (Eq. (4)), the gradient of Eq. (7) can be used as an intrinsic reward for maximising it. Differentiating it gives:

$$r_d^i(s, a) = (1 - (\|\psi^i - \psi^{j_i^*}\|_2 / \ell_0)^3) \phi(s, a) \cdot (\psi^i - \psi^{j_i^*}). \quad (8)$$

As we can see, once a set of policies satisfies a diversity degree of $\|\psi^i - \psi^{j_i^*}\|_2 = \ell_0$, the intrinsic reward is zero, which allows the policies to focus on maximizing the extrinsic reward.

To maximize this reward, we learn an intrinsic value function v_d (in addition to the policy and value functions) by first computing an intrinsic value target $z_d(s_t)$ and then adding a loss $(z_d - v_d)^2$ to Eq. (3). To compute the target we only take into account diversity rewards for the player that is currently playing. This is implemented by zeroing out the reward in all of the opponent transitions and using an even number of TD steps N_{TD} when computing the intrinsic value target $z_d(s_t) = \sum_i^{N_{\text{TD}}/2} r_d(s_{t+2i}) + v_d(s_{t+N_{\text{TD}}})$ (as black and white alternate moves in chess).

Similarly, during MCTS, the value estimate v in Eq. (2) is modified in two ways. Firstly, each value estimate is replaced by a combined value estimate: $v = \lambda^i v^i + (1 - \lambda^i) v_d^i$. Secondly, in each MCTS simulation, we accumulate all the intrinsic rewards from the root to the leaf and add them to v . We also zero out the intrinsic reward in states that correspond to the turn of the other player as we did for learning. Furthermore, during MCTS the same latent (associated with the player in the root) is used throughout the entire search, *i.e.*, the planning agent searches against itself. This latent conditions the value, prior and intrinsic value as well as the intrinsic reward. Therefore, the planning player does not have access to privileged information about the opponent.

Sub-additive planning. To choose a player from its team to solve a chess puzzle, AZ_{db} produces N_{players} moves, one for each player. The first method, which we call *max-over-latents*, is an oracle that selects the player who has successfully solved the puzzle. If at least one player has solved the puzzle, the score is one. Teams that propose diverse and creative ideas that are also useful have high *max-over-latents* scores.

The *max-over-latents* method is not always feasible, because we don't always know the solution and therefore cannot select the player who solved it. In these cases, we require a method that chooses a player without seeing the solution. We limit our discussion to sub-additive planning methods that involve two steps. (i) a *planning* step, in which each player $i \in \Pi = \{0, \dots, N_{\text{players}} - 1\}$ performs MCTS and returns the following search-statistics:

$$\eta^{\text{players}} = \{N^i(s, a), Q^i(s, a), U^i(s, a)\}_{i=0}^{N_{\text{players}}-1},$$

where N stands for visitation counts, Q for value estimates, and U for exploration terms (see the background section for more details). (ii) a *selection* step, in which a player is chosen based on the

search statistics η^{Players} . In other words, sub-additive planning uses a heuristic to decide which player has the best chances to solve the puzzle, similar to how humans use heuristics as a way of reducing the complexity of decision making [Gigerenzer \(1991\)](#). These two steps can be written together as:

$$i, \text{ score}_i \leftarrow \text{sub-additive-fn}(\eta^{\text{Players}}),$$

where `sub-additive-fn` outputs a player’s index i and its score in a puzzle score_i . Let $V^j(s) = \max_a Q^j(s, a)$, the player’s selection is based on one of the following rules:

$$\begin{aligned} \text{(Visit)} : i &\leftarrow \arg \max_j \max_a N^j(s, a). \\ \text{(Value)} : i &\leftarrow \arg \max_j V^j(s). \\ \text{(LCB)} : i &\leftarrow \arg \max_j \max_a Q^j(s, a) - U^j(s, a). \\ \text{(Gap)} : i &\leftarrow \arg \min_{j \in \{j: V^j(s) \geq V^j(s) - \text{Gap}(s)\}} V^j(s). \end{aligned} \tag{9}$$

These equations correspond (top to bottom) for selecting a player using visit counts, values, and Lower-Confidence-Bound (LCB) values ([Tirinzi et al., 2023](#)). The LCB is obtained by switching the sign of the exploration term in the PUCT action-selection formula in Eq. (1). The Gap method defines a set of candidate policies – all the policies whose value is within a $\text{Gap}(s)$ from the optimal value – and selects the policy with the lowest value in this set. The gap is usually a small number, so in most cases, this action selection is effectively the same as the Value. We set $\text{Gap}(s)$ to be $\max_{a,j} U^j(s, a)$ in puzzles, and $2/\text{seeds}$ in chess matchups (more details in the experiments). Selecting the $\arg \min$ over the candidates can be viewed as a technique to reduce over estimation ([Van Hasselt et al., 2016](#)).

Note that the cumulative sum of solutions in the team (*max-over-latents*) always upper bounds the performance of any sub-additive planning method (it also explains why these methods are called sub-additive). In addition to that, a good sub-additive method should perform better than the best player in the team. In the case of AZ_{db} this is usually player 0, which is trained without intrinsic reward ($\lambda^0 = 1$). For these reasons, we will always compare sub-additive planning with *max-over-latents* and the performance of player 0. Lastly, we note that in addition to the three sub-additive planning methods in Eq. (9), we tried other approaches such as selecting players based on their raw value estimates, on their entropy of the prior policy etc. The performance of these methods was lower than that of player 0.

4. Experiments

Setup. We trained AZ and AZ_{db} while keeping almost all the hyperparameters identical to those used in ([Silver et al., 2018b](#)). We now explain which hyperparameters were changed in each of our training configurations. Firstly, AZ’s state representation includes history states, which are not available in puzzles (which typically only have a position to solve and not the preceding states). To address this, we trained AZ with history dropout, which randomly removes history states with a probability of 0.2.

The *full configuration* trains a league of 10 players. It uses a batch size of 4096 during training and is trained for 1.6M training steps. The agent is trained with a distributed actor-learner setup. It plays a number of asynchronous self-play games and these games are stored in a replay buffer, along with the policy and raw value targets for each state transition from the game. These targets are used to train their respective estimates. In this configuration, the agent only stores 10% of state transitions from those games (randomly chosen) in the replay buffer and discards the rest. The replay buffer

can store at most 1M positions. At each training step, the agent samples a batch of 4096 samples to make one training update. AZ_{db} uses PSRO-NASH matchmaking algorithm for response diversity and diversity rewards as described in the methods section. More details can be found in the Appendix.

The *fast configuration* is identical to the full configuration, except for the following three changes: the league is of size 5, the batch size during training is reduced to 1024 and all positions generated from self-play games are stored in the replay buffer, *i.e.*, no positions are discarded. This configuration allows an experiment with AZ_{db} to complete in less than a day compared to a week with the full configuration, but has significantly weaker performance in both puzzles and matches. This configuration allows us to conduct a number of careful ablations on AZ_{db} . This configuration uses 1 million MCTS simulations in each position, unless otherwise specified.

To provide a proper AZ baseline, we retrained AZ using these configurations and provide its results when relevant. Concretely, AZ and AZ_{db} in this paper play the same number of games and learn with the same algorithm. For AZ_{db} , those games are divided uniformly between the players. The winrate of AZ against reference players in this paper is the same as in (Silver et al., 2018b). We use the PSRO matchmaker throughout all of our experiments. In Appendix B we present a detailed study comparing the different matchmakers. However, we could not identify a single matchmaker that outperformed the other matchmakers across the puzzle data sets, and the decision to use PSRO is arbitrary. One potential explanation for this is that chess is a transitive game, and therefore, does not require clever matchmaking techniques. Our experiments are designed to answer the following questions:

1. Do AZ_{db} 's policies play chess differently from each other?
2. What makes puzzles hard for AZ and why?
3. Do diversity bonuses emerge in AZ_{db} when solving puzzles and playing chess?

Do AZ_{db} policies play different chess?

To answer this question, we examined the piece occupancies of the different AZ_{db} policies and assessed their opening diversity. Recall that the expected feature vectors $\psi = \mathbb{E}_{s \sim d_\pi} [\phi(s)]$ represent the features that a policy π observes on average when playing chess. The features ϕ represent the chess board and also include castling rights. Our diversity objective is designed to discover policies with different associated vectors ψ (Eq. (8)). Therefore, we now inspect these vectors to gain insights on the behavior of AZ_{db} 's policies.

In Fig. 2 we present standard deviation (across players) of vectors ψ . We also present the raw vectors for all the players and a mean subtracted visualization in the Appendix (Fig. 17 & Fig. 18).




Figure 2 | Left: Standard deviation across players of the piece occupancies of AZ_{db} . Black rectangles highlight the starting position of each piece. Right: average Castling Rights of different players. x-axis corresponds to castling to the Queen (O-O-O) or King (O-O) side as White (W) and Black (B), y-axis corresponds to players. Color units (in the color bars) correspond to occupancy in % multiplied by the average number of moves in a game (65). For orientation, the chess coordinate system is displayed on the left top.

Figures show notable differences in the behavior of the players. Firstly we can see high variability in castling rights on Fig. 2, right. Players 2 & 4 have lower occupancy for castling, which means that they either castle earlier or move the king or rook so they lose their castling rights earlier in the game. Inspecting Fig. 18 in Appendix we can indeed see that when player 2 plays as white, its king occupies g1 for more moves (the square where it gets to after castling kingside). Player 3 on the other hand, castles later, and occasionally castles queenside. In Fig. 18 we see that it spends more time on e1.

Inspecting the pawn structure in Fig. 17, we can see that most players push the e and d pawns to take control over the center. Unlike the other players, player 5 favours the Fianchetto and pushes the g pawn (g2 – g3) to get its kingside bishop to g2 (Fig. 17). We can also see most players push the a and h pawns more than the b, f, g pawns and that the c-pawn is pushed fairly frequently too (the c2 square is often lighter than other squares).

There is also variability across players in the average time that each piece gets to stay on the board (see the black numbers above each sub-figure in Fig. 17). For example, the white queen stays on the board 48% of the game with player 0 as white and 52% with player 4 as black. The rook is the piece that survives the most, 62% with player 3 but only 60% with other players. We can also see that rooks tend to stay on the first rank and the second most common rank for them is the 7th (with lower occupancy). One exception is player 2 that has a rook at the 3rd rank.

To study the **opening diversity** of AZ_{db} we performed two analyses. Firstly, we study the diversity in 8 popular chess opening positions in Fig. 3. Inspecting the results, we can see that AZ_{db}'s players play different moves. All of these moves are played by chess GMs at the highest levels, although some of them are considered niche (eg. 6. Bd3 in the Najdorf, 4. c3 in the Caro-Kann advanced variation). We report the percentage in which each move was played by GMs (black numbers) as well as their relative win rate (blue numbers).




Figure 3 | Opening diversity with AZ_{db}. We selected 8 popular openings where humans tend to play different moves and present with green arrows the moves chosen by AZ_{db} policies with LCB action selection and 100k simulations. For each move, we present in black numbers (top) the % in which this move was played by chess GMs (taken from the Lichess Masters database), and in blue numbers (below them) the relative win percentage (% wins - % losses).

For a quantitative evaluation, we measured the number of different moves played by AZ_{db}'s policies in 3233 openings from the Encyclopaedia of Chess Openings using [lichess-org chess-openings](https://lichess.org/chess-openings) (we excluded irregular openings (A00)). AZ_{db} discovered 1.26 ± 0.01 , 2.45 ± 0.02 , and 1.65 ± 0.01 moves on average in each position, for 1, 400, 100k simulations respectively. Note that AZ_{db} discovered the most opening moves when it used 400 simulations, which is the number of simulations it used for training. We further discuss this finding at the end of this section.

Why are puzzles hard for AZ? In the previous section, we observed that the players in the AZ_{db} league play chess differently from each other. We would like to evaluate these players on challenging puzzles and examine if diversity bonuses emerge. But before we do so, we examine why chess puzzles make a challenge for computer chess agents like AZ.

For this goal, we have collected two data sets of challenging puzzles that can trick modern chess engines. The *Challenge set* includes 15 multi-step puzzles that are challenging to modern chess engines, and the *Penrose set* includes 53 challenging puzzles including fortresses and Penrose positions to measure the value accuracy of AZ. More details about these data sets can be found in the Appendix.

In the first row of the top panel in Table 1 we can see the solve rate of AZ with 1M MCTS simulations in these data sets; it only solved 11.76% of the Challenge set and 3.64% of the Penrose set. Although AZ is a very strong chess player, these puzzles confuse it. We now attempt to understand what makes these puzzles difficult for AZ through a few experiments and visualizations. Note that for all the experiments in this section, we used the fast configuration of AZ (see setup at the beginning of this section).

Our first hypothesis is that these puzzle sets present a hard generalization challenge for AZ, highlighting its weaknesses and blind spots. To verify this prediction, we allow AZ to begin self-play games from puzzle positions. It is important to note that the solution is not provided to AZ; only the starting position of the puzzle. By presenting the position to AZ, it can learn to solve the puzzle by trial and error and potentially improve at doing so via its RL pipeline. In the second row of the top panel in Table 1 we can see the solve rate of AZ when it begins self-play games from the puzzle positions. We can see that it solved most of the Penrose set and significantly improved its performance on the Challenge set (from 11.76% to 76.47%).

We further explore if this generalization challenge is in distribution generalization or OOD generalization in the middle panel of Table 1. For in distribution generalization, we randomly

	Challenge Set	Penrose Set
(1) AZ	11.76%	3.64%
(2) Self-play from puzzles	76.47%	96.23%
(3) In distribution	—	88.24%
(4) Out of distribution	—	11.76%
(5) + intermediate positions	88.23%	96.23%
(6) + exploration	94.11%	96.23%
(7) + half-move clock	94.11%	100%

Table 1 | Self-play from puzzle positions. From top to bottom. (1) AZ. (2) AZ with self-play from puzzle positions. (3) Randomly splitting the Penrose set to train and test sets. (4) Training on the first Penrose position & testing on the second (OOD). (5-7) AZ with self-play from puzzle positions and additional exploration techniques.

split the Penrose set to train and test sets. We allow AZ to begin self-play from the train set and evaluate its performance on the test set. In the OOD case, we split the Penrose set between variations of the two Penrose positions and evaluate on one of them (details in Appendix). We can see that AZ generalizes pretty well in distribution while it struggles to generalize to OOD positions.

However, it remains unclear what prevents AZ from reaching 100% solve rate when it begins self-play games from these puzzle positions. Recall that many of these puzzles involve a long sequence of actions and an incorrect action at any point results in the agent failing to solve the puzzle. Thus, we hypothesize that these positions also present an exploration challenge to AZ. To verify this second hypothesis, we performed three more experiments which we present in the lower panel of Table 1. We only briefly explain them here and refer the reader to Appendix C.3 for more details. (1) we extended the puzzle positions from the Challenge set to also include intermediate positions (76.47% \rightarrow 88.23% on the Challenge set), (2) we change AZ’s action selection mechanism to sample actions from a softmax distribution over the MCTS visits (88.23% \rightarrow 94.11%) and (3) we manually augmented the half-move counter (100% in the Penrose set). Each one of these changes have improved AZ’s performance on the challenging puzzles until it finally solve all but one of them.

We note that these puzzles were collected and composed by chess players deliberately to trick computers and it’s not likely that AZ had seen such positions when it learned to play chess by playing against itself. Our findings in this section suggest that the challenges for AZ are OOD generalization and exploration. In the next section we will inspect if different players in the AZ_{db} league generalize differently and solve more puzzles as a result. However, we would like to emphasize that there is no fundamental or incomputable problem to solve these puzzle positions. The challenges are OOD generalization and exploration which are well known challenges in AI.

AZ’s understanding of puzzles. In Fig. 4 and Fig. 19 in the Appendix, we visualize AZ’s policy $p(s)$, raw value $v(s)$, and MCTS value estimates $Q(s, a)$ for four puzzle positions. We compare the estimates of AZ with the fast configuration with the variation of AZ that begins self-play games from the puzzle positions. Fig. 4 shows two Penrose positions from the Penrose set and Fig. 19 in the Appendix shows two steps from a puzzle from the Challenge set (recall that these are multi-step puzzles and to get a score of one the agent has to find the correct solution in all the steps). The right column in these figures shows three bar plots: the prior policy’s probability p of legal actions from the given position (top), the raw values v (center), and the MCTS values Q (bottom). We use blue arrows and blue x-ticks to indicate correct actions, and red for incorrect ones.

Inspecting the first row in Fig. 4, we see one of the Penrose positions. It is white’s turn to play, and there are only two legal actions available (Kb2 and Kb1). We can see that the prior is roughly uniform across the legal actions, which is fine since both of them are correct. However, when we inspect the raw values (shown in black bars) we see that AZ estimates the white player to win when it plays either of these two actions. After MCTS, the value estimate reduces and predicts a draw, which is the correct evaluation. The self-play from puzzles agent (blue) accurately predicts the position to be a draw in both raw and MCTS values indicating better understanding.

The second puzzle in Fig. 4 is the original Penrose position (white’s turn to play). The solution to this puzzle is to move the white king to any of its neighboring white squares (blue color) so that it can eventually draw the game. Any other move leads to a losing position (red color). Inspecting the figure, we see similar findings to the previous puzzle: AZ is confused regarding the correct move to play. The prior puts an equal probability on any of the legal actions (c.f. prior barplot), raw value predicts a win (c.f. raw value barplot), and MCTS value predicts a loss (c.f. MCTS value barplot). The self-play from puzzles agent, on the other hand, have accurate raw and MCTS values for this position (a draw for the correct actions and a loss for the incorrect actions), demonstrating understanding.



Figure 4 | *Left*: Two Penrose puzzle positions from the Penrose set. *Right*: Visualizations of prior probability from a subset of legal actions, and its corresponding raw value and MCTS value estimates for AZ and AZ trained on puzzle sets. Correct and Incorrect moves are colored in Blue and Red respectively.

To summarize, in Fig. 4 (and Fig. 19 in the Appendix), we see that AZ fails to understand positions from the Penrose and the Challenge sets when trained with the fast configuration. AZ’s ability to understand difficult chess positions (in terms of prior, raw value, and MCTS value) improves when it is allowed to train from those puzzle positions, which also translates to better performance on those puzzle sets.

Diversity bonuses in AZ_{db}. In the previous sections we have seen that AZ_{db} trains a league of strong chess players that play chess differently. We have also seen that chess puzzles are a challenging problem for strong chess agents like AZ. In this section we examine if diversity bonuses emerge in AZ_{db} when solving chess puzzles. We then let AZ_{db} to play chess games against AZ from different opening positions and see if diversity bonuses emerge there and conclude with ablative analysis on diversity bonuses. From this point on, the puzzle sets are not available to any of the agents and they are not allowed to start self-play games from them.

Diversity bonuses in solving puzzles. In this section we compare the performance of AZ and AZ_{db} on four data sets. These include the Challenge and Penrose sets from before, and two larger scale data sets. The Strategic test suite (STS) composed of 1000 multiple choice puzzles and the Lichess data set, composed of 500 multi-step puzzles from lichess.org. More details about the data sets can be found in the Appendix.

We would like to see if diversity bonuses emerge in the upper limit of AZ’s computational rationality. Therefore, in this section we use the full configuration of AZ_{db} and allow it to search for 100M simulations per position – the largest search budget we could afford without running out of memory due to the size of the search tree. We also use the LCB action selection rule unless mentioned otherwise.

Inspecting Fig. 5, we can see AZ_{db}’s player 0 performs roughly as well as AZ on most of the data




Figure 5 | Puzzle solve rates.

sets. More importantly, we can observe strong diversity bonuses on each one of the data sets: AZ_{db} manages to solve many more puzzles as a team via sub-additive planning and *max-over-latents*. In particular, it solves more than twice more puzzles from Challenge set. In Penrose set, it solves the two notoriously hard Penrose positions and more generally, 56.6% of the puzzles. In Lichess and STS we observe that *max-over-latents* solves 3% more puzzles than player 0, and that sub-additive planning is 0.5% better. Note that these data sets are relatively easier (taken from human games), and AZ_{db} is already solving them at high percentage, thus the relative gains are actually quite significant.

Diversity bonuses in chess matchups.

In this section, we assess the impact of diversity bonuses on the performance of AZ_{db} in chess matches. We matched each player in the AZ_{db} league against AZ and let them play chess from 74 opening positions. For each opening, AZ played one game as black and one as white. The opening positions were carefully chosen by GM Matthew Sadler to favor one player by giving them a strong opening advantage (up to 1500 centipawns). This was done in order to produce more decisive matches between agents, with one player winning rather than the match ending in a draw. The games were played in a bullet time control, where each agent used 400 MCTS simulations to select a move in a match. This roughly translates to 100 milliseconds of thinking time per move. Recall that the same amount of MCTS simulations were used in training AZ_{db}.

Table 2 | Summary of chess matchups between AZ_{db} and AZ. The standard error, in winrate, is 0.34 (for each Player) up to quantization error of 0.01.

	Player 0	<i>Sub-additive planning</i>	<i>Max-over-latents</i>
Winrate	54.25	57.18	58.6
Elo	29.6	50.3	60.3

Table 2 presents a summary of the games in terms of overall win rate and relative Elo score for player 0, *Max-over-latents* & *Sub-additive planning*. Fig. 21 in the Appendix additionally presents the number of wins, draws and losses (as black and as white) for all the players. These statistics are obtained by averaging the results from each position over 500 random seeds. The win rate and relative Elo improvements show that most of the players in the AZ_{db} league are stronger than AZ.

AZ_{db} (white) vs AZ (black)

French Winawer C16	0.88	0.88	0.85	0.84	0.84	0.88	0.82	0.81	0.83	0.82	0.83	0.86
Benoni Wall A56	0.7	0.7	0.67	0.7	0.65	0.68	0.59	0.66	0.66	0.68	0.52	0.66
Blumenfeld E10	0.45	0.45	0.34	0.32	0.31	0.27	0.41	0.43	0.45	0.4	0.4	0.29
Magnus Sicilian B23	0.57	0.57	0.48	0.57	0.49	0.46	0.44	0.44	0.5	0.41	0.46	0.46
Trompovsky A45	0.3	0.3	0.19	0.15	0.3	0.26	0.15	0.02	0.09	0.16	0.12	0.15
	max	SA	I=9	I=8	I=7	I=6	I=5	I=4	I=3	I=2	I=1	I=0

Figure 6 | Results from a subset of chess openings between AZ_{db} and AZ (scores averaged from 500 random seeds) where AZ_{db} played as White and AZ as Black. Results from all openings are presented in the Appendix.

This is likely due to implicit diversity bonuses, such as seeing more diverse data during training and playing against each other. In particular, player 0, the strongest player, achieves a win rate of 54.25% and a relative Elo of 29.6 compared to AZ. It makes sense that player 0 is the strongest in the league, as it is a “special player” that is trained without the diversity intrinsic reward.

We can also see that the *max-over-latents* method performs better than player 0. But unfortunately, *max-over-latents* is not practical for selecting a player to play in an opening position, since we do not know the outcome of the game in the opening (when we select the player to play). On the other hand, sub-additive planning is practical, and we now discuss how we implemented it. For each opening and seed, we calculated the average score of each player across all other seeds (leaving the current seed out) and selected the player with the highest score. We then reported the score of this player against the current hold-out seed (more details in the Appendix). Table 2 shows the results for the Gap sub-additive planning method (Eq. (9)), where the gap is defined to be $2/500$. This value assures that the candidates are all the policies that can become optimal if we add one more game in which they win and the optimal policy loses. As we can see, this improves AZ_{db} ’s win rate to 57.18%, which translates to a 50.3 ELO improvement over AZ. The results for the other sub-additive rules were 40 for the value and visit counts, and 41 for the LCB, *i.e.*, they all improve by around 10 Elo over player 0, but the gap method was better than all of them.

In Fig. 6 we further inspect these results. We report the average score (a win is scored as 1, draw is 0 and loss is -1) of the players in AZ_{db} when playing as white in a subset of openings. The results on all the opening positions can be found in the Appendix (Fig. 22). We can see that the performance of the players in AZ_{db} vary across openings. In *French Winawer C16*, all the players perform roughly the same. However, in other openings, there are specific players that perform significantly better than others. For example, player 0 is strong in *Benoni Wall A56*, but is weak in comparison to the other players in *Blumenfeld E10*. While player 0 is the overall best player, they are not the best player in any specific opening. They are also not the best player when playing as white or black. Table 3 in the Appendix shows the average win rate of each player when playing as white and black. Player 0 won 68.27% as white and 40.21% as black. However, Player 2 had the highest win rate as white (69.86%), and player 9 as black (40.38%). These results suggest that different players specialize in different openings and that sub-additive planning can be used to select the better player in each opening. Further, Fig. 20 in the Appendix shows how sub-additive planning performs as a function of the number of seeds.

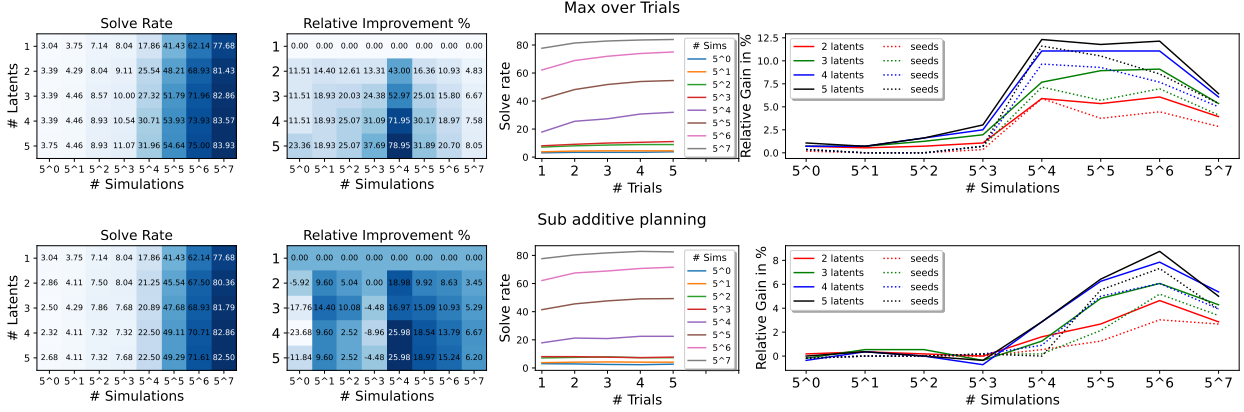


Figure 7 | Scaling laws with AZ_{db} . Top: Max over trials, Bottom: Sub-additive planning. Left to right: (1) Solve rate in % on Lichess puzzles with a different number of simulations and latents. (2) Relative gains in % on increasing the number of latents for each simulation budget. (3) scaling with the number of latents for different simulation budgets. (4) Relative gains with a different number of trials – latent in solid lines, seeds in dashed lines – at different simulation budgets.

Analysis of diversity bonuses in AZ_{db} . In the puzzle evaluation section, we observed that diversity bonuses emerge at the computational boundaries of AZ_{db} . In this section, we analyse what components of AZ_{db} are the most important for diversity bonuses and if diversity bonuses emerge at other compute budgets. We focus our evaluation on the Lichess data set and use the fast configuration for that. In Fig. 7 we study how AZ_{db} scales with different simulation budgets and team sizes. The top figure presents results for *max-over-latents* and the bottom figure shows results for sub-additive planning based on LCB (Eq. (9)). Most importantly, we observe diversity bonuses for all compute budgets and team sizes. In the leftmost table, we present the absolute solve rate in %. We can see that AZ_{db} 's performance improves monotonically with the number of simulations and the number of trials, implying that larger teams solve more puzzles together. We can also see that on the third sub-figure, which presents the same data in a different manner. For each column in the first table (simulation budget) we draw a line that presents the solve rate in % as a function of the number of trials (x-axis). We can see that the diversity bonuses keep increasing as we increase the number of trials for each simulation budget.

The second table from the right shows the relative gains computed in %¹ of AZ_{db} from having more players in the team. Interestingly, the highest relative gains are achieved when the number of simulations is $5^4 = 625$, which is the simulation budget that is closest to the one we use in training (see the discussion in the next section). On the rightmost sub-figure, we compare a diverse team of agents with a more homogeneous team. For the diverse team, we use different latents as before (solid lines) and for the homogeneous group, we use the best latent in the group (latent 0) and allow it different trials of search (with different seeds, in dashed lines). We can see that across different group sizes, simulation budgets, and for both *max-over-latents* and sub-additive planning, a diverse team outperforms the homogeneous one.

We further study diversity bonuses in a few variations of AZ_{db} and compare different sub-additive planning methods (counts, value, and LCB) in Eq. (9). We also compare AZ_{db} with a variation that only uses the extrinsic reward and extrinsic value function at test time (dashed lines). Inspecting Fig. 8, we can see that using the intrinsic reward and value function (solid lines) receives higher diversity bonuses across different action selection rules and simulation budgets. We also see higher

¹we divide each row by the first row (player 0), subtract 1 and multiply by 100

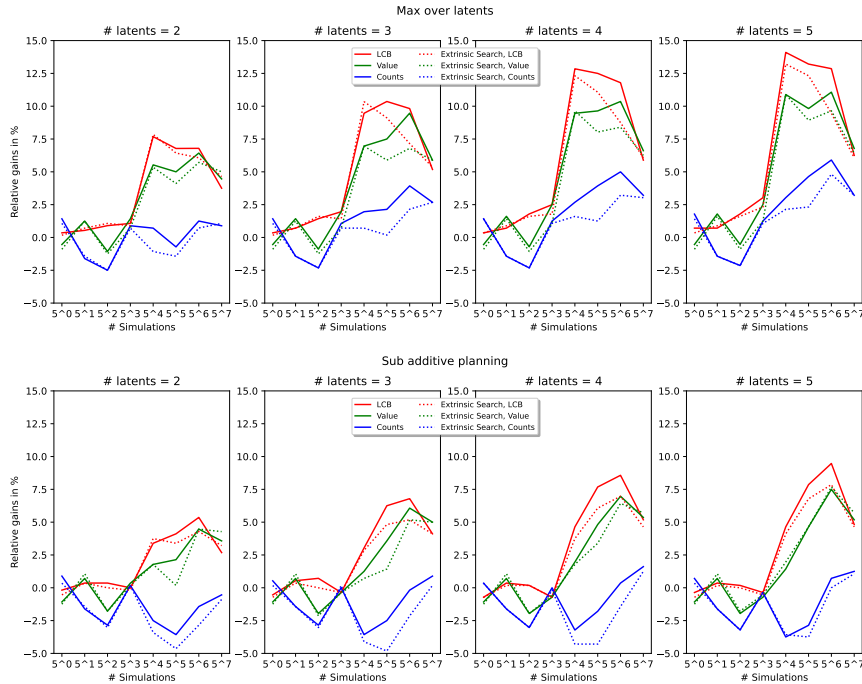


Figure 8 | An ablation of action selection rules (color) and search type (solid/dashed) with *max-over-latents* (top) and sub-additive planning (bottom). Sub figures correspond to the number of trials.

diversity bonuses when selecting actions based on value and LCB rather than counts.

Discussion. AZ_{db} benefited from diversity bonuses regardless of the amount of simulations it was given. However, it gained the most from diversity in solving puzzles when it was given 625 simulations (see Fig. 7), and it discovered the most unique openings when it was given 400 simulations. One possible explanation for these findings is the *spinning top hypothesis* (Czarnecki et al., 2020), which states that the largest diversity of a population of players is observed at a medium level of play, and at the low and high levels of play there is less room for diversity. Indeed, Sanjaya et al. (Sanjaya et al., 2022) observed that in chess, human play exhibits the most diversity in the medium levels of play (1600 ELO) compared to the higher and lower levels. A different possible explanation is that AZ_{db} 's diversity overfits to the number of simulations it uses for training (400), in the sense that it achieves higher diversity bonuses when it uses the same number of simulations it used for training compared to other simulation budgets. Note that AZ 's and AZ_{db} 's performance improves monotonically with the number of simulations and we do not suggest that the performance overfits to the number of simulations, only that the diversity does. To better understand this explanation, recall that the expected features ψ in our diversity objective are estimated from the average games played by AZ_{db} , and since these games are played with 400 simulations, it is possible that the diversity objective itself overfits to the behavior of AZ_{db} 's policies with 400 simulations.

5. Related work

Chess has been a long-standing subject of AI research, and is often called the “Drosophila of AI” (Kasparov, 2018). The first chess program was written by Alan Turing on slips of paper, before the invention of the computer. Achieving human-level performance in chess was known as the “chess Turing test”, and computer scientists made significant progress towards this goal until Deep Blue finally beat the world chess champion in 1996. In his book “Deep Thinking” (Kasparov, 2017) Kasparov tells his side of the story of his loss to Deep Blue. He describes the experience of strategizing

against an opponent that is tireless and never makes mistakes, and the mistakes he made that led to his defeat. While some have expressed concerns that chess will become less popular as machines become stronger than humans, chess is now more popular than ever due to an increase in online play during the COVID-19 pandemic, and the success of the Netflix series “The Queen’s Gambit” and streamers in bringing attention to the sport ([WUFT](#)).

In 2018, AZ made a significant impact on the world of chess by learning to play chess on its own, without any human knowledge. The games between AZ and SF drew a lot of attention from the chess community, on which Kasparov wrote that “Deep blue was the end, AlphaZero is a beginning”. In fact, almost all modern chess engines today have followed AZ’s lead and use a neural network score function instead of a human-designed one. Professional chess players also changed their playing style after the AZ vs SF games were introduced ([González-Díaz and Palacios-Huerta, 2022](#)), prototyped new variants of chess using AZ ([Tomašev et al., 2020](#)), and investigated when and where human chess concepts are represented in AZ ([McGrath et al., 2022](#)).

Despite the strength of chess engines, it is uncertain whether we are near optimal play, as evidenced by the results of computer chess competitions like the [TCEC](#), where the strength of the top engines continue to increase with each season. Another piece of evidence is that modern chess engines tend to make blunders when constrained to use a limited amount of search. A famous example is the winning move in game eight of the Classical World Chess Championship 2004 between Peter Leko and Vladimir Kramnik ([chessbase](#), [Agadmator’s Chess Channel](#)): “*Grand Master (GM) Vladimir Kramnik relied a bit too much on the engine in his Marshall Gambit preparation. The engine initially told Kramnik he was winning, but Leko let his engine calculate more deeply and realized that his mating attack would triumph over Kramnik’s passed pawn*” ([Copeland, 2021a](#)). While this example is from 2004 and nowadays chess engines can easily solve it, we can find similar cases on “chess twitter” on a regular basis. For example, GM Anish Giri explains that for similar reasons “*top players use cloud engines and not the free stuff available online*” ([Twitter](#)).

In order to beat superhuman computers, chess players developed [anti-computer tactics](#), that exploit blind spots in how computers reason about chess. Chess players have been using these tactics against chess programs in their early days, but successful examples became fewer as machines became better. One notable exception is the achievements of GM Hikaru Nakamura against modern chess engines. In 2008, Nakamura played against the top engine at the time, Rybka, and managed to close the position, offer sacrifices, and convince the engine to overpress and eventually lose ([Copeland, 2021b](#)); more recently Nakamura used similar tactics against [Mittens](#) and secured a draw. Another example of the limitations of computer chess programs is that they cannot solve certain positions ([Copeland, 2021a](#); [Doggers, 2017](#); [Penrose, 2020](#)). To address this, scientists have developed puzzle-specific chess programs like [Crystal](#) and worked on specific methods for specific positions like fortresses ([Guid and Bratko, 2012](#)).

In this work, we investigated whether creative problem-solving mechanisms identified in humans can make machines more creative and help them solve such positions. This was inspired by Kasparov’s statement that “creativity has a human quality: it accepts the notion of failure” ([Kasparov, 2017](#)). We focused on diverse policy discovery, since diverse teams are known to be more creative and robust to failures.

Although diverse policy discovery has not been extensively studied in chess, there are some related mentions. For example, chess engine designers make different choices that can lead to different playing styles. Lc0 and AZ use MCTS, while SF uses alpha-beta pruning. Even the rules of chess have some references to diversity, such as the tension between the bishop and the knight ([Shead, 2020](#)). Chess opening theory is also diverse, as many opening positions have multiple possible variations. As an example, different AZ seeds prefer different moves in the Ruy Lopez opening ([McGrath et al.,](#)

2022). GM Magnus Carlsen has also commented on the importance of finding ideas that are missed by other players in the [Lex Fridman Podcast](#): “Now it’s all about finding ideas that are missed by agents, either that they are missed entirely or missed at low depth. These are used for surprising your opponent with lines in which you have more knowledge”.

In AI, the concept of diversity dates back to Turing’s cultural search in “Intelligent Machinery” (Turing, 1948): “it is necessary for a man to be immersed in an environment of other men, whose techniques he absorbs, ..., he may then perhaps do little search on his own to make new discoveries which are passed to other men”. Turing listed cultural search as one of three types of search in which machines can potentially learn. Since then, diversity had been studied in different sub topics. The [Quality-Diversity web page](#) provides an excellent overview of the work done in the evolutionary learning community (Cully and Demiris, 2017; Mouret and Clune, 2015; Pugh et al., 2016). In RL, research had been conducted on intrinsic rewards for diversity (Eysenbach et al., 2019; Gregor et al., 2017), on repulsive and attractive forces (Flet-Berliac et al., 2021; Liu et al., 2017; Vassiliades et al., 2017; Zahavy et al., 2023), and on learning a behavior that maximizes a general utility function (Geist et al., 2022; Hazan et al., 2019; Zahavy et al., 2021b; Zhang et al., 2020). There had also been a significant amount of study on diversity in the multi-agent community (Balduzzi et al., 2019; Czarnecki et al., 2020; Liu et al., 2022b; Perez-Nieves et al., 2021; Vinyals et al., 2019). We refer the reader to the Appendix for further discussion about these works.

Lastly, sub-additive planning had been investigated in AI as well as in other fields. In psychology, heuristics is the process by which humans use mental short cuts to arrive at decisions. Heuristics are simple strategies that humans, animals, organizations, and even machines use to quickly form judgments, make decisions, and find solutions to complex problems. Heuristics are also considered to be an evidence for bounded rationality in human decision making. In RL, a commonly used heuristic is General Policy Improvement (GPI; Barreto et al., 2017b), which is similar to sub-additive planning based on value. However, in special situations GPI is sub optimal, and can be improved (Zahavy et al., 2020b). Similarly, our results suggest that the LCB and Gap methods performed better than value.

6. Summary

This paper investigated whether AI can benefit from diversity bonuses, inspired by how teams of diverse thinkers outperform homogenous teams in challenging tasks. We tested this hypothesis in the game of chess, the “drosophila of AI”, which poses computational challenges for AI systems due to the large number of possibilities that have to be calculated. We proposed AZ_{db} , a league of AlphaZero (AZ) agents that are modeled via a latent-conditioned architecture, and trained its players to play chess differently via behavioral diversity techniques and observed that they choose different pawn structures, castle at different times, and prefer different moves in the opening. In a chess tournament against AZ, we observed that different players in AZ_{db} specialize in different openings, and, that selecting a player from AZ_{db} via sub-additive planning leads to an improvement in Elo points.

We also found that diversity bonuses are present in solving puzzles regardless of the computational budget. In particular, they persisted even when the computational budget of a single search player was pushed to its upper computational rationality limit of 100M simulations (using more simulations resulted in out-of-memory problems). In this setting, we observed that AZ_{db} achieved twice the solve rate of AZ on the hardest of our data sets (the Challenge set), suggesting that diversity bonuses emerge even in this limit. We note that physical limitations bound humans, and some hardware constraints bound most AI agents. While it should be possible to improve these memory constraints, there should always be some limit on which we can achieve with existing hardware, setting a different bounded rationality limit (Lewis et al., 2014; Wolfram, 2023).

Some of the hardest chess puzzles we considered were created by Sir Roger Penrose. Penrose

argued that human creativity is not entirely algorithmic and cannot be replicated by a sufficiently complex computer (Penrose, 1994; Penrose and Mermin, 1990; Wikipedia, 2023). By understanding why humans can solve puzzles that AI cannot, we may be able to uncover the non-algorithmic aspects of human thinking and gain a better understanding of human consciousness (Doggers, 2017). While Fig. 4 does show that AZ does not understand these positions, it also suggests that AZ can solve almost any puzzle using a simple computational process: learning from trial and error via RL by starting chess games from puzzle positions and playing against itself. In other words, Fig. 4 demonstrates that the Penrose positions are computable in the sense presented in (Penrose, 1994) and implied by the finiteness of chess, but also, that there are indeed gaps between human and machine thinking. Through careful analysis, we suggested that the challenge for AZ is that it does not see similar positions when it learns to play chess by playing against itself, and therefore, struggles to generalize to these unseen, OOD positions. In contrast, each policy in AZ_{db} is intrinsically motivated to experience different games, making it less likely that a specific position will be OOD for any of its policies. Indeed, in Fig. 5 we observed that some of the players in the AZ_{db} team were able to solve the Penrose positions without training on them, while other players, including AZ, were not. This suggests that incorporating human-like creativity and diversity into AZ can improve its ability to generalize. On the other hand, there were still many puzzle positions that AZ_{db} failed to generalize to as a team, suggesting that there is still a gap between human and machine thinking. We hope that our paper will inspire future research of these gaps.

7. Acknowledgements

We would like to thank Nolan Bard, Daniel J. Mankowitz, Khimya Khetarpal, Siqi Liu, Ulrich Paquet and Thomas Hubert for discussions and feedback on this work; Hussain Masoom for strategic advice and support in organization; Will Dabney and Jack Parker-Holder for reviewing the manuscript; Julian Schrittwieser for reviewing code and advising; and Grand Master Matthew Sadler for discussions and feedback on chess.

A. Additional related work on QD

QD optimization is a type of evolutionary algorithm that aims at generating large collections of diverse solutions that are all high-performing (Cully and Demiris, 2017; Pugh et al., 2016, QD). It comprises two main families of approaches: MAP-Elites (Cully et al., 2015; Mouret and Clune, 2015) and novelty search with local competition (Lehman and Stanley, 2011, NSLC), which both distinguish and maintain policies that are different in the behavior space. The main difference is how the policy collection is implemented, either as a structured grid or an unstructured collection, respectively. The behavior space is either handcrafted based on domain knowledge (Cully et al., 2015; Tarapore et al., 2016) or learned from data (Cully, 2019). Some other works combine evolutionary QD with RL (Nilsson and Cully, 2021; Pierrot et al., 2022; Tjanaka et al., 2022). Further references can be found on the [QD webpage](#).

In RL, intrinsic rewards have been used for discovering diverse skills. The most common approach is to define diversity in terms of the discriminability of different trajectory-specific quantities and to use these ideas to maximize the Mutual information between states and skills (Baumli et al., 2021; Eysenbach et al., 2019; Gregor et al., 2017; Sharma et al., 2020) or eigenvectors of the graph Laplacian (Klissarov and Machado, 2023). Other works implicitly induce diversity to learn policies that maximize the set robustness to the worst-possible reward (Kumar et al., 2020; Zahavy et al., 2021a), and some add diversity as a regularizer when maximizing the extrinsic reward (Badia et al., 2020; Hong et al., 2018; Masood and Doshi-Velez, 2019; Pathak et al., 2017; Peng et al., 2020; Sharma et al., 2020; Sun et al., 2021; Zhang et al., 2020).

Parker-Holder et al. (Parker-Holder et al., 2020) measure diversity between behavioral embeddings of policies. The advantage is that the behavioral embeddings can be defined via a differential function directly on the parameters of the policy, which makes the diversity objective differentiable w.r.t to the diversity measure. Other works use successor features to discover and control policies (Alver and Precup, 2022; Barreto et al., 2017a, 2018, 2019; Nangue Tasse et al., 2020; Tasse et al., 2021; Zahavy et al., 2021a). In this line of work, there is a sequence of rounds, and a new policy is discovered in each round by maximizing a stationary reward signal (the objective for each new policy is a standard RL problem). The rewards can be random vectors, one hot vectors, or an output of an algorithm (for example, an algorithm can be designed to increase the diversity of the set). There is also a growing interest in techniques for solving MDPs with general utilities (Abbeel and Ng, 2004; Belogolovsky et al., 2021; Geist et al., 2022; Hazan et al., 2019; Mutti et al., 2022; Shani et al., 2022; Zahavy et al., 2020a, 2021b; Zhang et al., 2020) which can be used for designing diversity objectives.

Repulsive and attractive forces have also been studied in related work. Vassiliades et al. (Vassiliades et al., 2017) suggested to use Voronoi tessellation to partition the feature space of the MAP-Elite algorithm to regions of equal size and Liu et al. (Liu et al., 2017) proposed a Stein Variational Policy Gradient with repulsive and attractive components. Flet-Berliac et al. (Flet-Berliac et al., 2021) introduced an adversary that mimics the actor in an actor-critic agent and then added a repulsive force between the agent and the adversary.

There has also been a significant amount of work in **multi-agent** learning on discovering diverse policies. A common approach in this literature is to encourage response diversity via matchmaking to increase the space of strategies (Balduzzi et al., 2019; Perez-Nieves et al., 2021), which is particularly useful in non-transitive games. Sanjaya et al. (Sanjaya et al., 2022) quantified the non-transitivity in chess through real-world data from human players over one billion chess games from Lichess and FICS. Their findings suggest that the strategy space occupied by real-world chess strategies demonstrates a spinning top geometry (Czarnecki et al., 2020). These findings are correlated with our observation that there is more diversity in AZ_{db} play style when it uses a medium number of simulations (400)

and less when it uses a high or low number of simulations.

Response diversity techniques are usually combined with population-based training (Jaderberg et al., 2017), which represents a collection of players with a separate set of parameters for each player. One important example is the Alpha League, which led to significantly improved AlphaStar’s performance in Starcraft 2 (Vinyals et al., 2019). More recently, response diversity has also been used with a latent conditioned architecture (Liu et al., 2022a,b), closer to our setup.

Another important use case for diversity in multi-agent learning is to first discover diverse high-quality policies and then train the best response policy that exploits all of the discovered policies (Lupu et al., 2021). In another example, Liu et al. (Liu et al., 2017) considered diversity in Markov games and proposed a state-occupancy-based diversity objective via the shared state-action occupancy of all the players in the game.

B. Response diversity

B.1. Technical details

We encourage response diversity in a league of players Π by training on matchups sampled from a probability function $\mathbb{P}(i, j, t^i)$. Intuitively, the matchmaking works by first uniformly sampling player i from Π . It then uniformly samples an indicator t^i to denote whether player i is playing as black or white, and the oponent player j is sampled conditionally as:

$$\mathbb{P}(i, j, t^i) = \mathbb{P}(j|i, t^i)\mathbb{P}_{\text{unif}}(i)\mathbb{P}_{\text{unif}}(t^i) \quad (10)$$

Note that chess is not a symmetric game. There is "a general consensus among players and theorists that the player who makes the first move (white) has an inherent advantage" (Wikipedia). Hence we condition the sampling of j on both i and t^i .

Several population learning frameworks can be expressed this way. In practice, $\mathbb{P}(j|i, t^i)$ is represented as the adjacency matrix of an interaction graph $\Sigma \in \mathbb{R}^{N \times N} = \{\sigma^i\}_{i=1}^N$, where each row σ^i corresponds to a mixture policy over Π , and the j^{th} element is the probability of i playing against j . For example, in Policy Space Response Oracle (PSRO; Lanctot et al., 2017), each agent is computed as the approximate best response to the Nash equilibrium over players in the current population for some fixed iteration order. Specifically, we consider the following matchmaking algorithms, also illustrated in Fig. 9:

- **Selfplay** each player only plays against itself, *i.e.*, $j = i$
- **Uniform** player j is sampled uniformly from Π , *i.e.*, $\mathbb{P}(j|i, t^i) = \mathbb{P}_{\text{unif}}(j)$
- **PSRO-NASH** player j is sampled from a mixture policy σ_b^i or a mixture policy σ_w^i , which are the Nash equilibria over $\{\pi_k\}_{k=0}^i$ playing as (b)lack and (w)hite respectively. For a real example of interaction graphs for a PSRO-NASH league, see Fig. 10.
- **PSRO-RECTIFIED** is the same as **PSRO-NASH**, except we respond to the “rectified” Nash as per (Balduzzi et al., 2019), *i.e.*, each player is trained against the Nash weighted mixture of players that it beats or ties, in order to discover strategic niches and enlarge the empirical gamescape.
- **Fictitious Play** or PSRO-UNIFORM or (Brown, 1951; Leslie and Collins, 2006): player i responds to the uniform mixture over the current population of players. This has been shown to perform well empirically.
- **PSRO-CYCLE** the first $n - 1$ players forms a strategic cycle and the last player responds to the Nash equilibrium over the entire league.

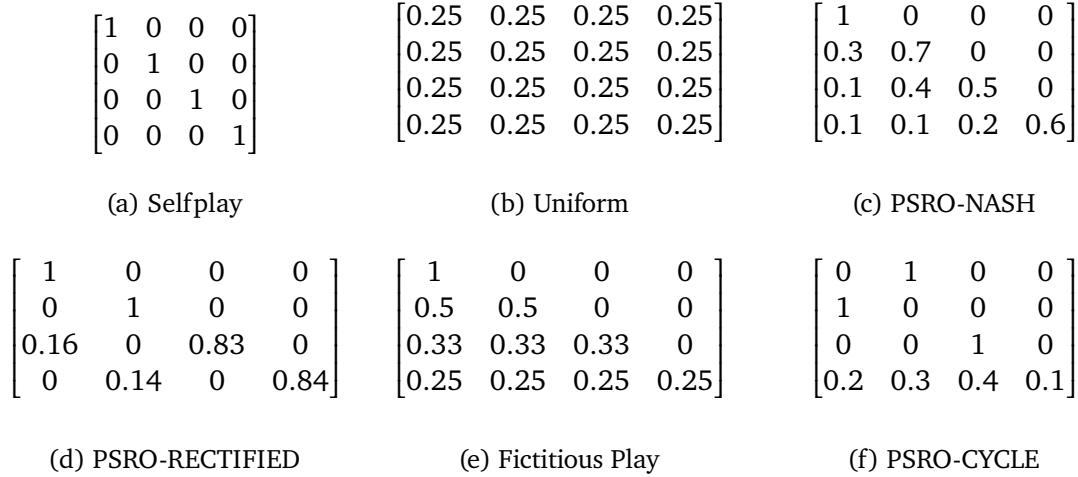


Figure 9 | Matchmaking Algorithms: (a) Selfplay: Each player only plays against itself. (b) Uniform: Each player is sampled uniformly at random. (c) PSRO-NASH: Player j is sampled from a Nash mixture policy over the current population (up to i) that is gradually updated over the course of training. (d) PSRO-RECTIFIED: Similar to PSRO-NASH, except player j is sampled from a “rectified” Nash mixture policy. (e) PSRO-UNIFORM: Player j is sampled from the uniform mixture over the current population of players. (f) PSRO-CYCLE: The first $n - 1$ players form a strategic cycle and the last player responds to Nash equilibrium over the entire league.

In contrast to the typical approach where a policy population is grown by iteratively training new policies against snapshots of existing policies, we train all players in the AZ_{db} league concurrently.

For PSRO variants, we compute separate interaction graphs Σ_b and Σ_w for playing as black and white respectively. During training, we track the outcome of matchups in a payoff matrix and compute Σ_w and Σ_b for the row (white) and column (black) players respectively. We also differ slightly from the standard PSRO setup and include the current player in the population it is responding to. This means player 0 is not treated as a fixed policy and essentially engages in standard AZ self-play. This works well empirically.

Practical Considerations. For PSRO variants where each player is potentially responding to a different mixture policy or matched against specific opponents to discover strategic niches, it is necessary to filter the experience generated from each matchup so that only i ’s experience (the exploiter) is used to update the agent, and not j ’s experience (the exploitee). This is because updating from both players’ experiences will effectively cancel out the diversity effect we are seeking when players are trained concurrently.

Information Hiding. For standard AZ self-play, MCTS uses the same agent to play both sides of the game during search. This is no longer true for AZ_{db} as each player in the league is conditioned on a latent variable l , and care needs to be taken to restrict the searching player from using privileged information about their opponent’s identity. Therefore during training, we ensure that each player can only plan with their own policy and values.




Figure 10 | Empirical PSRO-NASH Interaction Graphs for a league with 5 players. Each grid represents a matchup between exploiter player i (row) and exploitee player j (column), lighter color implies higher probability (white is probability 1 and black is 0). Since we only match an opponent j with player i for $j \leq i$, the interaction graph is lower diagonal (but includes the diagonal and therefore allows self-play). We can see that most players get to play self-play games (have probability mass on the diagonal) but also get to play against other players.

B.2. Matchmaker comparisons

We now compare the six matchmaking mechanisms proposed above in four experiments (Fig. 11-14). In each experiment, we trained six agents using the fast configuration. These experiments differ in two aspects: we trained the agents with and without the diversity intrinsic reward (when there is no intrinsic reward, there is only response diversity), and with and without "information hiding", *i.e.*, hiding the privileged information about the opponent during planning. The columns correspond to different matchmakers, grouped in triplets, where for each matchmaker we present the results for latent 0, max-over-latents and sub additive planning (left-to-right).

The rows correspond to the data sets in the main paper, where we additionally include a version of STS with a 0/1 loss (the agent gets a score of 1 if it finds any of the propose solutions, regardless of their internal scores). In the Penrose set, we include two thresholds: 0.25 as in the main paper and 0.1.

Inspecting the results, we could not identify a significant advantage to any of the match making mechanisms. The different matchmakers perform roughly the same, in particular, on the larger and more statistically significant data sets (STS, Lichess). Furthermore, different matchmakers were better than others in different configurations. We could not make an informed decision regarding which matchmaker to use given these results. Nevertheless, we believe that presenting these results here may be of benefit for other researchers interested in match making ideas, in particular in chess and other transitive games.

Diversifying AI:
Towards Creative Chess with AlphaZero

STS	95.13	96.52	95.29	96.14	97.06	96.15	96.01	97.18	96.13	95.71	96.81	95.94	95.43	96.46	95.73	95.3	96.68	95.66
STS 0/1	99.83	99.83	99.75	99.92	99.92	99.92	99.92	100.0	99.83	100.0	100.0	99.83	99.92	100.0	99.92	99.92	100.0	100.0
Challenge	26.67	33.33	33.33	20.0	26.67	20.0	20.0	20.0	13.33	26.67	40.0	33.33	26.67	26.67	26.67	6.67	13.33	6.67
Penrose_1	5.66	5.66	5.66	9.43	9.43	9.43	9.43	9.43	9.43	9.43	9.43	9.43	11.32	11.32	11.32	7.55	7.55	7.55
Penrose_25	18.87	18.87	18.87	16.98	16.98	16.98	11.32	11.32	11.32	20.75	20.75	20.75	16.98	16.98	16.98	16.98	16.98	16.98
Lichess	86.61	88.04	87.14	87.68	88.39	88.04	87.68	89.11	89.11	86.96	89.29	88.75	88.21	88.75	88.75	87.14	88.04	88.04
	AZ _{db} _PSRO_Recified (latent=0)	AZ _{db} _PSRO_Recified (max-latents)	AZ _{db} _PSRO_Recified (latents-SA)	AZ _{db} _Fictitious_Play (latent=0)	AZ _{db} _Fictitious_Play (max-latents)	AZ _{db} _Fictitious_Play (latents-SA)	AZ _{db} _PSRO_CYCLE (latent=0)	AZ _{db} _PSRO_CYCLE (max-latents)	AZ _{db} _PSRO_CYCLE (latents-SA)	AZ _{db} _Self_Play (latent=0)	AZ _{db} _Self_Play (max-latents)	AZ _{db} _Self_Play (latents-SA)	AZ _{db} _PSRO (latent=0)	AZ _{db} _PSRO (max-latents)	AZ _{db} _PSRO (latents-SA)	AZ _{db} _Uniform (latent=0)	AZ _{db} _Uniform (max-latents)	AZ _{db} _Uniform (latents-SA)

Figure 11 | No diversity, without privileged information.

STS	95.24	98.13	96.03	95.27	97.74	95.79	95.91	98.17	95.95	95.11	98.06	95.96	95.64	98.04	95.91	95.78	98.03	95.95
STS 0/1	99.83	100.0	99.83	99.75	100.0	99.92	99.92	100.0	99.92	99.75	100.0	99.83	100.0	100.0	99.83	99.83	99.92	99.92
Challenge	20.0	33.33	20.0	6.67	26.67	13.33	13.33	33.33	26.67	13.33	26.67	13.33	0.0	20.0	0.0	13.33	33.33	20.0
Penrose_1	5.66	9.43	5.66	5.66	7.55	5.66	9.43	13.21	13.21	5.66	5.66	5.66	7.55	9.43	7.55	5.66	5.66	5.66
Penrose_25	9.43	9.43	9.43	9.43	11.32	11.32	26.42	35.85	32.08	11.32	13.21	11.32	13.21	20.75	13.21	9.43	11.32	9.43
Lichess	86.25	91.61	88.93	86.25	88.21	87.32	85.0	89.11	87.14	87.14	89.29	87.86	87.68	90.89	89.29	86.25	89.82	89.29
	AZ _{db} _PSRO_Recified (latent=0)	AZ _{db} _PSRO_Recified (max-latents)	AZ _{db} _PSRO_Recified (latents-SA)	AZ _{db} _Fictitious_Play (latent=0)	AZ _{db} _Fictitious_Play (max-latents)	AZ _{db} _Fictitious_Play (latents-SA)	AZ _{db} _PSRO_CYCLE (latent=0)	AZ _{db} _PSRO_CYCLE (max-latents)	AZ _{db} _PSRO_CYCLE (latents-SA)	AZ _{db} _Self_Play (latent=0)	AZ _{db} _Self_Play (max-latents)	AZ _{db} _Self_Play (latents-SA)	AZ _{db} _PSRO (latent=0)	AZ _{db} _PSRO (max-latents)	AZ _{db} _PSRO (latents-SA)	AZ _{db} _Uniform (latent=0)	AZ _{db} _Uniform (max-latents)	AZ _{db} _Uniform (latents-SA)

Figure 12 | With diversity, without privileged information.

Diversifying AI:
Towards Creative Chess with AlphaZero

STS	95.6	97.0	95.92	95.57	96.54	95.87	95.3	96.68	95.88	95.71	96.81	95.94	95.34	96.76	95.7	95.81	97.07	95.98
STS 0/1	99.83	100.0	99.92	100.0	100.0	100.0	99.92	100.0	99.92	100.0	100.0	99.83	100.0	100.0	100.0	100.0	100.0	100.0
Challenge	5.88	11.76	11.76	17.65	17.65	11.76	23.53	29.41	29.41	23.53	35.29	29.41	5.88	11.76	5.88	11.76	11.76	11.76
Penrose_1	7.55	7.55	7.55	9.43	9.43	9.43	13.21	13.21	13.21	9.43	9.43	9.43	5.66	5.66	5.66	9.43	9.43	9.43
Penrose_25	11.32	11.32	11.32	20.75	22.64	20.75	28.3	28.3	28.3	20.75	20.75	20.75	18.87	18.87	18.87	16.98	16.98	16.98
Lichess	87.14	88.57	88.39	88.21	89.29	88.75	86.25	87.68	86.79	86.96	89.29	88.75	86.07	87.68	86.96	88.39	89.82	88.93
	AZ _{db} _PSRO_Recified (latent=0)	AZ _{db} _PSRO_Recified (max-latents)	AZ _{db} _PSRO_Recified (latents-SA)	AZ _{db} _Fictitious_Play (latent=0)	AZ _{db} _Fictitious_Play (max-latents)	AZ _{db} _Fictitious_Play (latents-SA)	AZ _{db} _PSRO_CYCLE (latent=0)	AZ _{db} _PSRO_CYCLE (max-latents)	AZ _{db} _PSRO_CYCLE (latents-SA)	AZ _{db} _Self_Play (latent=0)	AZ _{db} _Self_Play (max-latents)	AZ _{db} _Self_Play (latents-SA)	AZ _{db} _PSRO (latent=0)	AZ _{db} _PSRO (max-latents)	AZ _{db} _PSRO (latents-SA)	AZ _{db} _Uniform (latent=0)	AZ _{db} _Uniform (max-latents)	AZ _{db} _Uniform (latents-SA)

Figure 13 | No diversity, with privileged information.

STS	94.85	97.82	95.48	94.67	98.04	95.14	95.34	98.08	95.29	95.71	96.81	95.94	95.41	97.81	95.71	95.72	97.74	96.07
STS 0/1	99.83	100.0	99.92	99.75	100.0	99.83	99.66	100.0	99.75	100.0	100.0	99.83	99.92	100.0	99.92	100.0	100.0	100.0
Challenge	17.65	17.65	11.76	0.0	0.0	0.0	5.88	17.65	11.76	23.53	35.29	29.41	5.88	23.53	11.76	11.76	11.76	11.76
Penrose_1	7.55	9.43	7.55	7.55	7.55	7.55	9.43	9.43	9.43	9.43	9.43	9.43	3.77	5.66	3.77	5.66	7.55	5.66
Penrose_25	16.98	18.87	16.98	9.43	9.43	9.43	11.32	15.09	15.09	20.75	20.75	20.75	18.87	20.75	20.75	9.43	9.43	9.43
Lichess	85.36	88.57	86.61	83.57	87.68	85.89	86.07	88.93	87.32	86.96	89.29	88.75	87.86	90.0	88.93	85.54	89.11	87.86
	AZ _{db} _PSRO_Recified (latent=0)	AZ _{db} _PSRO_Recified (max-latents)	AZ _{db} _PSRO_Recified (latents-SA)	AZ _{db} _Fictitious_Play (latent=0)	AZ _{db} _Fictitious_Play (max-latents)	AZ _{db} _Fictitious_Play (latents-SA)	AZ _{db} _PSRO_CYCLE (latent=0)	AZ _{db} _PSRO_CYCLE (max-latents)	AZ _{db} _PSRO_CYCLE (latents-SA)	AZ _{db} _Self_Play (latent=0)	AZ _{db} _Self_Play (max-latents)	AZ _{db} _Self_Play (latents-SA)	AZ _{db} _PSRO (latent=0)	AZ _{db} _PSRO (max-latents)	AZ _{db} _PSRO (latents-SA)	AZ _{db} _Uniform (latent=0)	AZ _{db} _Uniform (max-latents)	AZ _{db} _Uniform (latents-SA)

Figure 14 | With diversity, with privileged information.

C. Diversity bonuses in AZ

It might surprise some readers that RL can benefit from creative-problem solving mechanisms. In principle, RL can solve any problem via trial and error. However, RL methods cannot find truly optimal policies in most problems of practical interest and hence find a particular (hopefully) near-optimal solution whose choice is determined by the nature of the agent’s bounds (memory, compute, training experience), as depicted in Fig. 15 (left). This singular sub-optimal choice can potentially generalise poorly when faced with situations it did not encounter during its training. In other words, a single policy will face Out of Distribution (OOD) situations w.r.t to its self-play distribution (the games that the agent plays against itself and gets to train on). Incorporating multiple policies with different self-play distributions (diverse d_{π}) extends the self-play distribution of AZ_{db} as a team, and we can therefore expect that some players will generalize better than others in some situations.

Another potential issue is that RL is designed to find an optimal policy from an initial state distribution. This is fine if one only cares about playing chess from the initial board, but people are also interested in knowing what the optimal move is from a variety of positions, potentially outside of the path of the optimal policy. Even with advanced exploration techniques, this fundamental problem remains: exploration is designed to explore the state space until the agent is confident enough about the optimal policy, and at this point, it becomes less likely for it to encounter sub-optimal states. Another issue is that the policy is often deterministic, and as such, will play the same move from each position, ignoring other potentially optimal moves. Modern chess engines, however, are required to do much more than play perfect chess from the initial board: they are being asked to recommend moves from sub-optimal positions, analyze games, and more.

In Fig. 15 (right), we depict one hypothetical optimal trajectory (assuming access to an optimal policy) starting from the initial board (state 0) and continuing to states 1, 4 and 6 with red arrows (self-play). Having a diverse team of agents can help in alleviating this issue and improving our data distribution in the following manner. In green color, we depict a different optimal trajectory in which the opponent selects a different optimal action in state 1 leading the game to states 5 & 7. The agent might miss such a trajectory when training with a self-play league and using other matchmakers can help to explore it. Alternatively, an agent may choose to play a different optimal action in state 0 leading it to visit 2, 8 & 9 by optimizing quality diversity (purple).

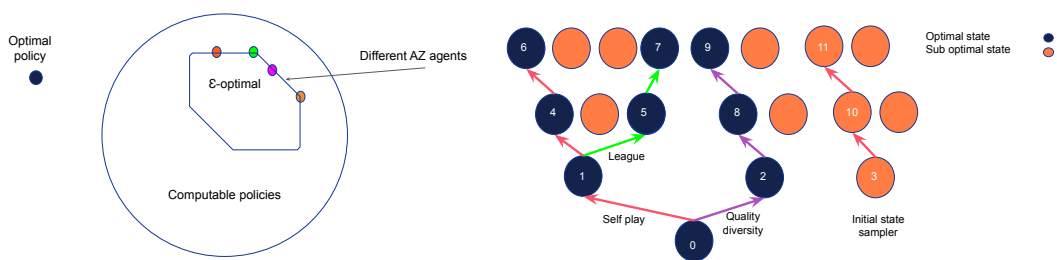


Figure 15 | Left: an illustration of the space of chess agents. The optimal policy is not computable; within the computable policies there are different ϵ -optimal policies. Right: An illustration of a state space. Blue colors represent optimal states, *i.e.*, states that can be reached with an optimal policy. Red arrows (starting from state 0) represent one possible optimal trajectory, whereas green and purple arrows represent alternative optimal trajectories. State 3 is a sub-optimal state that would not be visited by an optimal policy.

C.1. Evaluation of external chess engines in puzzle

The goal of this work is to study diversity bonuses in AZ. Our focus was not on comparing or making any claims about the strength of a particular chess engine compared to others. Nevertheless we believe it is informative to present the solve rate of reference chess engines, as it provides further evidence of the challenges in solving these puzzles. We evaluated Stockfish 8 & 15, and Crystal 4.1 with the following configurations: Level = 20, Number of threads = 16, Search time / move = 20 mins, Hash table size = 1024 MB.

We would like to emphasize that our methods are not specific to chess and can be applied to other domains.

STS	97.06	98.08	98.18	97.2	96.86	99.55	97.37
Challenge	46.67	53.33	40.0	33.33	20.0	53.33	20.0
Penrose	32.08	56.6	58.49	30.19	50.94	56.6	50.94
Lichess	97.5	97.68	97.32	94.64	95.0	98.04	95.54
	SF ₈	SF ₁₅	Crystal4.1	AZ	AZ _{db-latent}	AZ _{db-MAX}	AZ _{db-SA}

Figure 16 | The solve rate of external chess engines across different data sets.

C.2. Visualizing the piece occupancies

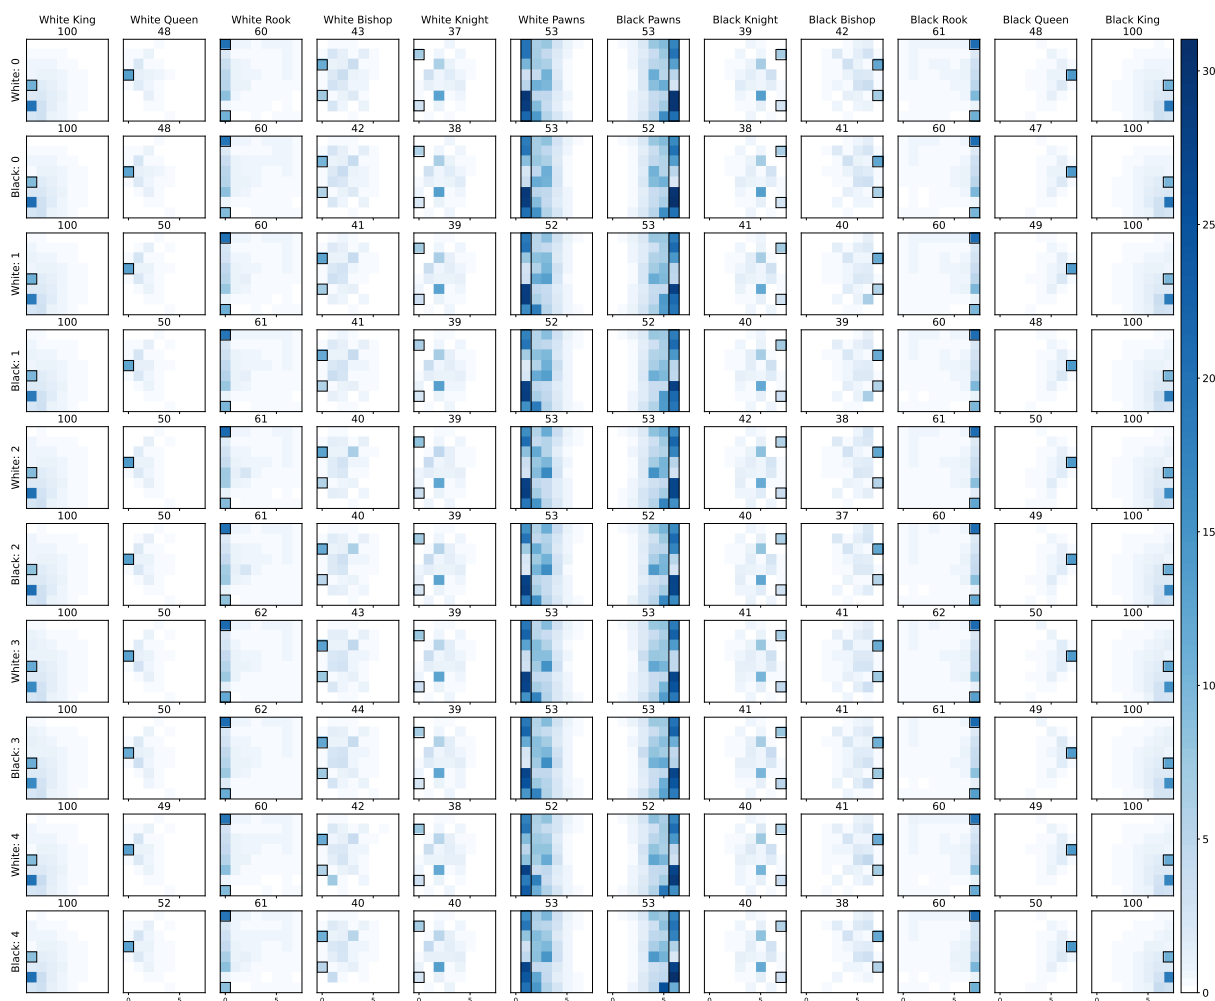


Figure 17 | Piece occupancies of AZ_{db}. Columns correspond to pieces: leftmost five as white and rightmost as black. Rows correspond to players playing as white (odd rows) and black (even rows). Numbers above sub figures indicate the sum of the occupancy over the board divided by the number of pieces (eg 1 for king, 8 for pawns) and indicate how long does a piece stays on the board on average. Black rectangles highlight the starting position of each piece. Color units (in the colorbars) correspond to occupancy in % multiplied by the average number of moves in a game (65). For orientation, the chess coordinate system is displayed in Figure 3 in the main paper, top-left.

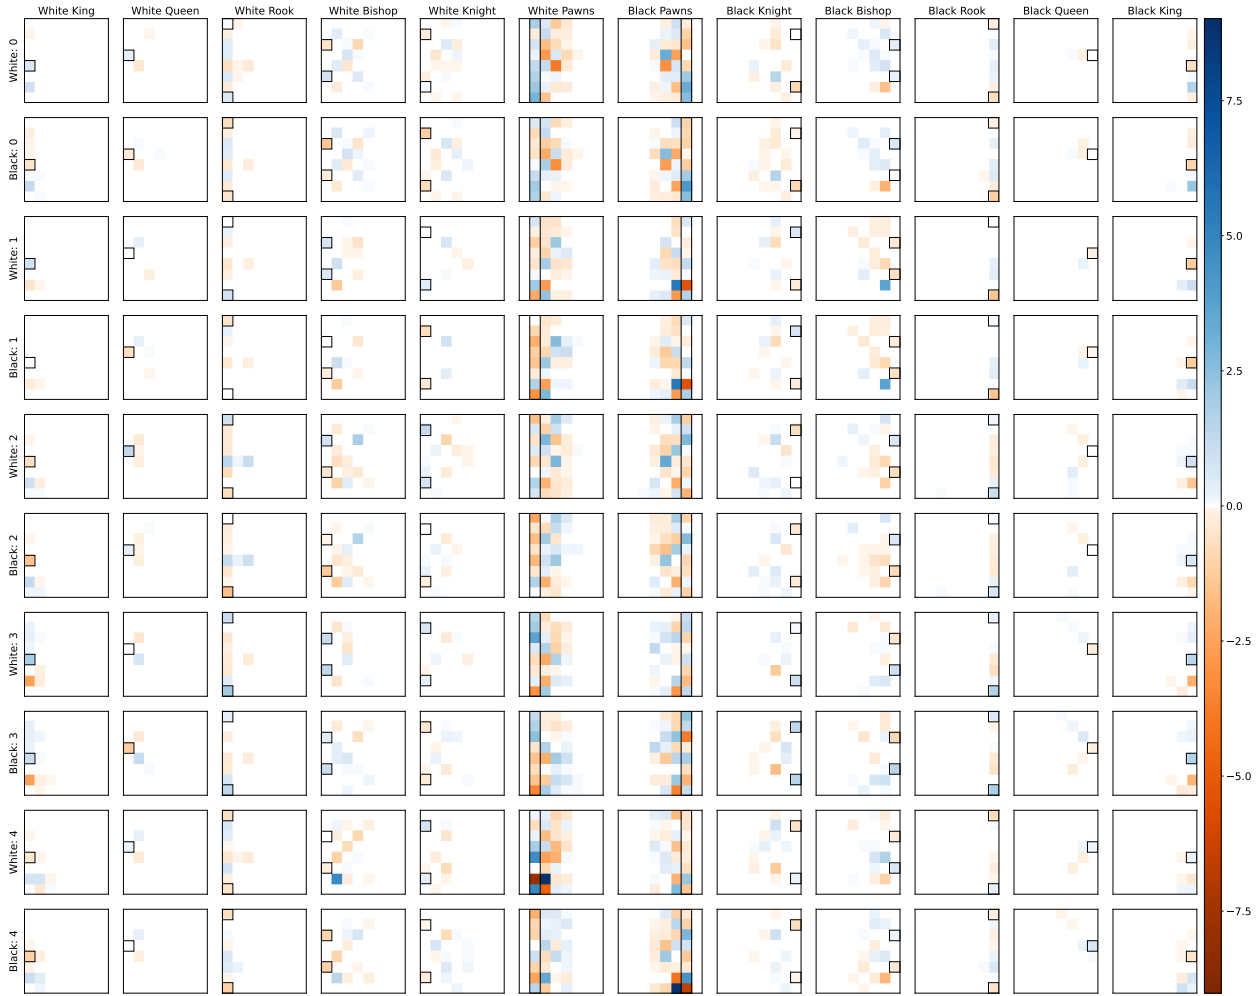


Figure 18 | Mean-subtracted piece occupancies of AZ_{db} : $\tilde{\psi}^i = \psi^i - \frac{1}{N_{Players}} \sum \psi^i$. Columns correspond to pieces: leftmost five as white and rightmost as black. Rows correspond to players playing as white (odd rows) and black (even rows). Blue color implies that a specific player has higher occupancy than the average, and red color implies lower occupancy. Black rectangles highlight the starting position of each piece. Color units (in the colorbars) correspond to occupancy in % multiplied by the average number of moves in a game (65). For orientation, the chess coordinate system is displayed in Figure 3 in the main paper, top-left.

C.3. Details on self-play from puzzle positions.

Table 1 in the main paper presents the performance of AZ on the challenging puzzles sets (Challenge and Penrose) when we allow AZ to begin self-play games from puzzle positions. Here we further explain the details for each one of the variations in the table.

Self-play from puzzle positions: This agent uses a *start-position sampler* that samples an initial board position from the following distribution: with $\mathbb{P} = 0.5$ it uses the standard initial chess board and with $\mathbb{P} = 0.5$ it uniformly samples a position from the Challenge and Penrose Sets. Note that we do not give the solutions to the puzzles to the agent in any form. AZ is allowed to train on self-play games that it plays from these positions.

+ *intermediate positions:* While the previous agent used the first position from each puzzle in the *start-position sampler*, this agent also samples intermediate positions from each puzzle in the Challenge set. Since those puzzles have a unique multi-step solution, adding intermediate positions allows the agent to play on more advanced steps of a puzzle and can help the agent with exploration.

+ *exploration:* This agent additionally uses an exploration mechanism to sample actions during training. The agent samples actions from a softmax distribution ($\tau = 10$; identical setup for exploration is used in AZ (Silver et al., 2018b)) constructed over the MCTS visit-counts $N^0(s, a)$ for the first 15 steps of an episode. After the first 15 steps, the agent continues to sample actions greedily from its visit-counts.

+ *half-move clock:* The half-move clock in chess counts how many moves have been played since the last move where a pawn advanced or a piece was captured. When this counter reaches 100 the game ends in a draw, which is known as the 50-move draw rule. The closer this counter is to 100, the more likely it is that a position will end up in a draw. Whenever we introduce positions in the start position matchmaker, we introduce them with the *half-move clock* set to 0. However, in this experiment, we also included puzzle positions with half-move clocks set from 0 to 80, in increments of 20. This data augmentation should make it easier for the agent to learn that some positions are leading to a draw.

Train/test splits. In the middle panel of Table 1, we conducted two more experiments to understand the generalization ability of AZ. The first experiment focused on in distribution generalization, *i.e.*, AZ used a start-position sampler that consisted of 36 positions from the Penrose set and was evaluated on the remaining 17 held-out positions. This split implies that AZ had access to some of the variations of each Penrose position in its training set and the others in its test set (see Tables 4-11). Inspecting the results, we can see that the agent achieves 88.24% (see (3) in Table 1 in the main paper) suggesting that it generalizes well in distribution (across variations of the same position). The second experiment focused on OOD generalization, where AZ used all the puzzles from the Penrose set except the original (position 2) Penrose puzzle position and its variations (training set = 36 puzzles; test set = 17 puzzles). This agent achieved a performance of 11.76% on the test set (see (4) in Table 1 in the main paper).

C.4. Additional results on AZ’s understanding of puzzles

In the main paper, we inspected AZ’s understanding of the Penrose positions. We now complement this analysis with a position from the Challenge set in Fig. 19.

The first and second rows in Fig. 19 shows the first and third steps respectively of the Challenge set puzzle. Inspecting the first step we can see that AZ has a higher probability for selecting the correct action Nxa2 and also an accurate raw value estimate for this move (draw). However, the raw value estimates for the other actions (which are not part of the solution) suggest that the agent can win if it played any of them, and after MCTS, AZ predicts that all actions are equally likely to lose the game.

In the third step of this puzzle (second row), it is white’s turn to play, black has the possibility of promoting its a2 pawn in the next move, and thus, it is important to pick the correct move here. We indicate some of the legal moves with arrows, where red arrows indicate wrong moves that lead to a loss and the blue arrow shows the unique, correct move (Kc1) that leads to a draw (if black promotes the a2 pawn to a queen, white responds with Bishop b1 that locks the position).

AZ again selects the wrong move, it has low prior probability for all the legal actions available from this position, and more importantly, its raw value estimates predict that it is a win for all of those legal moves! When AZ trained on puzzles with self-play, it correctly predicts the action (in terms of prior probability) and also understands the position correctly (measured in terms of raw and MCTS values).

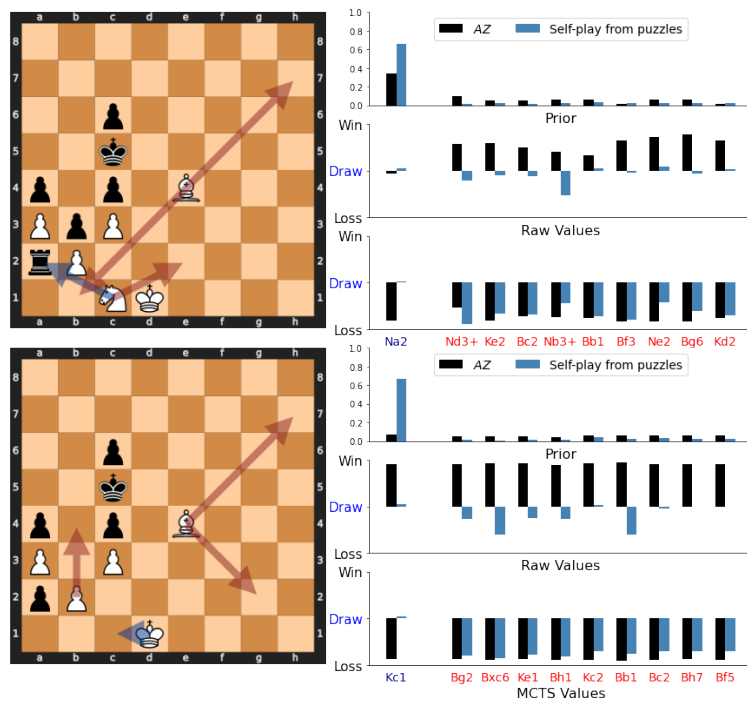


Figure 19 | *Left*: Two puzzle positions from the Challenge set. *Right*: Visualizations of prior probability from a subset of legal actions, and its corresponding raw value and MCTS value estimates for AZ and AZ trained on puzzle sets. Correct and Incorrect moves are colored in Blue and Red respectively.

C.5. Sub-additive planning in chess matchups

In this ablation experiment, we looked at the performance of sub-additivity as a function of the number of seeds in matchups. We looked at the averaged scores from matchups across opening positions generated using a certain number of random seeds and, from those matchups, defined a look-up table that mapped from an opening position and AZ_{db} 's player color to the player from the team with the highest score against AZ. This look-up table was used to select a player on hold-out matchups from a distinct set of 30 seeds. This answers the question of whether such a sub-additive approach can generalize to a hold-out set of matchups and more broadly answers the question of whether different players specialize to specific opening positions. The results from this study are presented in Fig. 20, where the average win rate (scores in $[-1, 1]$ are transformed to $[0, 1]$) against AZ as a function of *training seeds*, i.e., the number of matchups that were used to define the look-up table for sub-additive player selection. The figure also shows in dashed horizontal lines the average win rate of the best latent (black) and *max-over-latents* (red). From the plot, we see that when players are selected sub-additively using the look-up table (constructed from 90 seeds), the win rate of AZ_{db} improves to 0.598 on the hold-out set of matchups (from 30 random seeds). This finding is similar to the one that was previously described in Figure 1 in the main paper (Right): sub-additive planning with players in AZ_{db} improves performance also in matches, while the previous experiments demonstrated this only in puzzles. The figure also shows that sub-additive planning in matchups improves as a function of training seeds (that are used to construct the look-up table).

For each opening and seed, we calculated the average score of each player across all other seeds (leaving one seed out) and defined a mapping from an opening to a player with the highest score. We then evaluated this mapping on the held-out seed of matchup. This process is repeated for each seed that is held-out for evaluation and the scores obtained from the evaluation are averaged and reported as the sub-additive planning scores in Fig. 21.




Figure 20 | Sub-additive planning in chess matchups against AZ. Reports the average winrate of AZ_{db} against AZ, generated from 500 random seeds. Also shows the winrate of the best latent (in Black) and *max-over-latents* (in Red) across those matchups.

=0	49.2	2.6	22.2	28.7	2.0	43.2	54.25	29.6
=1	48.6	2.6	22.8	28.3	2.1	43.6	53.51	24.4
=2	50.4	2.6	21.0	27.5	2.2	44.4	54.22	29.4
=3	49.4	2.4	22.2	27.0	2.0	45.0	53.08	21.4
=4	50.0	2.6	21.4	26.9	2.0	45.0	53.53	24.6
=5	49.2	2.7	22.1	28.0	2.1	44.0	53.76	26.2
=6	49.9	2.8	21.4	27.8	2.2	44.1	54.11	28.6
=7	50.0	2.6	21.4	27.3	2.1	44.5	53.88	27.0
=8	49.9	2.6	21.5	27.3	2.1	44.6	53.76	26.1
=9	48.4	2.8	22.8	28.8	2.2	43.0	53.84	26.8
SA	52.2	2.7	19.1	30.0	2.2	41.8	57.18	50.3
max	52.9	2.6	18.5	31.4	2.2	40.4	58.6	60.3
	# wins (W)	# draws (W)	# loses (W)	# wins (B)	# draws (B)	# loses (B)	winrate %	Elo

Figure 21 | Summary of chess matchups between AZ_{db} and AZ. The first ten rows reports the average of the match outcomes (i.e., wins, draws and loses) for each player in the team. The last two rows reports the same for sub-additive planning and *max-over-latents*. In addition to game outcomes, we also report the average winrate and relative Elo improvement over AZ.

Diversifying AI:
Towards Creative Chess with AlphaZero

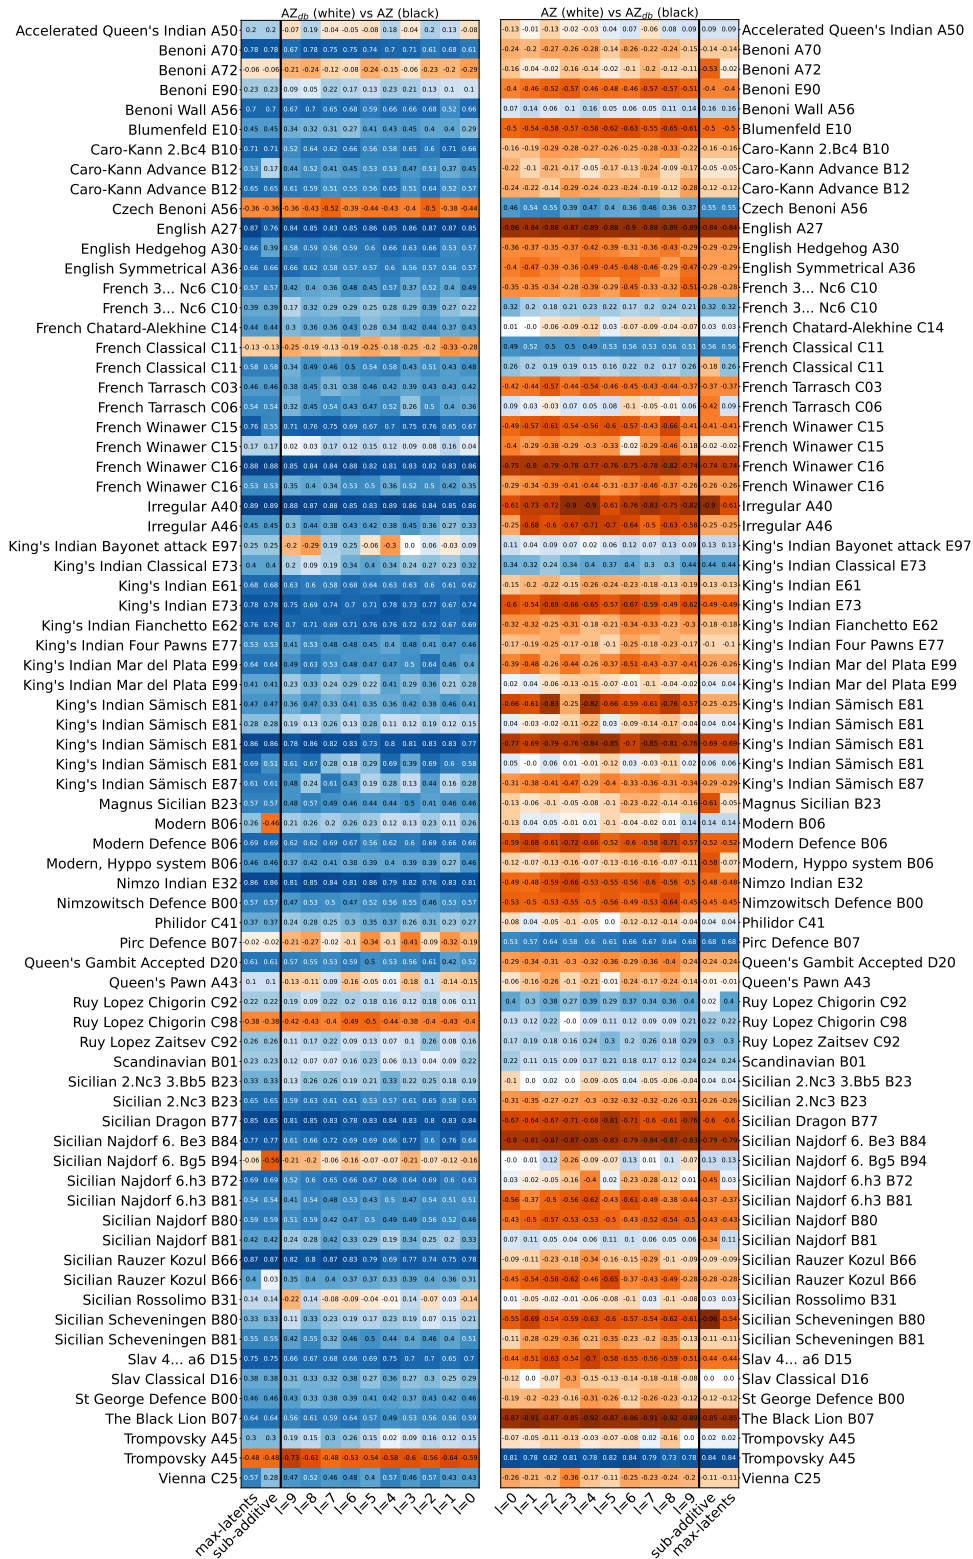


Figure 22 | Results from different chess openings between AZ_{db} and AZ (scores averaged from 30 random seeds). In the figure on the left, AZ_{db} played as White and AZ as Black; the colors were flipped for the figure on right.

Table 3 | Average winrates of each player in AZ_{db} when playing as White and as Black against AZ.

AZ _{db} team	as White	as Black
Player 0	68.27%	40.21%
Player 1	67.39%	39.63%
Player 2	69.86%	38.57%
Player 3	68.31%	37.83%
Player 4	69.26%	37.79%
Player 5	68.32%	38.19%
Player 6	69.23%	38.98%
Player 7	69.36%	38.38%
Player 8	69.24%	38.26%
Player 9	67.30%	40.38%

D. Puzzle sets

The *Lichess* data set is composed of 560 puzzles that we selected from a larger data set provided by lichess.org. The complete data set includes 3M puzzles collected from 300M analysed games from the Lichess database and re-analyzed with SF. We selected 560 puzzles from them by creating a subset of the 20K hardest puzzles (rated by "ELO" score) and then further selecting from them the positions that were unsolved by SF15 with a low number of simulations (10k). Each puzzle in the set requires multiple steps; the agent receives a score of 1 if it predicts the correct unique action in all the steps and 0 otherwise.

The *Strategic Test Suite (STS)* is a data set of 1088 single-step, multi-choice chess puzzles, designed to evaluate chess engine's long term understanding of strategical and positional concepts. The suite was created by Dann Corbit and Swaminathan Natarajan, but we adapt the more recent [STS-Rating \(v6\)](#) by Ferdinand Mosca. Actions receive scores $\in [0, 1000]$, we normalize the score for each action by the score of the best action (1000), so the score in each puzzle is $\in [0, 1]$.

The next two data sets were collected and created by our team. They include puzzles from different sources and we refer the reader to the tables in the end of the Appendix for more details.

The *Challenge set* includes 15 multi-choice & step puzzles that are designed to be challenging to modern chess engines. A score of one is given if the agent predicts the correct unique action in each step of the puzzle and 0 otherwise. We deal with a few cases of ambiguous solutions via a simple concept of "or lists", if the agent finds one of the possible unique multi-step solutions to the puzzle it gets a score of 1 in this puzzle. Sometimes the opponent has different ways to respond to the agent's moves, in this case, we apply an "and list" and require the agent to find the unique multi-step solution on all the variations.

The *Penrose set* includes 53 challenging single-step puzzles, where the goal is to predict the value of the position (in $[0, 0.5, 1]$). A score of one is given if the value prediction (after a search) is within a threshold (of 0.25 or 0.1) from the correct value and 0 otherwise. The set includes three sources of puzzles: (1) The Penrose position(s): a chess puzzle proposed by Sir Roger Penrose to "*learn more about the uniqueness of the human brain*" ([Doggers, 2017](#)), can be seen in Table 7, top. We also include a second position by Penrose with a similar idea that we found [here](#) (Table 4, top). (2) A set of variations of the two Penrose positions as well as other puzzles created by our team (Tables 4-11). (3) Fortress positions from ([Guid and Bratko, 2012](#)), selected from ([Dvoretsky, 2010](#)) (Tables 19-20) and blocked positions composed by us (Tables 12-18). Together with the Lichess data set, these puzzle sets present a diverse set of challenges for chess engines.

The following tables present positions from the Penrose and Challenge sets respectively. The tables include the starting position of each puzzle, the FEN, the solution, the player to play and comments on each position. The comment provides context on the position and sometimes explain why it may be hard for engines to correctly evaluate. This analysis was performed using a web version of SF, and we emphasize that it is not using SF at its full strength. The goal of these comments is to provide context on why the position is hard to chess programs and we use SF due to its popularity.

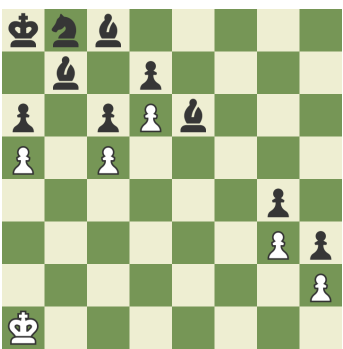
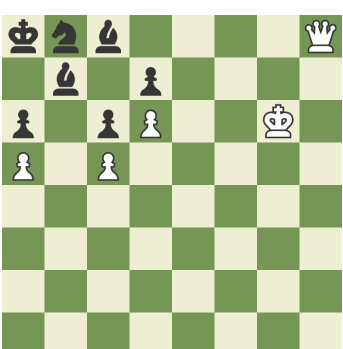
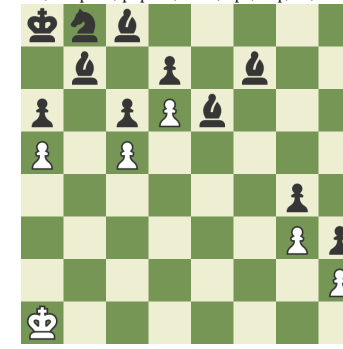
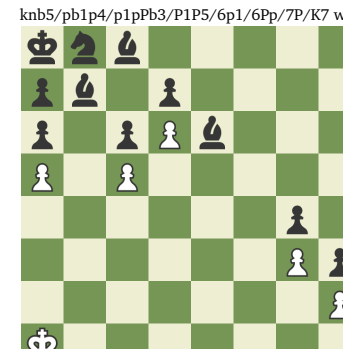
Board	Solution	To Play	Comment
<p>knb5/1b1p4/p1pPb3/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Draw	White	White can draw in several ways, for example, by repeating Kb2-a1 to draw via the three-fold repetition rule or the 50 move rule. There is no way for white or black to make progress. This position would be even be drawn is white picks up the pawns on g4 and h3. The reason is that black can sacrifice the bishop for one of the pawns. The resulting position is a draw because it's a fortress.
<p>knb4Q/1b1p4/p1pP2K1/P1P5/8/8/8/8 b - -</p> 	Draw	Black	A position that arises from a calculation sequence in the baseline puzzle. This position is a fortress, and white cannot make further progress.
<p>knb5/1b1p1b2/p1pPb3/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Draw	White	This puzzle is even more difficult for SF, as there is an additional piece meaning more possible moves. However, for humans, it's clear that the fundamental aspects are the same.
<p>knb5/pb1p4/p1pPb3/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Draw	White	Interestingly enough, this position is easier for Stockfish to evaluate. This is likely because the pieces in the left top of the board have no moves, and SF recognizes that stalemate (no moves left, resulting in a draw) is a likely outcome.

Table 4 | Penrose set positions, Penrose variation 1, positions 1-4


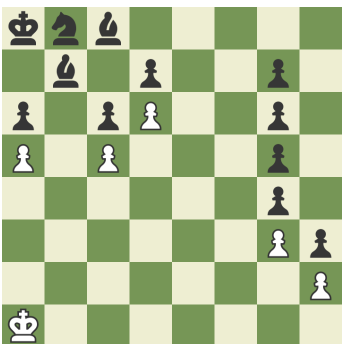
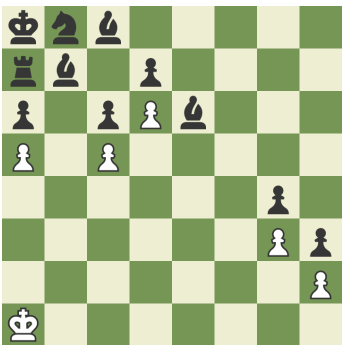
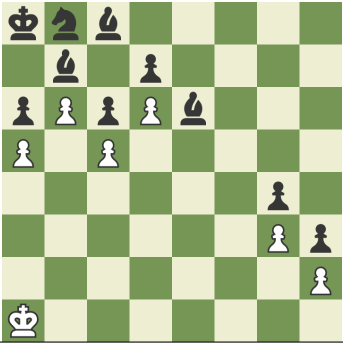
Board	Solution	To Play	Comment
<p>knb5/1b1p4/p1pPp3/P1P2p2/3K2p1/6Pp/7P/8 w - -</p> 	Draw	White	This also seems to be an easier position for SF to understand, likely because it's easier to calculate the relevant variations (contrary to the above example which requires some 'pattern' recognition).
<p>knb5/1b1p2p1/p1pP2p1/P1P3p1/6p1/6Pp/7P/K7 w - -</p> 	Draw	White	Like the position above, this seems to be an easier position for SF to understand. It's a combination of the two reasons above – fewer variations to calculate, and white has no way to infiltrate black's position.
<p>knb5/rb1p4/p1pPb3/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Draw	White	Similar to the position with the added pawn on a7 (three rows up), this position seems to be easier for SF despite the large material increase.
<p>knb5/1b1p4/pPpPb3/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Draw	White	This seems to be an easier position to understand (than the original puzzle), suggesting that perhaps stalemates are likely an easier concept to grasp than fortresses.

Table 5 | Penrose set positions, Penrose variation 1, positions 5-8

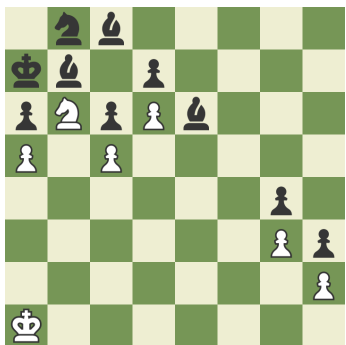
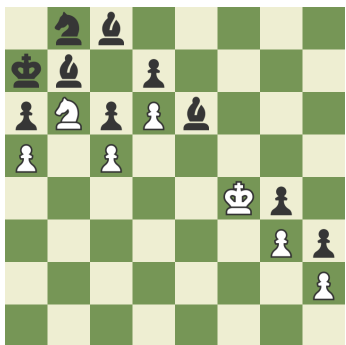
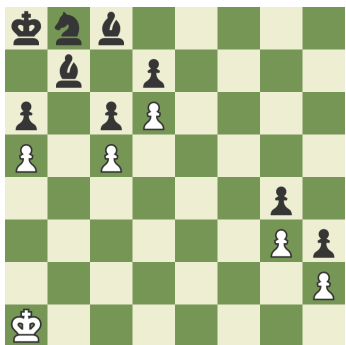
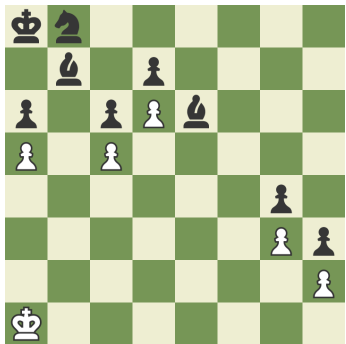
Board	Solution	To Play	Comment
<p>1nb5/kb1p4/pNpPb3/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Win	White	<p>Knight a4 keeps the position locked and allows the white king to fight with the knight against the black bishop on e6 and pick up the pawns.</p>
<p>1nb5/kb1p4/pNpPb3/P1P5/5Kp1/6Pp/7P/8 w - -</p> 	Win	White	<p>Here white has a concrete way to win: 1. Kg5 Bb3 2. Kxg4. After that, white can carefully advance the g-pawn, and promote the g-pawn and/or pick up the h-pawn. The knight on b6 is important for two reasons: [1] it prevents black from having a fortress and [2] allows white to control when black can play Ka7-Ka8. SF misevaluates this as a draw.</p>
<p>knb5/1b1p4/p1pP4/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Win	White	<p>This position seems to be easy for SF to find the correct move. White can win by walking the king over to the right side on the board, picking up both pawns and then using the two queens to mate black (or using the pieces to take black's pieces and the pawns on d7 and c6).</p>
<p>kn6/1b1p4/p1pPb3/P1P5/6p1/6Pp/7P/K7 w - -</p> 	Lose	White	<p>The black king can escape by moving the b7-bishop to a8. From there its an easy win.</p>

Table 6 | Penrose set positions, Penrose variation 1, positions 9-12

Board	Solution	To Play	Comment
<p>8/p7/kpP5/qrp1b3/rpP2b2/pP4b1/P3K3/8 w - -</p>	Draw	White	The Penrose position. As white has control over the white squares, and black has control over the black squares, neither side can make any progress to win the game. As such, this position is a draw.
<p>8/p7/kpP2b2/qrp1b3/rpP2b2/pP4b1/P3K3/8 w - -</p>	Draw	White	Compared to the original penrose position, we have further increased the number of bishops leading to a larger material advantage for black. However, the additional pieces do not help black – black cannot use them to make further progress. As such, the positions is still a draw.
<p>5b2/p5b1/kpP2b1b/qrp1b3/rpP2b2/pP4b1/P3K3/8 w - -</p>	Draw	White	In this position, there are even more bishops compared to the position above. However, the position is still a draw. This position is more difficult to evaluated with calculation alone, as there are more possible move sequences. However, if the engine understands that neither side can make any progress towards winning, then it can correctly evaluate the position.
<p>8/p7/kpP5/qrp1b3/rpP5/pP6/P3K3/8 w - -</p>	Draw	White	Similar to the original position, however black only has one bishop. SF recognizes that this position is a draw, despite the fact that black has large material advantage.

Table 7 | Penrose set positions, Penrose variation 2, positions 1-4

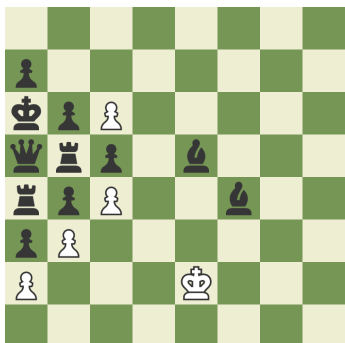
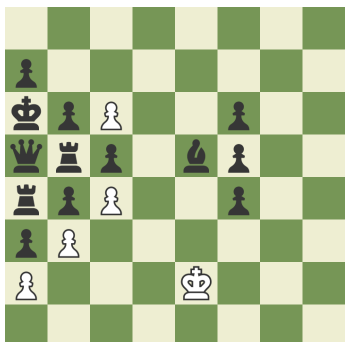
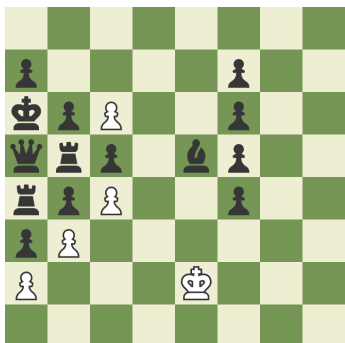
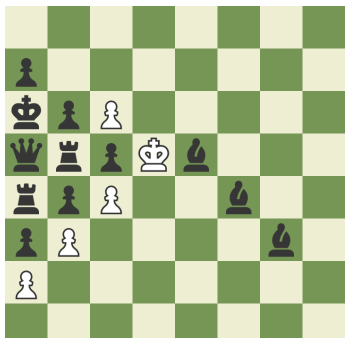
Board	Solution	To Play	Comment
<p>8/p7/kpP5/qrp1b3/rpP2b2/pP6/P3K3/8 w - -</p> 	Draw	White	Similar to the original position, however black has two bishops. However, SF gets this wrong.
<p>8/p7/kpP2p2/qrp1bp2/rpP2p2/pP6/P3K3/8 w - -</p> 	Draw	White	SF understands that this is a draw. This is surprising given that the evaluation of the previous positions. Black's material advantage is the same as the position above.
<p>8/p4p2/kpP2p2/qrp1bp2/rpP2p2/pP6/P3K3/8 w - -</p> 	Draw	White	SF does not understand that this is a draw, suggesting that there is a material threshold beyond which point engines may become confused.
<p>8/p7/kpP5/qrpKb3/rpP2b2/pP4b1/P7/8 w - -</p> 	Draw	White	The king's position in the center of the board makes SF evaluate the position even worse for white (compared to the original Penrose position). However, the king's position does not affect the result.

Table 8 | Penrose set positions, Penrose variation 2, positions 5-8

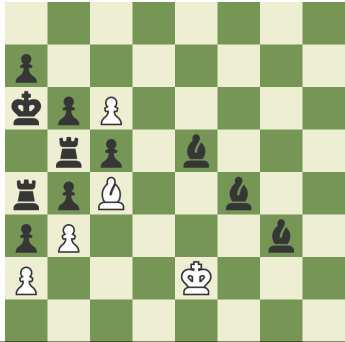
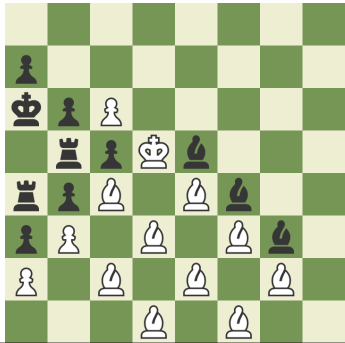
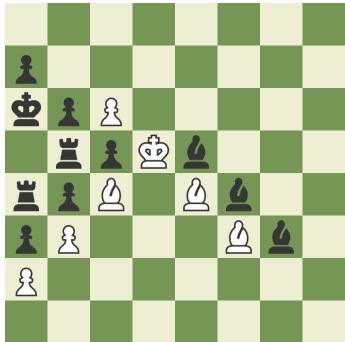
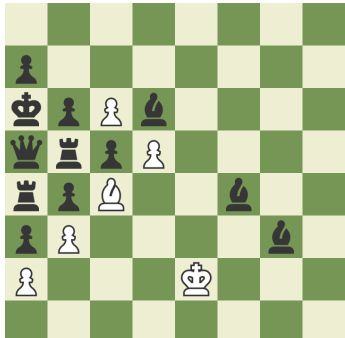
Board	Solution	To Play	Comment
<p>8/p7/kpP5/1rp1b3/rpB2b2/pP4b1/P3K3/8 w - -</p> 	Draw	White	Black has more space on the king side, compared to the original Penrose position, however it does not help Black. This position still evaluates to a draw.
<p>8/p7/kpP5/1rpKb3/rpB1Bb2/pP1B1Bb1/P1B1B1B1/3B1B2 w - -</p> 	Draw	White	In this position, we change the Penrose position to add more material for white, such that white has a material advantage. However, it is still a draw – white only controls the white squares. As black has control over the black squares, white cannot progress. SF thinks that white is winning.
<p>8/p7/kpP5/1rpKb3/rpB1Bb2/pP3Bb1/P7/8 w - -</p> 	Draw	White	White and black get extra bishops compared to the original Penrose position, however the bishops are on opposite colors and the position is still a draw. SF correctly evaluates this position as a draw.
<p>8/p7/kpPb4/qrpP4/rpB2b2/pP4b1/P3K3/8 w - -</p> 	Draw	White	This position is a draw, however black needs to be accurate in their play. SF does not see the danger in the position, and plays incorrectly (resulting in a draw). We elaborate on this in the two positions below.

Table 9 | Penrose set positions, Penrose variation 2, positions 9-12

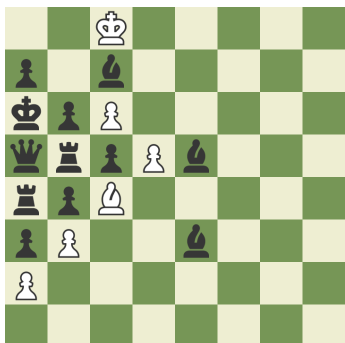
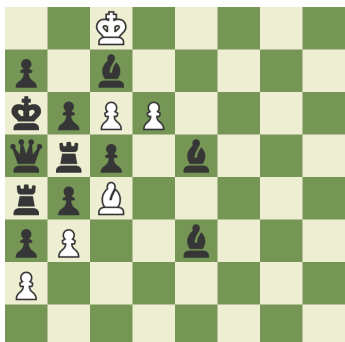
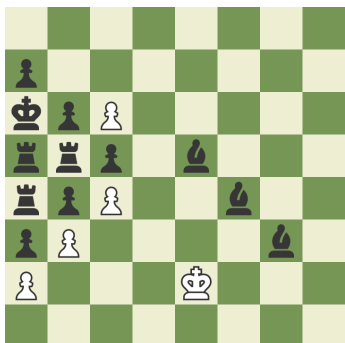
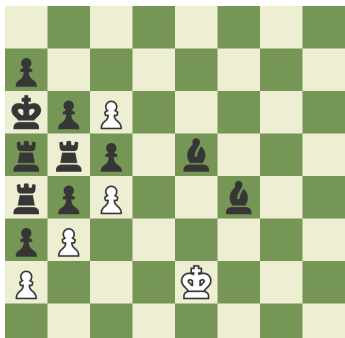
Board	Solution	To Play	Comment
<p>2K5/p1b5/kpP5/qrpPb3/rpB5/pp2b3/P7/8 w - -</p> 	Win	White	Here, SF still thinks black is winning. However, white has mate in 5 with 1. d6.
<p>2K5/p1b5/kpPP4/qrp1b3/rpB5/pP2b3/P7/8 b - -</p> 	Lose	Black	Here, white is winning and SF's evaluation jumps to winning for white. White can now use the e6 square for its bishop to mate black.
<p>8/p7/kpP5/rrp1b3/rpP2b2/pP4b1/P3K3/8 w - -</p> 	Draw	White	SF thinks black is winning. However having a rook, rather than a queen as in the original Penrose position, does not affect the evaluation.
<p>8/p7/kpP5/rrp1b3/rpP2b2/pP6/P3K3/8 w - -</p> 	Draw	White	Similar to the position above, SF thinks black is winning.

Table 10 | Penrose set positions, Penrose variation 2, positions 13-16

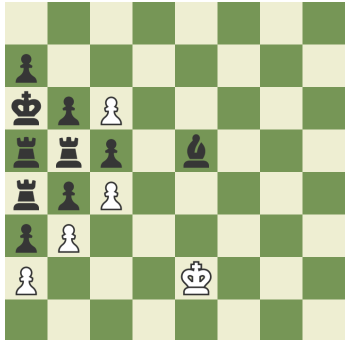
Board	Solution	To Play	Comment
<p>8/p7/kpP5/rrp1b3/rpP5/pP6/P3K3/8 w - -</p> 	<p>Draw</p>	<p>White</p>	<p>SF recognizes this position is a draw – the trend when comparing against the previous two positions is interesting because it suggests that there’s a limit to material advantage.</p>

Table 11 | Penrose set positions, Penrose variation 2, position 17

Board	Solution	To Play	Comment
<p>n1B4r/r1p2k2/BpPp3p/pP1PpBpP/P2bPpPb/ 1R1RbPb1/R1K1R1R1/1R3R2 w - -</p>	Draw	White	As the pawn structure is locked, and White is unable to remove any of the Black pawns by force, this is a trivial draw. Should White capture one of Black's dark-squared bishops, Black should avoid recapturing, and keep the blockade.
<p>n1B4r/r1p2k2/BpPp3p/pP1PpBpP/P2bPpPb/ 1R1RbPb1/R1K1B1R1/1R3R2 w - -</p>	Draw	White	As the pawn structure is locked, and White is unable to remove any of the Black pawns by force, this is a trivial draw. Should White capture one of Black's dark-squared bishops, Black should avoid recapturing, and keep the blockade.
<p>n1B4r/r1p2k2/BpPp3p/pP1PpBpP/P1RbPpPb/ 1R1RbPbB/R1K3R1/1R3R1R w - -</p>	Draw	White	As the pawn structure is locked, and White is unable to remove any of the Black pawns by force, this is a trivial draw. Should White capture one of Black's dark-squared bishops, Black should avoid recapturing, and keep the blockade.

Table 12 | Penrose set positions, Blocked set variation 1, positions 1-3. These represent the same theme (blockade), albeit with a slight variation in the total amount of material on each side, and therefore – the branching factor in search.

Board	Solution	To Play	Comment
<p>2rBb3/1kBbP3/1BbP3B/BbP1B3/bP5B/ P1B1B3/5K1B/B1B1B3 w - -</p> 	Draw	White	Despite the overwhelming material advantage, White is unable to penetrate on the light squares controlled by Black light-squared bishops. The position is also out of distribution in that it would not normally be possible to have as many bishops on the board.
<p>3Bb2R/1kBbP3/1BbP3B/BbP1B3/bP5B/ P1B1B3/5K1B/2B1B3 w - -</p> 	Draw	White	Despite the overwhelming material advantage, White is unable to penetrate on the light squares controlled by Black light-squared bishops. The position is also out of distribution in that it would not normally be possible to have as many bishops on the board. The addition of the rook makes it possible to exchange one bishop off, but doesn't change the evaluation.
<p>3Bb2R/1kBbP3/nBbP3B/BbP1B3/bP5B/ P1B1B3/5K1B/2B1B3 w - -</p> 	Draw	White	Despite the overwhelming material advantage, White is unable to penetrate on the light squares controlled by Black light-squared bishops. The position is also out of distribution in that it would not normally be possible to have as many bishops on the board. A further addition of the Knight on a6 should not affect the evaluation.

Table 13 | Penrose set positions, Blocked set variation 2, positions 1-3

Board	Solution	To Play	Comment
<p>5k2/5Bp1/5pP1/1K2pP2/4P3/1B3B2/8/5B2 w - -</p> 	Draw	White	White has no way of challenging the dark squares and the pawn on g7. Black should just move the king Kf8-e7-f8, and the game would end in a draw. If White vacates the a2-g8 diagonal, the Black king could be forced to g8 or h8, but that wouldn't change the outcome. Different amounts of extra material test for whether having extra light-squared bishops affects the engine's judgement.
<p>5k2/5Bp1/5pP1/1K2pP2/4P3/8/8 w - -</p> 	Draw	White	White has no way of challenging the dark squares and the pawn on g7. Black should just move the king Kf8-e7-f8, and the game would end in a draw. If White vacates the a2-g8 diagonal, the Black king could be forced to g8 or h8, but that wouldn't change the outcome. Different amounts of extra material test for whether having extra light-squared bishops affects the engine's judgement.
<p>B4k2/5BpB/2B2pP1/1K2pP1B/4P3/3B1B1B/B7/8 w - -</p> 	Draw	White	White has no way of challenging the dark squares and the pawn on g7. Black should just move the king Kf8-e7-f8, and the game would end in a draw. If White vacates the a2-g8 diagonal, the Black king could be forced to g8 or h8, but that wouldn't change the outcome. Different amounts of extra material test for whether having extra light-squared bishops affects the engine's judgement.

Table 14 | Penrose set positions, Blocked set variation 3, positions 1-3

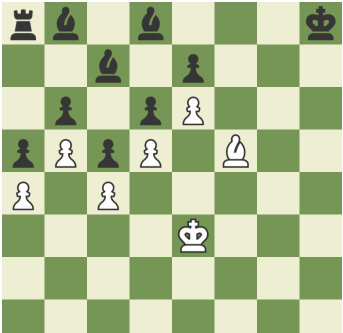
Board	Solution	To Play	Comment
<p>rb1b3k/2b1p3/1p1pP3/pPpP1B2/P1P5/4K3/8/8 w - -</p> 	Draw	White	<p>By getting the king to g5, White blocks the activation of the Black king, which cannot make further progress. Thanks to the light-squared bishop, which will penetrate the light squares on the kingside, Black will not be able to safely get the rook out, and the dark-squared bishops cannot escape. Sample line to illustrate the point: 1. Kf4 Ba7 2. Bg6 Rb8 3. Be8 Rb7 4. Bc6</p>

Table 15 | Penrose set positions, Blocked set variation 4, position 1

Board	Solution	To Play	Comment
<p>2b1rr2/1p1p1r1k/1PpPp2r/2P1Pp1r/5Pp1/3K2Pp/7P/8 w - -</p>	Draw	White	Despite the extra material, the bishop on c8 makes it impossible for Black to utilize any of the extra pieces, resulting in a fully blocked position, and therefore a trivial draw.
<p>2b1rr2/1p1p1r1k/1PpPp2r/2P1Pp1r/1R3Pp1/2RK2Pp/5B1P/2B5 w - -</p>	Draw	White	Same as the previous position, with some extra pieces for White. Black has sufficient extra material to defend against any potential exchange on b7, which is the only additional dynamic option introduced by the White rooks. Therefore, the evaluation is still a trivial draw, as neither side can make progress.
<p>2b1rr2/1p1p1r1k/1PpPp2r/2P1Pp1r/1R3Pp1/3K2Pp/5B1P/8 w - -</p>	Draw	White	As above, though with less material for White, to better evaluate how the extent of material differences may affect engine evaluation.

Table 16 | Penrose set positions, Blocked set variation 5, positions 1-3



Board	Solution	To Play	Comment
<p>2b2r2/1p1p3k/1PpPp3/2P1Pp1r/5Pp1/3K2Pp/7P/8 w - -</p> 	Draw	White	The bishop on c8 makes it impossible to utilize Black rooks, resulting in a trivial draw.
<p>2b2r2/1p1p3k/1PpPp3/2P1Pp2/5Pp1/3K2Pp/7P/8 w - -</p> 	Draw	White	The bishop on c8 makes it impossible to utilize Black rook, resulting in a trivial draw.

Table 17 | Penrose set positions, Blocked set variation 5, positions 4-5

Board	Solution	To Play	Comment
<p>6b1/2k3P1/4b1PQ/3b2P1/b1b3PQ/ 1b2p1PQ/b3P1P1/4KBRR w - -</p>	Draw	White	White is unable to utilize the overwhelming material advantage. White can exchange some extra material on the g8 square, though counting the pieces should make it obvious that at the end of all those exchanges, Black will remain with a bishop on g8. The bishop on a4 also makes it impossible to activate White's king.
<p>6bQ/1b1b2P1/4b1PQ/3b2P1/2b3PQ/ 1b2p1PQ/b3P1P1/2k1KBRR w - -</p>	Draw	White	As above, piece counting with respect to possible captures on g8 makes it clear that White's pieces on the h-file will not be able to activate. Compared to the previous position, there is a larger material advantage here, but the evaluation is the same.
<p>6bQ/1b1b2PR/2b1b1PQ/3b2P1/2b3PQ/ 1b2p1PQ/b3P1P1/2k1KBRR w - -</p>	Draw	White	As above, piece counting with respect to possible captures on g8 makes it clear that White's pieces on the h-file will not be able to activate. Compared to the previous position, there is a larger material advantage here, but the evaluation is the same.

Table 18 | Penrose set positions, Blocked set variation 6, positions 1-3



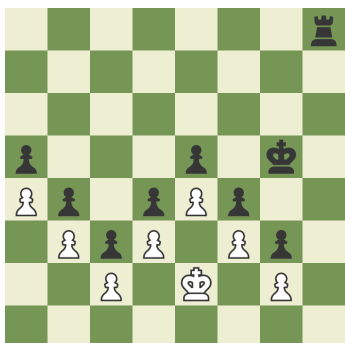
Board	Solution	To Play	Comment
<p>2k5/2p5/1q1p4/pPpPp1pp/N1P1Pp2/P4PbP/KQ4P1/8 w - -</p> 	Win	White	A. Petrosian-Hazai, 1970. Black is strategically lost, and tries his last chance by offering a queen sacrifice. The trap worked, but white should have instead fought over the a5 pawn and win.
<p>2k5/8/1p1p4/pPpPp1pp/2P1Pp2/P4PbP/KQ4P1/8 w - -</p> 	Draw	White	Black threatens to seal the kingside with ...h4 but White still has a window of opportunity to open entry lines into the black position: 1.h4 gxh4 when the h3-square is freed for the white queen. However after 2.Qb1 h3!! 3.gxh3 h4, the door is sealed shut once more!
<p>7r/8/p3p1k1/Pp1pPp2/1Pp1Pp1/2P1K1P1/8 b - -</p> 	Draw	Black	Black's problem is that the rook can only invade by entering White's first rank (e.g. via 1...Rh1) but that also turns out to be stalemate!

Table 19 | Penrose set positions, Detecting Fortresses in Chess, Dvoretzky, positions 1-3

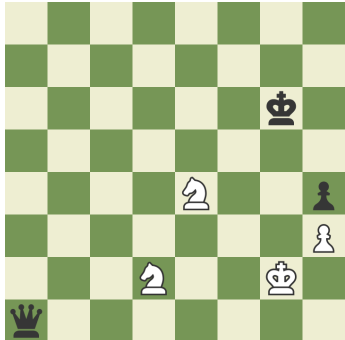
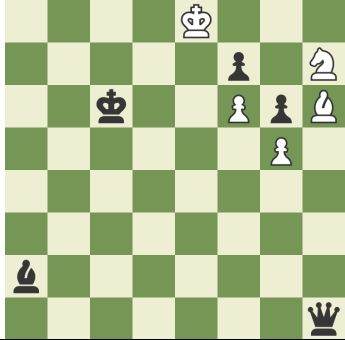
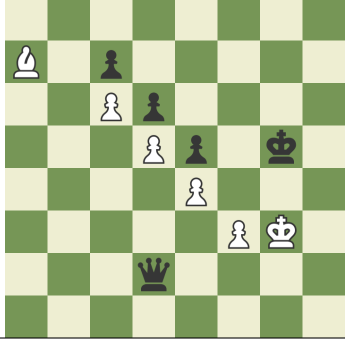
Board	Solution	To Play	Comment
<p>8/8/6k1/8/4N2p/7P/3N2K1/q7 b - -</p> 	Draw	Black	P. Pechenkin, 1953. The black king cannot go further than e2 and f4, while the queen cannot come close enough to the white king to deny him of the free squares in the corner.
<p>4K3/5p1N/2k2PpB/6P1/8/8/b7/7q b - -</p> 	Draw	Black	Nothing can stop the white king from heading to the fortified security of the g7 square!
<p>8/B1p5/2Pp4/3Pp1k1/4P3/5PK1/3q4/8 b - -</p> 	Draw	Black	Chekhover, 1948. Black king breakthrough on the king side is difficult, and it cannot march to the queen side since the white bishop controls the a7 square.

Table 20 | Value set positions, Detecting Fortresses in Chess, Dvoretsky, positions 4-6


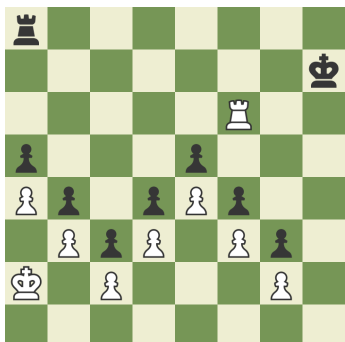
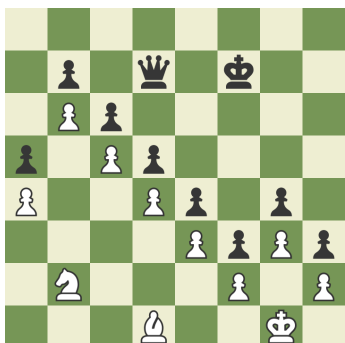
Board	Solution	Comment	
<p>3B4/1r2p3/r2p1p2/bkp1P1p1/1p1P1PPp/ p1P4P/PPB1K3/8 w - -</p> 	Draw	White	William Rudolph vs. Unknown player. White sacrifices the bishop, marches its pawns while delivering checks, and closes the position.
<p>r7/7k/5R2/p3p3/Pp1pPp2/ 1PpP1Pp1/K1P3P1/8 w - -</p> 	Draw	White	Hasek vs. Unknown player. White sacrifices its rook and builds a fortress.
<p>8/1p1q1k2/1Pp5/p1Pp4/P2Pp1p1/ 4PpPp/1N3P1P/3B2K1 w - -</p> 	Draw	White	One of those puzzles where the solution is still difficult to reconstruct even if you've solved it before! White has a choice of pieces to sacrifice and where to sacrifice them to create a fortress and the solution is the least obvious of all!

Table 21 | Challenge set positions, positions 1-3

Board	Solution	Comment	
<p>8/3P3k/n2K3p/2p3n1/1b4N1/2p1p1P1/8/3B4 w - -</p> 	Win	White	<p>Solution to a truly remarkable study, Grandmasters and Engines Couldn't Solve This Puzzle. Then Came The Magician.</p>
<p>4b1k1/P4pPp/1R3P1P/2r5/8/1P6/1K6/8 w - -</p> 	Win	White	<p>Sokolsky vs. Ravinsky, USSR 1938. Source: Tim Krabbe's blog, Under-promotion In Games. White promotes thr a7 pawn to a bishop on its way to a win. Other promotions end up with a stalemate.</p>
<p>n2Bqk2/5p1p/Q4KP1/p7/8/8/8/8 w - -</p> 	Win	White	<p>Most chess engines struggle with this beautiful mating pattern in which Black has many delaying moves to calculate, but a human can clearly see they will not help (Copeland, 2021a).</p>

Table 22 | Challenge set positions, positions 4-6

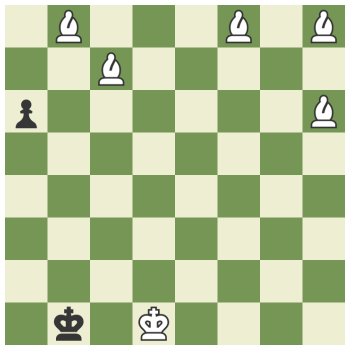
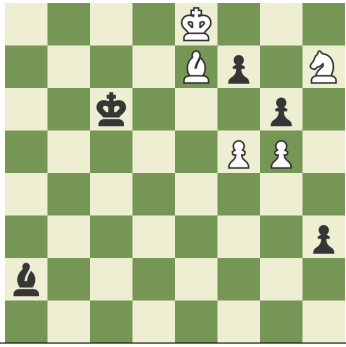
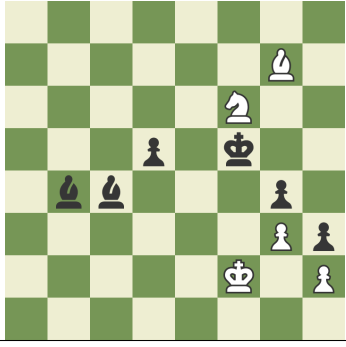
Board	Solution	Comment	
<p>1B3B1B/2B5/p6B/8/8/8/1k1K4 w - -</p> 	Win	White	<p>Troitsky, 1916. To win, white must force the black king to the a1 square, the black pawn to a2, the white king has to control the b1 square (via c1 or c2), and have a bishop checking the king on the long diagonal. This can be achieved via a sequence of maneuvers, zugzwangs (that force the black pawn to march to a2) and a discovered check. Otherwise, white cannot mate black with the same color bishops. Engines miss this long mating sequence.</p>
<p>4K3/4Bp1N/2k3p1/5PP1/8/7p/b7/8 w - -</p> 	Draw	White	<p>It's clear that White's cannot stop the h-pawn or create a passed pawn of his own. The dilemma is to find the most solid setup on the kingside to be able to hide the white king on g7 and keep the g5-pawn safe.</p>
<p>8/6B1/5N2/3p1k2/1bb3p1/6Pp/5K1P/8 w - -</p> 	Draw	White	<p>The resulting fortress position is not unfamiliar to human players but the added material makes it strangely difficult to visualise. It seems particularly strange that 2 mighty black bishops cannot drive the white king away from the kingside pawns.</p>

Table 23 | Challenge set positions, positions 7-9


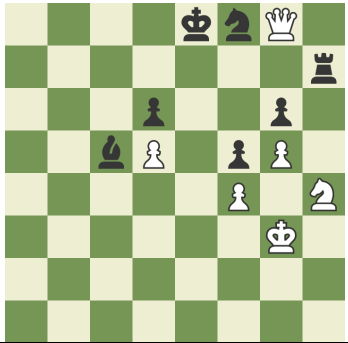

Board	Solution	Comment	
<p>7r/p3k3/2p5/1pPp4/3P4/PP4P1/3P1PB1/2K5 w - -</p> 	Draw	White	Chekhover, 1947. The Black Rook can infiltrate White's position and start capturing pawns. Yet White has a fortress like resource in his dispose.
<p>4knQ1/7r/3p2p1/2bP1pP1/5P1N/6K1/8/8 b - -</p> 	Draw	Black	White's queen is exerting a lot of pressure on Black's position but 1...Rxh4 2.Kxh4 Bd4! denies the queen access to any squares! Without any queen mobility, White is powerless to break the black fortress.
<p>7k/4p3/5p2/1p1p1p2/1PpP1PpB/1pP3P1/7P/6NK w - -</p> 	Draw	White	White's line to draw 1.Bxf6+ ef 2.h4 b2 3.Kg2 b1(Q) 4.Kf2 followed by Ne2.

Table 24 | Challenge set positions, positions 10-12

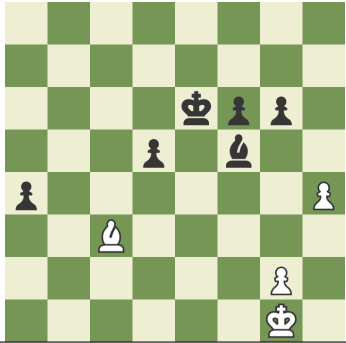
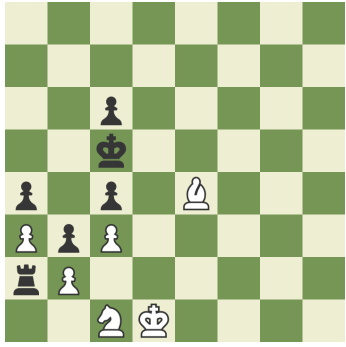
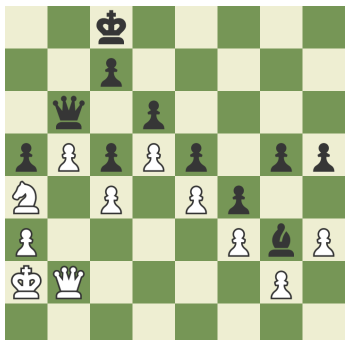
Board	Solution	Comment
<p data-bbox="252 275 558 297">8/8/4kpp1/3p1b2/p6P/2B5/6P1/6K1 b - -</p> 	<p data-bbox="627 450 683 472">Win</p>	<p data-bbox="774 450 847 472">Black</p> <p data-bbox="919 353 1382 566">Topalov-Shirov, Linares 1998 (Round 10). Here, the best move for black is to continue with Bh3. However, the uncoventional sacrifice is difficult for engines to see, and therefore they miss the win.</p>
<p data-bbox="252 707 558 730">8/8/2p5/2k5/p1p1B3/PpP5/rP6/2NK4 w - -</p> 	<p data-bbox="627 882 699 904">Draw</p>	<p data-bbox="774 882 847 904">White</p> <p data-bbox="919 667 1382 1126">Prophylaxis, the prevention of an opponent's ideas, is one of the most difficult things in chess. It's no surprise then that engines may also struggle with this concept. Prophylaxis is naturally most difficult when the opponent's ideas are well-hidden. In the following position, Black has a dangerous idea to play ...Ra2 and win the b2-pawn. White has one clever defense that requires anticipating the strength of ...Ra2 (Copeland, 2021a).</p>
<p data-bbox="212 1144 596 1178">2k5/2p5/1q1p4/pPpPp1pp/N1P1Pp2/P4PbP/KQ4P1/8 w - -</p> 	<p data-bbox="627 1323 683 1346">Win</p>	<p data-bbox="774 1323 847 1346">White</p> <p data-bbox="919 1211 1382 1458">Arshak B Petrosian vs. Laszlo Hazai. GM Laszlo Hazai offered a queen sacrifice, and GM Arshak Petrosian took the bait. Many chess engines fall for the same trap, while a careful play for white leads to a win (Copeland, 2021a).</p>

Table 25 | Challenge set positions, positions 13-15

References

- P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM, 2004.
- S. Alver and D. Precup. Constructing a good behavior basis for transfer using generalized policy updates. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=7IWGzQ6gZ1D>.
- A. P. Badia, P. Sprechmann, A. Vitvitskyi, D. Guo, B. Piot, S. Kapturowski, O. Tieleman, M. Arjovsky, A. Pritzel, A. Bolt, and C. Blundell. Never give up: Learning directed exploration strategies. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=Sye57xStvB>.
- D. Balduzzi, M. Garnelo, Y. Bachrach, W. Czarnecki, J. Perolat, M. Jaderberg, and T. Graepel. Open-ended learning in symmetric zero-sum games. In *International Conference on Machine Learning*, pages 434–443. PMLR, 2019.
- A. Barreto, W. Dabney, R. Munos, J. J. Hunt, T. Schaul, H. P. van Hasselt, and D. Silver. Successor features for transfer in reinforcement learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017a. URL <https://proceedings.neurips.cc/paper/2017/file/350db081a661525235354dd3e19b8c05-Paper.pdf>.
- A. Barreto, W. Dabney, R. Munos, J. J. Hunt, T. Schaul, H. P. van Hasselt, and D. Silver. Successor features for transfer in reinforcement learning. In *Advances in neural information processing systems*, pages 4055–4065, 2017b.
- A. Barreto, D. Borsa, J. Quan, T. Schaul, D. Silver, M. Hessel, D. Mankowitz, A. Zidek, and R. Munos. Transfer in deep reinforcement learning using successor features and generalised policy improvement. In *International Conference on Machine Learning*, pages 501–510. PMLR, 2018.
- A. Barreto, D. Borsa, S. Hou, G. Comanici, E. Aygün, P. Hamel, D. Toyama, S. Mourad, D. Silver, D. Precup, et al. The option keyboard: Combining skills in reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- K. Baumli, D. Warde-Farley, S. Hansen, and V. Mnih. Relative variational intrinsic control. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35:6732–6740, May 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16832>.
- S. Belogolovsky, P. Korsunsky, S. Mannor, C. Tessler, and T. Zahavy. Inverse reinforcement learning in contextual mdps. *Machine Learning*, 110(9):2295–2334, 2021.
- G. W. Brown. Iterative solution of games by fictitious play. *Act. Anal. Prod Allocation*, 13(1):374, 1951.
- A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, P. Schuh, K. Shi, S. Tsvyashchenko, J. Maynez, A. Rao, P. Barnes, Y. Tay, N. Shazeer, V. Prabhakaran, E. Reif, N. Du, B. Hutchinson, R. Pope, J. Bradbury, J. Austin, M. Isard, G. Gur-Ari, P. Yin, T. Duke, A. Levskaya, S. Ghemawat, S. Dev, H. Michalewski, X. Garcia, V. Misra, K. Robinson, L. Fedus, D. Zhou, D. Ippolito, D. Luan, H. Lim, B. Zoph, A. Spiridonov, R. Sepassi, D. Dohan, S. Agrawal, M. Omernick, A. M. Dai, T. S. Pillai, M. Pellat, A. Lewkowycz, E. Moreira, R. Child, O. Polozov, K. Lee, Z. Zhou, X. Wang, B. Saeta, M. Diaz, O. Firat, M. Catasta, J. Wei,

- K. Meier-Hellstern, D. Eck, J. Dean, S. Petrov, and N. Fiedel. Palm: Scaling language modeling with pathways, 2022.
- S. Copeland. 10 positions chess engines just don't understand, 2021a. URL <https://www.chess.com/article/view/10-positions-chess-engines-just-dont-understand>.
- S. Copeland. Hikaru nakamura crushes computer in blitz chess - top 10 of the 2000s - rybka vs. nakamura, 2008, 2021b. URL <https://www.chess.com/blog/SamCopeland/hikaru-nakamura-crushes-computer-in-blitz-chess-top-10-of-the-2000s-rybka-vs-nakamura-2008>.
- A. Cully. Autonomous skill discovery with quality-diversity and unsupervised descriptors. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 81–89, 2019.
- A. Cully and Y. Demiris. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, 22(2):245–259, 2017.
- A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.
- W. M. Czarnecki, G. Gidel, B. Tracey, K. Tuyls, S. Omidshafiei, D. Balduzzi, and M. Jaderberg. Real world games look like spinning tops. *Advances in Neural Information Processing Systems*, 33: 17443–17454, 2020.
- C. L. Dezsø and D. G. Ross. Does female representation in top management improve firm performance? a panel data investigation. *Strategic management journal*, 33(9):1072–1089, 2012.
- P. Doggers. Will this position help understand human consciousness?, 2017. URL <https://www.chess.com/news/view/will-this-position-help-to-understand-human-consciousness-4298>.
- M. Dvoretzky. *Dvoretzky's endgame manual*. SCB Distributors, 2010.
- B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine. Diversity is all you need: Learning skills without a reward function. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SJx63jRqFm>.
- A. Fawzi, M. Balog, A. Huang, T. Hubert, B. Romera-Paredes, M. Barekatin, A. Novikov, F. J. R Ruiz, J. Schrittwieser, G. Swirszcz, et al. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature*, 610(7930):47–53, 2022.
- Y. Flet-Berliac, J. Ferret, O. Pietquin, P. Preux, and M. Geist. Adversarially guided actor-critic. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=_mQp5cr_iNy.
- M. Geist, J. Pérolat, M. Laurière, R. Elie, S. Perrin, O. Bachem, R. Munos, and O. Pietquin. Concave utility reinforcement learning: The mean-field game viewpoint. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS '22*, page 489–497, Richland, SC, 2022. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450392136.
- G. Gigerenzer. From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological review*, 98(2):254, 1991.

- J. González-Díaz and I. Palacios-Huerta. Alphazero ideas, 2022.
- K. Gregor, D. J. Rezende, and D. Wierstra. Variational intrinsic control. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Skc-Fo4Yg>.
- M. Guid and I. Bratko. Detecting fortresses in chess. *Elektrotehniški Vestnik*, 79(1-2):35–40, 2012.
- J. Hadamard. *The mathematician’s mind: The psychology of invention in the mathematical field*, volume 109. Princeton University Press, 2020.
- E. Hazan, S. Kakade, K. Singh, and A. Van Soest. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, pages 2681–2691. PMLR, 2019.
- Z.-W. Hong, T.-Y. Shann, S.-Y. Su, Y.-H. Chang, T.-J. Fu, and C.-Y. Lee. Diversity-driven exploration strategy for deep reinforcement learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 10510–10521, 2018.
- V. Hunt, D. Layton, S. Prince, et al. Diversity matters. *McKinsey & Company*, 1(1):15–29, 2015.
- M. Jaderberg, V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, K. Simonyan, C. Fernando, and K. Kavukcuoglu. Population based training of neural networks, 2017.
- G. Kasparov. *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*. John Murray, 2017.
- G. Kasparov. Chess, a *drosophila* of reasoning. *Science*, 362(6419):1087–1087, 2018. doi: 10.1126/science.aaw2221. URL <https://www.science.org/doi/abs/10.1126/science.aaw2221>.
- G. K. Kasparov and D. King. *Kasparov against the world: the story of the greatest online challenge*. KasparovChess Online, Incorporated, 2000.
- D. M. Katz, M. J. Bommarito, S. Gao, and P. Arredondo. Gpt-4 passes the bar exam, 2023.
- M. Kemmerling, D. Lüticke, and R. H. Schmitt. Beyond games: A systematic review of neural monte carlo tree search applications, 2023.
- M. Klissarov and M. C. Machado. Deep laplacian-based options for temporally-extended exploration. *International Conference on Machine Learning (ICML)*, 2023.
- S. Kumar, A. Kumar, S. Levine, and C. Finn. One solution is not all you need: Few-shot extrapolation via structured maxent rl. *Advances in Neural Information Processing Systems*, 33:8198–8210, 2020.
- M. Lanctot, V. Zambaldi, A. Gruslys, A. Lazaridou, K. Tuyls, J. Pérolat, D. Silver, and T. Graepel. A unified game-theoretic approach to multiagent reinforcement learning. *Advances in neural information processing systems*, 30, 2017.
- J. Lehman and K. O. Stanley. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 211–218, 2011.
- D. S. Leslie and E. J. Collins. Generalised weakened fictitious play. *Games and Economic Behavior*, 56(2):285–298, 2006.

- R. L. Lewis, A. Howes, and S. Singh. Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science*, 6(2):279–311, 2014.
- S. Liu, M. Lanctot, L. Marris, and N. Heess. Simplex neural population learning: Any-mixture Bayes-optimality in symmetric zero-sum games. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 13793–13806. PMLR, 17–23 Jul 2022a. URL <https://proceedings.mlr.press/v162/liu22h.html>.
- S. Liu, L. Marris, D. Hennes, J. Merel, N. Heess, and T. Graepel. NeuPL: Neural population learning. In *International Conference on Learning Representations*, 2022b. URL https://openreview.net/forum?id=MIX3fJkl_1.
- X. Liu, H. Jia, Y. Wen, Y. Hu, Y. Chen, C. Fan, Z. Hu, and Y. Yang. Towards unifying behavioral and response diversity for open-ended learning in zero-sum games. *Advances in Neural Information Processing Systems*, 34:941–952, 2021.
- Y. Liu, P. Ramachandran, Q. Liu, and J. Peng. Stein variational policy gradient. In *33rd Conference on Uncertainty in Artificial Intelligence, UAI 2017*, 2017.
- A. Lupu, B. Cui, H. Hu, and J. Foerster. Trajectory diversity for zero-shot coordination. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 7204–7213. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/lupu21a.html>.
- A. Mandhane, A. Zhernov, M. Rauh, C. Gu, M. Wang, F. Xue, W. Shang, D. Pang, R. Claus, C.-H. Chiang, C. Chen, J. Han, A. Chen, D. J. Mankowitz, J. Broshear, J. Schrittwieser, T. Hubert, O. Vinyals, and T. Mann. Muzero with self-competition for rate control in vp9 video compression, 2022.
- M. Masood and F. Doshi-Velez. Diversity-inducing policy gradient: Using maximum mean discrepancy to find a set of diverse policies. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5923–5929. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/821. URL <https://doi.org/10.24963/ijcai.2019/821>.
- T. McGrath, A. Kapishnikov, N. Tomašev, A. Pearce, M. Wattenberg, D. Hassabis, B. Kim, U. Paquet, and V. Kramnik. Acquisition of chess knowledge in alphazero. *Proceedings of the National Academy of Sciences*, 119(47):e2206625119, 2022.
- C. Metz. Alphago versus lee sedol, 2016. URL <https://www.wired.com/2016/03/googles-ai-wins-pivotal-game-two-match-go-grandmaster/>.
- J. Mouret and J. Clune. Illuminating search spaces by mapping elites. *CoRR*, abs/1504.04909, 2015. URL <http://arxiv.org/abs/1504.04909>.
- M. D. Mumford. Where have we been, where are we going? taking stock in creativity research. *Creativity research journal*, 15(2-3):107–120, 2003.
- M. Mutti, R. D. Santi, P. D. Bartolomeis, and M. Restelli. Challenging common assumptions in convex reinforcement learning. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=13S0tUMqynI>.
- G. Nangue Tasse, S. James, and B. Rosman. A boolean task algebra for reinforcement learning. *Advances in Neural Information Processing Systems*, 33:9497–9507, 2020.

- O. Nilsson and A. Cully. Policy gradient assisted map-elites. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 866–875, 2021.
- OpenAI. Gpt-4 technical report, 2023.
- A. Osborn. *Applied imagination-principles and procedures of creative writing*. Read Books Ltd, 2012.
- S. E. Page. *The diversity bonus: How great teams pay off in the knowledge economy*. Princeton University Press, 2019.
- J. Parker-Holder, A. Pacchiano, K. M. Choromanski, and S. J. Roberts. Effective diversity in population based reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pages 2778–2787. PMLR, 2017.
- Z. Peng, H. Sun, and B. Zhou. Non-local policy optimization via diversity-regularized collaborative exploration, 2020.
- R. Penrose. *Shadows of the Mind*, volume 4. Oxford University Press Oxford, 1994.
- R. Penrose and N. D. Mermin. *The emperor’s new mind: Concerning computers, minds, and the laws of physics*, 1990.
- S. R. Penrose. Sir roger penrose and dr. stuart hameroff: Consciousness and the physics of the brain, 2020. URL <https://www.youtube.com/watch?v=xGbgDf4HCHU>.
- N. Perez-Nieves, Y. Yang, O. Slumbers, D. H. Mguni, Y. Wen, and J. Wang. Modelling behavioural diversity for learning in open-ended games. In *International Conference on Machine Learning*, pages 8514–8524. PMLR, 2021.
- T. Pierrot, V. Macé, F. Chalumeau, A. Flajolet, G. Cideron, K. Beguir, A. Cully, O. Sigaud, and N. Perrin-Gilbert. Diversity policy gradient for sample efficient quality-diversity optimization. In *ICLR Workshop on Agent Learning in Open-Endedness*, 2022.
- J. K. Pugh, L. B. Soros, and K. O. Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.
- C. D. Rosin. Multi-armed bandits with episode context. *Annals of Mathematics and Artificial Intelligence*, 61(3):203–230, 2011.
- M. Sadler and N. Regan. *Game changer*. New in Chess, 2019.
- R. Sanjaya, J. Wang, and Y. Yang. Measuring the non-transitivity in chess. *Algorithms*, 15(5):152, 2022.
- J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- L. Shani, T. Zahavy, and S. Mannor. Online apprenticeship learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8240–8248, 2022.
- A. Sharma, S. Gu, S. Levine, V. Kumar, and K. Hausman. Dynamics-aware unsupervised discovery of skills. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=HJgLZR4KvH>.

- S. Shead. How deepmind boss demis hassabis used chess to get billionaire peter thiel to ‘take notice’ of his ai lab, 2020. URL <https://www.cnbc.com/2020/12/07/deepminds-demis-hassabis-used-chess-to-get-peter-thiels-attention.html>.
- M. Shin, J. Kim, B. van Opheusden, and T. L. Griffiths. Superhuman artificial intelligence can improve human decision-making by increasing novelty. *Proceedings of the National Academy of Sciences*, 120(12):e2214840120, 2023.
- D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018a. doi: 10.1126/science.aar6404. URL <https://www.science.org/doi/abs/10.1126/science.aar6404>.
- D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018b.
- D. Silver, S. Singh, D. Precup, and R. S. Sutton. Reward is enough. *Artificial Intelligence*, 299:103535, 2021. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2021.103535>. URL <https://www.sciencedirect.com/science/article/pii/S0004370221000862>.
- J. Singh and L. Fleming. Lone inventors as sources of breakthroughs: Myth or reality? *Management science*, 56(1):41–56, 2010.
- K. Singhal, T. Tu, J. Gottweis, R. Sayres, E. Wulczyn, L. Hou, K. Clark, S. Pfohl, H. Cole-Lewis, D. Neal, M. Schaeckermann, A. Wang, M. Amin, S. Lachgar, P. Mansfield, S. Prakash, B. Green, E. Dominowska, B. A. y Arcas, N. Tomasev, Y. Liu, R. Wong, C. Semturs, S. S. Mahdavi, J. Barral, D. Webster, G. S. Corrado, Y. Matias, S. Azizi, A. Karthikesalingam, and V. Natarajan. Towards expert-level medical question answering with large language models, 2023.
- P. D. Stokes. *Creativity from constraints: The psychology of breakthrough*. Springer Publishing Company, 2005.
- H. Sun, Z. Peng, B. Dai, J. Guo, D. Lin, and B. Zhou. Novel policy seeking with constrained optimization, 2021. URL <https://openreview.net/forum?id=iEcqwosBEgx>.
- D. Tarapore, J. Clune, A. Cully, and J.-B. Mouret. How do different encodings influence the performance of the map-elites algorithm? In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pages 173–180, 2016.
- G. N. Tasse, S. James, and B. Rosman. Generalisation in lifelong reinforcement learning through logical composition. In *International Conference on Learning Representations*, 2021.
- F. Timbers, N. Bard, E. Lockhart, M. Lanctot, M. Schmid, N. Burch, J. Schrittwieser, T. Hubert, and M. Bowling. Approximate exploitability: learning a best response. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3487–3493, 2022.
- A. Tirinzoni, A. Al-Marjani, and E. Kaufmann. Optimistic pac reinforcement learning: the instance-dependent view. In *International Conference on Algorithmic Learning Theory*, pages 1460–1480. PMLR, 2023.

- B. Tjanaka, M. C. Fontaine, J. Togelius, and S. Nikolaidis. Approximating gradients for differentiable quality diversity in reinforcement learning. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '22*, page 1102–1111, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392372. doi: 10.1145/3512290.3528705. URL <https://doi.org/10.1145/3512290.3528705>.
- N. Tomašev, U. Paquet, D. Hassabis, and V. Kramnik. Assessing game balance with alphazero: Exploring alternative rule sets in chess, 2020.
- A. Turing. Intelligent machinery. *Philosophia Mathematica*, page 395, 1948.
- H. Van Hasselt, A. Guez, and D. Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- V. Vassiliades, K. Chatzilygeroudis, and J.-B. Mouret. Using centroidal voronoi tessellations to scale up the multi-dimensional archive of phenotypic elites algorithm, 2017.
- O. Vinyals, I. Babuschkin, J. Chung, M. Mathieu, M. Jaderberg, W. M. Czarnecki, A. Dudzik, A. Huang, P. Georgiev, R. Powell, et al. Alphastar: Mastering the real-time strategy game starcraft ii. *DeepMind blog*, 2:20, 2019.
- T. T. Wang, A. Gleave, N. Belrose, T. Tseng, J. Miller, M. D. Dennis, Y. Duan, V. Pogrebniak, S. Levine, and S. Russell. Adversarial policies beat professional-level go ais. *arXiv preprint arXiv:2211.00241*, 2022.
- Wikipedia. Roger penrose, 2023. URL https://en.wikipedia.org/wiki/Roger_Penrose.
- S. Wolfram. *The Second Law: Resolving the Mystery of the Second Law of Thermodynamics*. Wolfram Media, 2023.
- T. Zahavy, A. Cohen, H. Kaplan, and Y. Mansour. Apprenticeship learning via frank-wolfe. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6720–6728, 2020a.
- T. Zahavy, A. Hasidim, H. Kaplan, and Y. Mansour. Planning in hierarchical reinforcement learning: Guarantees for using local policies. In *Algorithmic Learning Theory*, pages 906–934, 2020b.
- T. Zahavy, A. Barreto, D. J. Mankowitz, S. Hou, B. O’Donoghue, I. Kemaev, and S. Singh. Discovering a set of policies for the worst case reward. In *International Conference on Learning Representations*, 2021a. URL <https://openreview.net/forum?id=PUkhWz65dy5>.
- T. Zahavy, B. O’Donoghue, G. Desjardins, and S. Singh. Reward is enough for convex mdps. *Advances in Neural Information Processing Systems*, 34:25746–25759, 2021b.
- T. Zahavy, Y. Schroecker, F. Behbahani, K. Baumli, S. Flennerhag, S. Hou, and S. Singh. Discovering policies with DOMiNO: Diversity optimization maintaining near optimality. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=kjkdzBW3b8p>.
- J. Zhang, A. Koppel, A. S. Bedi, C. Szepesvari, and M. Wang. Variational policy gradient method for reinforcement learning with general utilities. *Advances in Neural Information Processing Systems*, 33:4572–4583, 2020.