
Improving the Strength of Human-Like Models in Chess

Saumik Narayanan

Washington University in St. Louis
saumik@wustl.edu

Kassa Korley

Microsoft Research
kassakorley@gmail.com

Chien-Ju Ho

Washington University in St. Louis
chienju.ho@wustl.edu

Siddhartha Sen

Microsoft Research
sidsen@microsoft.com

Abstract

Designing AI systems that capture human-like behavior has attracted growing attention in applications where humans may want to learn from, or need to collaborate with, these AI systems. Many existing works in designing human-like AI have taken a supervised learning approach that learns from data of human behavior, with the goal of creating models that can accurately predict human behavior. While this approach has shown success in capturing human behavior, it also suffers from the drawback of mimicking human mistakes. Moreover, existing models only capture a snapshot of human behavior, leaving the question of how to improve them largely unanswered. Using chess as an experimental domain, we investigate the question of teaching an existing human-like model to be stronger using a data-efficient curriculum, while maintaining the model’s human similarity. To achieve this goal, we extend the concept of curriculum learning to settings with multiple labeling strategies, allowing us to vary both the *curriculum* (dataset) and the *teacher* (labeling strategy). We find that the choice of teacher has a strong impact on both playing strength and human similarity; for example, a teacher that is too strong can be less effective at improving playing strength and degrade human similarity more rapidly. We also find that the choice of curriculum can impact these metrics, but to a smaller extent. Finally, we show that our strengthened models achieve human similarity at higher-level datasets, suggesting human-like improvement.

1 Introduction

AI systems are growing increasingly capable at solving tasks, making decisions, and assisting humans in a wide number of domains. In games such as Chess, Go, and Poker, AI has demonstrated clear superiority over human performance (18; 3). In domains such as marketing, transportation, medicine, law, hiring, and finance, AI decision making is capable enough to be deployed alongside humans in many real-life scenarios, and in some cases, AI performance is sufficient to take over entirely.

More recently, researchers have looked into developing AI with the explicit goal of replicating human decision making, in contrast to optimizing for AI performance. In domains such as autonomous driving, developing human-like AI can result in increased safety for both humans and AI, by providing an understanding of human driving tendencies (6). In situations where AI is assisting or educating humans, humans exhibit a higher level of trust when the AI is human-like (22; 11; 9), and higher levels of satisfaction when interacting with human-like AI (1; 17; 8). In gaming, human-like AI can offer opportunities to practice against models of real opponents (15; 13), and is generally more enjoyable to play with in competitive or cooperative play (24; 19).

Modern efforts to create human-like AI are improving rapidly, but these systems naturally lead to reduced task performance as a result of focusing on human similarity. For example in chess, researchers have used a supervised learning approach on human game data to create models that match human decision-making at different skill levels (14); these models can even be fine-tuned to match the decisions of an individual player and distinguish between individual playing styles (15; 16). Since these models strive to match human behavior at different skill levels, they also adopt the human tendencies and mistakes made at those levels. Moreover, each of these models only represents a snapshot of human behavior, with no investigation of how to improve the models.

In this paper, we attempt to address these questions by developing a methodology for improving the strength of human-like models, while retaining their human similarity. Specifically, given a human-like model at a particular skill level, our goal is to improve the model’s strength in a data-efficient manner, without degrading its human similarity by too much. Thus, we evaluate these human-like models along a trade-off of three metrics: strength, efficiency, and human similarity (Section 3.2).

Chess is an ideal domain for our study due to the availability of AI systems at both ends of the strength vs. human similarity spectrum. Chess engines like Stockfish and Leela (the open-source replication of AlphaZero (18)) dominate the best human chess players in strength, while Maia Chess (14) provides a framework for training models that can accurately mimic human behavior at different skill levels.

In devising a methodology for strengthening human-like models, we considered two main approaches. The first is to leverage the reinforcement learning framework of AlphaZero, but with an additional constraint to match human moves. However, is difficult to integrate into non-RL, human-like models. Additionally, this process is far removed from any real human learning process.

Instead, we chose a methodology that more closely mimics how a human might learn: starting with an existing human-like model, which we call the *student*, we fine-tune the model using a customized training curriculum. The core idea behind our approach is quite similar to the concept of curriculum learning, which looks at selecting the most efficient ordering of the data for training a model (2). However, we extend this concept in an important way, by not just looking at data selection, but also considering the *strategy for labeling the data*. That is, we distinguish between the input dataset, which we call the *curriculum*, and the labeling strategy, which we call the *teacher*. In the case of chess, the curriculum corresponds to a set of chess positions, and the teacher could be any chess model that suggests which move to make in each position. Whereas traditional curriculum learning assumes a single, “correct” label, in our work we consider chess engines and human-like models of varying strength, which may have different opinions about which move to make in a given position. In our evaluation, we explore what happens to learning when we vary the curriculum and the teacher.

Our results lead to several interesting conclusions. We find that the choice of teacher has a strong impact on both playing strength and human similarity, and it is not the case that the strongest teacher is the best choice. For example, a teacher like Stockfish is too strong and different when the student is a human-like model of novice strength. Instead, using a stronger human-like model as a teacher leads to much better outcomes in strength, human similarity, and training efficiency. We also find that the choice of curriculum can impact these metrics, but to a smaller extent. For example, training on a curriculum of positions the student made mistakes on provides only a marginal benefit over training on random positions the student encountered. Finally, we show that our strengthened models achieve human similarity on datasets corresponding to their strengthened level of play (Section 4.4), suggesting that our curriculum training methodology is improving these models in human-like steps.

2 Related Work

Chess has been a useful domain for measuring the capability of AI, and researchers for half a century have worked on developing stronger chess engines. The first breakthrough occurred in 1997, when Deep Blue defeated the reigning world chess champion in a match (4). By the mid-2000s, similar alpha-beta engines could consistently beat the strongest humans on standard computer hardware, and this trend continues with Stockfish, currently the strongest chess engine in the world.

In 2018, the release of AlphaZero showed that chess engines trained using deep reinforcement learning could rival traditional alpha-beta engines (18). Recently, the Maia engine was released, which repurposes the AlphaZero architecture to use supervised learning on human games to mimic human play directly, instead of trying to learn optimal play (14). Maia was trained on a massive public

human chess dataset provided by the website Lichess.org, and showed the ability to mimic human behavior at a variety of different skill levels. Further work expanded the capability of Maia to mimic individual players (15) as well as distinguish between styles (16). Overall, the chess contains some of the most sophisticated work on human-like AI, further supporting its position as the “Drosophila of AI” (12). However, research on human-like AI extends far beyond chess. Researchers have developed human-like AI for games like Go and Diplomacy (7) and for autonomous vehicles (6).

In general, optimizing AI towards human similarity will necessarily lead to reduced task performance, especially in scenarios where AI achieves superhuman performance. In this paper, we choose to handle this trade-off by extending curriculum learning (2) to support varying both the curriculum and the labeling strategy (teacher). When training an existing human-like model on a particular curriculum and teacher, we adopt the fine-tuning methodology of McIlroy et al. (15), which itself draws on recent methods from the transfer learning literature (20; 23; 10). While we intend this to be a domain-agnostic framework for efficiently developing strong and human-like AI systems, our choice of chess allows us to draw from existing chess pedagogical work to inform experiments. There is strong empirical evidence supporting the notion that human chess training can build knowledge to improve future skill (21). There are many avenues that human chess players can take to study, but there is a wide disagreement as to which of these is the most optimal path for training (5).

3 Methodology

The goal of this work is to investigate data-efficient approaches for improving the strength of a given human-like model, while maintaining its human-like property. In this section, we first describe the procedure we use to train human-like models in the domain of chess. We then define our performance metrics, which allow us to evaluate the trade-off between data efficiency, model strength, and human similarity when improving a given model. Finally, we introduce our curriculum training approach, which extends the concept of curriculum learning to multiple teachers, enabling us to examine the impact of varying both the curriculum (dataset) and the teaching advice (label strategy).

3.1 Training Human-Like Models in Chess

To develop human-like models, we replicate and extend the Maia framework developed by McIlroy et al. (14). We acquire a database of human chess moves from the online platform Lichess.org, where each move is labeled according to the player’s Elo (skill) rating. We then group the moves by Elo rating, and use supervised learning to train a neural network to predict what move a human would play in a given chess position. For example, we train a model on moves made by players rated between 1150 to 1249, and we denote this model as Maia-1200. The neural network we use is a slightly modified version of Maia’s neural network, which itself is adapted from Leela’s architecture. Notably, we significantly optimized the preprocessing pipeline and trained a broader set of models, going up to Maia-2600. Further information about our training procedure can be found in the appendix.

3.2 Metrics

Given a human-like model, we aim to improve its strength in a data-efficient manner, while maintaining its human similarity. We discuss the metrics we use to evaluate these performance measures.

Data efficiency: Data size. For data efficiency, we measure the amount of data used to train or update the model. Note that while we have access to an abundant amount of data in the domain of chess, if we are to generalize our research to other domains with practical data limitations, minimizing the amount of data usage will play an important role (Section 5).

Human similarity: Move-matching accuracy. To evaluate whether the model maintains its human-like playing style, we evaluate how well the updated model can predict human behavioral data. In general, we perform this evaluation specifically on the behavioral data used to train the original model, as this provides a direct way to measure the deviation from the original model. However, in the case of chess, we also have access to human data at a broad range of skill levels, so we also present results to evaluate our updated models against this additional data (Section 4.4).

To provide more insight into this metric, we trained a set of human-like models (following the procedure in Section 3.1) from Maia-1200 to Maia-2600, and measure their move-matching accuracy

on the datasets with player ratings from 1200 to 2600. The results, shown in Figure 1a, demonstrate that the Maia- K model is the most accurate in predicting the data of players of rating K . In particular, it is not the case that higher-rated models are better at predicting lower-rated datasets. This result replicates the findings of McIlroy et al. (14) and provides evidence that move-matching accuracy is a reasonable measure to evaluate whether the updated model is still human-like.

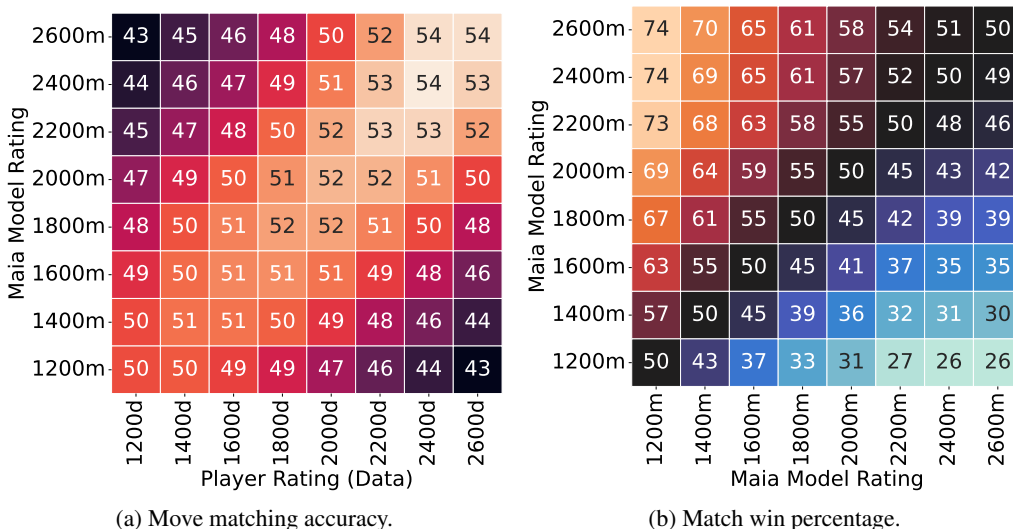


Figure 1: Illustrating the metrics of move-matching accuracy and match win percentage.

Strength: Match win percentage. To evaluate the strength of the updated model, we have the updated model play against the original model in a large number of full-length chess games (until checkmate or a draw is reached), and measure the match win percentage of the updated model. This win percentage is then used as a proxy for measuring the model strength.

To provide more insight into this metric, we have the models from Maia-1200 to Maia-2600 play against each other. The match win percentages (based on 5120 games played between each pair of models) are shown in Figure 1b. We can see that as we increase the distance between two models, the win rate of the higher-rated model increases. We also see that as ratings go up, the win rates tend to be closer to 50%. In general, these win rates correlate with the general notion of chess Elo ratings, but the magnitudes are not the same. In the standard Elo system, a 200 point difference should correspond to a 75% win rate, but in our models it corresponds to between 51% and 58% win rate. Jacob et. al (7) showed that higher-rated Maia models trained using standard imitation learning were suboptimal, and claimed that this would lead to weaker play than expected. Our results support this claim, by showing that the difference between (for example) a Maia-1800 and Maia-1600 is smaller than the difference between an 1800-rated human and a 1600-rated human.

3.3 Curriculum Training

We formally introduce the notion of *curriculum training* as the data-efficient fine-tuning of an existing human-like model, with the dual objectives of improving the model’s strength and preserving the human similarity of the original model.

In curriculum training, we start with a human-like model M , which we call the *student*. Our goal is to choose a pair of *curriculum* X and *teacher* Y to update the given model. A curriculum is specified by the input distribution X . For example, X could be a distribution of chess positions that players of a certain rating have encountered in their games. A teacher is specified by the labeling strategy Y , which specifies the labels for each input x drawn from X . For example, Y could be the move a particular human would make in the given position, or the move a superhuman engine would make.

After specifying the curriculum and teacher (X, Y) , we draw data points $\{(x, y)\}$ from the provided distribution and use them to update the model M . Following the curriculum learning approach, we update the model by performing gradient descent using this data. Note that our curriculum training

notion extends the standard curriculum learning notion developed by Bengio et. al (2) to settings with multiple teachers (labeling strategies). To evaluate the performance of curriculum training, we vary the amount of data used to update the model M and measure the strength (match win percentage) and human similarity (move-matching accuracy) of the updated model, as per Section 3.2.

In this work, we use the domain of chess to evaluate the effectiveness of curriculum training. We pre-trained a set of Maia models to serve as the initial human-like model M and the potential teachers Y . We vary the choice of teachers Y to be different ratings of human players or human-like models, or a superhuman chess engine, and examine their impact. We also examine the impact of choosing curriculum X based on the potential mistakes that the student might make, using the assistance of a superhuman chess engine, stronger-rated Maia models, or a human expert.

4 Experiments

In this section, we examine the effectiveness of curriculum training in the domain of chess. We adopt the Maia-1200 model as our student (i.e., the initial human-like model we wish to improve), and update it using different designs of curriculum X and teacher Y , varying both X and Y (Section 4.1), varying only Y (Section 4.2), and varying only X (Section 4.3). We evaluate the effectiveness of curriculum training on our three performance measures: data-efficiency, human similarity, and model strength. We then extend our evaluation of human similarity to cover datasets of higher-rated players, to measure the progress of the student with respect to higher-rated human-like models (Section 4.4).

4.1 Experiment 1: Higher-Rated Player Data

We first examine the most natural design of the curriculum and teacher pair: updating the initial human-like model Maia-1200 using data from higher-rated human players. Specifically, let X_R be the distribution of positions encountered by R -rated players, and Y_R be the moves made by R -rated players in these positions. (Incidentally, this is the dataset used to train the Maia- R models.) In this experiment, we set our student M as Maia-1200, and we vary the rating R from 1400 to 3000 and use (X_R, Y_R) as the (curriculum, teacher) to update the student. To investigate the efficiency of this design, we gradually increase the amount of data used to update the student model. For each resulting model (updated using a different amount of data and different (X_R, Y_R)), we measure its human similarity (i.e., move matching accuracy) and strength (i.e., match win percentage).

For each design from $R = 1400$ to 3000, we fine-tune the student for up to 10,000 batches of data drawn from (X_R, Y_R) , where a batch is a set of 1024 data points. Due to a lack of data for higher-rated players, the size of the training curriculum is smaller for $R = 2600, 2800,$ and 3000 (with 7300, 900, and 100 batches, respectively). We evaluate win percentage with a match size of 1024 games, and move-matching accuracy using 10 batches. We compute these evaluations every 10 batches between batch-0 and batch-100, and every 100 batches between batch-100 and batch-10,000. We repeat each set of evaluations 10 times, for a total of 10,240 games and 100 batches per evaluation, and report the moving average (with a window size of 50) of these measurements.

Experiment results. The results, as displayed in Figure 2, demonstrate the trade-offs of our three performance metrics for different pairs of curriculum and teacher (X_R, Y_R) . First, if our goal is purely to increase strength without concerns for human similarity or efficiency, training using the highest-rated curriculum leads to the strongest model.

If our objective is to increase the model’s strength while maintaining the model’s human similarity above a selected threshold (without concern for data efficiency), using the closest-rated data for fine-tuning leads to the optimal outcome. The drawback is that training on low rated data will eventually hit a plateau—not only do we require an order of magnitude more data to achieve the same performance, but the overall gains are limited. Figure 1a suggests that training a 1200-student on 1400-rated data can only achieve 57% win rate, even with an infinite amount of data, while higher-rated curricula easily surpasses this limit.

If our objective is to maximize strength gains with a constraint on data usage, we see an interesting and non-trivial result: at many points during fine-tuning, the optimal curriculum is not the highest-rated or lowest-rated one, but at a point in between. For example, at batch-30, the 2000-rated curriculum is optimal, while at batch-100 and batch-300, the 2400-rated curriculum is optimal. Past this point, it appears that the 2600-rated curriculum is optimal at all batch sizes. However, it would not be

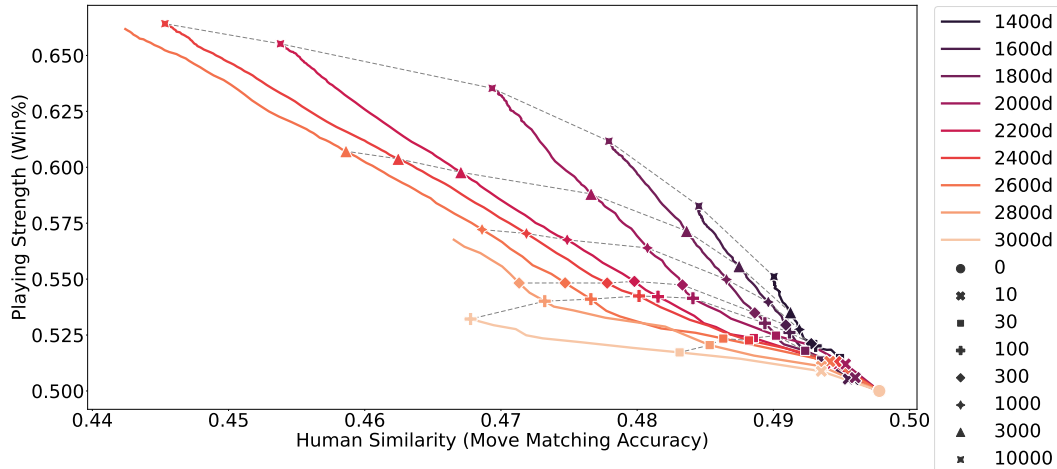


Figure 2: The results of using higher-rated player data to fine-tune a lower-rated student model. Each curve represents the change in human similarity and strength when fine-tuning the Maia-1200 model on increasing amounts of data from players of a particular rating. For each curve, we label the points that are trained using different numbers of batches, from 10 batches to 100,000 batches.

surprising if the trend continues: if we had access to the 3000-rated curriculum at higher batch sizes, the high-rated model might perform worse than a 2600-rated or 2800-rated curriculum.

Overall, these results suggest that the strength of a teacher matters significantly when training models. Note that in this experiment, we have simultaneously updated the curriculum X and the teacher Y . In order to separately examine the effects of curricula and the effects of teachers on the fine-tuning process, we run two additional experiments to investigate these factors separately.

4.2 Experiment 2: Effect of Teachers

To understand the effect of teachers on training, we hold X constant and vary Y . For these results, we chose to fix X as the set of board positions that the student was initially trained on. In this case, X would be the distribution of board positions encountered by 1200-rated players.

However, when we fix the input distribution X , we are unable to vary Y as in the previous experiment, as we might not have data on which moves, say 1400-rated players, would actually make in our selected X distribution. As an approximate solution, we use Maia models that we have trained to serve as the teacher Y . Thus, we teach Maia-1200 using a curriculum where X is the set of positions encountered by 1200-rated humans, and Y are the moves that Maia-1400 would make in these positions. We also introduce an additional possible teacher, Stockfish, the strongest chess engine in the world. This allows us to approximate much higher rated humans, since no models higher than Maia-2600 are available. Using this curriculum design, we ran another experiment with the same training procedure and evaluation methods as Experiment 1.

Experiment results. The results are displayed in Figure 3. We see that, similar to the earlier experiment, move-matching accuracy is monotonically decreasing with respect to teacher rating. We also see the same non-trivial result that the most data-efficient curriculum is often not the highest or lowest-rated teacher, but somewhere in the middle. In particular, the most efficient teacher is Maia-2200, Maia-2400, or Maia-2600, and these three points tend to be clustered closely together. Interestingly, we see that Stockfish performs noticeably worse than most of the human-rated models. Overall, these results confirm our observation in Experiment 1 that the teacher has a significant effect in fine-tuning a human-like model.

4.3 Experiment 3 - Effect of Curricula

We now hold the teacher Y constant and vary X to examine the effect of curricula. From the previous experiment, we see that Maia-2200, Maia-2400, and Maia-2600 are the most efficient teachers for




Figure 3: The results of varying the teacher Y to Stockfish (a superhuman chess engine) and Maia models of different ratings, while fixing the input distribution X to the student’s positions.

retaining human similarity and improving playing strength. Accordingly, we choose, the middle of these, Maia-2400, as the teacher for the next set of experiments.

To choose our possible learning curricula, we use domain knowledge in chess to develop curricula based on fixing mistakes made by the student. We list three categories of mistakes, and choose one representative curriculum for each category to analyze.

- Mistakes identified by a *superhuman chess engine*. We choose the position of blunders, which is defined as a move which lowers expected win probability by over 15% as predicted by Stockfish.
- Mistakes identified by a *human expert*. We collaborated with an International Master who analyzed 80 simulated games played by Maia models at various strengths, and determined that Maia tends to play suboptimally in positions where Maia should “connect the rooks”. As this is also a common problem for human chess players, we selected this as our representative curriculum.
- Mistakes identified by a *human-like model*. The final category is positions where different level Maia models disagree. While this is not conclusive evidence of a mistake, it is a possible indication that the lower-level Maia is playing suboptimally. We selected positions where Maia-1200 and Maia-1400 differ, which we denote as 1400-inflection positions.

Using these curricula and the Maia-2400 teacher, we ran another experiment with the same training procedure and evaluation methods as Experiments 1 and 2. For context, we also present a baseline curriculum which is just a random set of positions encountered at the student’s rating level.

Experiment results. The results are displayed in Figure 4. If our objective is to maximize strength gains with a constraint on the data usage, both the blunder curriculum and the 1400-inflection curriculum are more efficient at teaching the base model than a random curriculum, while the rooks curriculum underperforms in this regard. One hypothesis is that the rooks curriculum (human expert recommendation) is too specific in the positions it tries to teach. On the other hand, the blunder and 1400-inflection curricula look at mistakes that the model makes across the entire distribution of general board states. If our objective is to increase the model strength while ensuring the model maintains human similarity, the blunder curriculum outperforms random while the 1400-inflection and rooks curriculum do not. However, the difference in performance is small across curricula.

4.4 Experiment 4 - Measuring Human Similarity

Until this point, we evaluated human similarity by using the 1200-rated dataset. Because the student was originally trained on this dataset, it makes sense to use this as a way to measure model deviation.

Thanks to the rating metadata included in the human chess dataset, we can go one step further. During curriculum training, the model will naturally degrade in its initial performance, but it might play more like a higher-rated model. To investigate, we took the 11 models in (Section 4.2) and

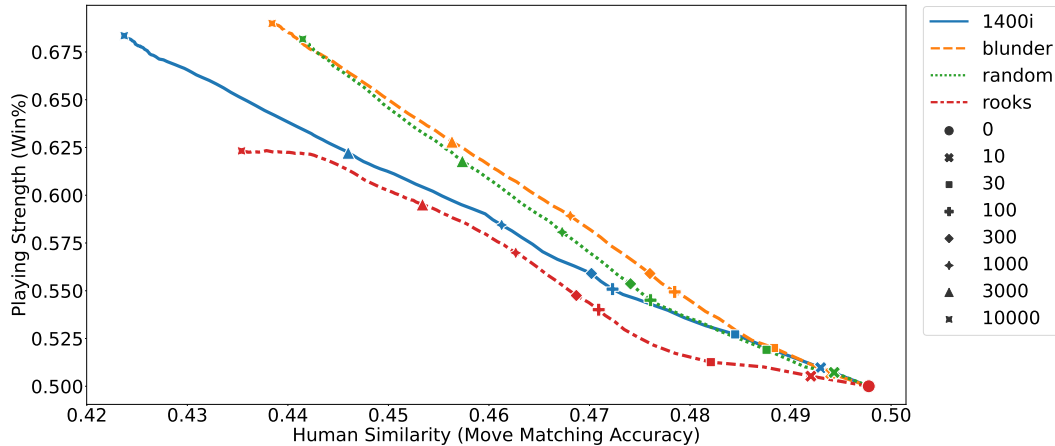


Figure 4: The results of fixing Y to Maia-2400 while varying X to positions identified by a superhuman chess engine (blunder), a human expert (rooks), and a human-like AI (1400-inflection).

(Section 4.3) after training 10,000 batches, and analyzed their move-matching accuracy against higher-rated datasets.

Experiment results. The results are displayed in Figure 5. Overall, we find that the updated models are more similar to higher-rated datasets than the original. There is a strong correlation between the rating of the teacher and the most similar dataset of the updated model. For instance, the model taught by Maia-1800 is most similar to the 1800-rated dataset. Interestingly, each of the models taught using targeted curricula and Maia-2400 teacher all perform most similar to the 2400-dataset. This provides more credence to our earlier observation that the teacher is much more important than the curriculum.

5 Conclusion and Discussion

Despite a growing effort to develop human-like models, surprisingly little attention has been paid to how these models can be improved. In this paper, we propose a generalized framework for improving a human-like model while maintaining its human similarity, using a novel extension of curriculum learning which considers both the data selection, and labeling method. We measure performance as a trade-off between model strength, its human similarity, and the data efficiency of the curriculum.

We evaluate this framework in the domain of chess through a series of experiments. At a high level, we find that the choice of teacher is critical for maximizing the efficiency of strength gains. In many cases, the best teacher is *not* the strongest teacher. Notably, Stockfish (the world’s strongest chess engine) is a worse teacher than even a low-rated human-like model. Our results also suggest that the choice of curriculum is less important than the choice of teacher.

It is important to note that any insights we gain from training human-like models do not necessarily apply to training real humans. These are, after all, neural network models, and not human brains. Nevertheless, the results in Section 4.4 suggest that our curriculum training methodology is able to improve human-like models in a way that ensures human similarity at their improved level of play, on a path similar to human learning. This could provide insights into training real human students.

Teachers	Curricula							
	1200	1400	1600	1800	2000	2200	2400	2600
stockfish	42	43	44	44	44	44	45	44
2600m	44	46	47	48	49	49	50	49
2400m	44	47	48	49	50	50	51	49
2200m	45	47	49	49	50	50	50	49
2000m	47	49	50	50	50	50	49	47
1800m	48	50	50	51	50	49	48	46
1600m	48	50	50	50	49	48	47	45
1400m	49	50	50	49	48	47	46	44
random	44	47	48	49	50	50	51	49
blunder	44	46	48	49	50	50	50	49
rooks	44	46	47	48	49	49	50	48
1400i	43	45	46	47	49	49	50	48

Figure 5: The results of analyzing how similar curriculum-updated models are to human datasets at a range of skill levels.

References

- [1] AMIGÓ, E., GIMÉNEZ, J., GONZALO, J., AND MÀRQUEZ, L. Mt evaluation: Human-like vs. human acceptable. In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions* (2006), p. 17–24.
- [2] BENGIO, Y., LOURADOUR, J., COLLOBERT, R., AND WESTON, J. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning* (2009), p. 41–48.
- [3] BROWN, N., AND SANDHOLM, T. Superhuman AI for multiplayer poker. *Science* 365, 6456 (2019), 885–890.
- [4] CAMPBELL, M., HOANE JR, A. J., AND HSU, F.-H. Deep blue. *Artificial intelligence* 134, 1–2 (2002), 57–83.
- [5] GOBET, F., AND JANSEN, P. J. Training in chess: A scientific approach. *Education and chess* (2006).
- [6] HECKER, S., DAI, D., AND VAN GOOL, L. Learning accurate, comfortable and human-like driving. *arXiv preprint arXiv:1903.10995* (2019).
- [7] JACOB, A. P., WU, D. J., FARINA, G., LERER, A., HU, H., BAKHTIN, A., ANDREAS, J., AND BROWN, N. Modeling strong and human-like gameplay with kl-regularized search. In *International Conference on Machine Learning* (2022), PMLR, p. 9695–9728.
- [8] JENNEBOER, L., HERRANDO, C., AND CONSTANTINIDES, E. The impact of chatbots on customer loyalty: A systematic literature review. *Journal of theoretical and applied electronic commerce research* 17, 1 (2022), 212–229.
- [9] KIM, J., MERRILL JR., K., XU, K., AND KELLY, S. Perceived credibility of an ai instructor in online education: The role of social presence and voice features. *Computers in Human Behavior* 136 (Nov 2022), 107383.
- [10] KORNBLITH, S., SHLENS, J., AND LE, Q. V. Do better imagenet models transfer better? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2019), pp. 2661–2671.
- [11] LI, L., LASSITER, T., OH, J., AND LEE, M. K. Algorithmic hiring in practice: Recruiter and hr professional’s perspectives on AI use in hiring. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (2021), p. 166–176.
- [12] MCCARTHY, J. Chess as the drosophila of ai. In *Computers, chess, and cognition*. Springer, New York, NY, 1990, pp. 227–237.
- [13] MCILROY-YOUNG, R., KLEINBERG, J., SEN, S., BAROCAS, S., AND ANDERSON, A. Mimetic models: Ethical implications of AI that acts like you. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (2022), p. 479–490.
- [14] MCILROY-YOUNG, R., SEN, S., KLEINBERG, J., AND ANDERSON, A. Aligning superhuman ai with human behavior: Chess as a model system. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery Data Mining* (2020), p. 1677–1687.
- [15] MCILROY-YOUNG, R., WANG, R., SEN, S., KLEINBERG, J., AND ANDERSON, A. Learning models of individual behavior in chess. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2022), p. 1253–1263.
- [16] MCILROY-YOUNG, R., WANG, Y., SEN, S., KLEINBERG, J., AND ANDERSON, A. Detecting individual decision-making style: Exploring behavioral stylometry in chess. *Advances in Neural Information Processing Systems* 34 (2021), 24482–24497.
- [17] RAGOT, M., MARTIN, N., AND COJEAN, S. Ai-generated vs. human artworks. a perception bias towards artificial intelligence? In *Extended abstracts of the 2020 CHI conference on human factors in computing systems* (2020), p. 1–10.

- [18] SILVER, D., HUBERT, T., SCHRITTWIESER, J., ANTONOGLU, I., LAI, M., GUEZ, A., LANCTOT, M., SIFRE, L., KUMARAN, D., AND GRAEPEL, T. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science* 362, 6419 (2018), 1140–1144.
- [19] SONI, B., AND HINGSTON, P. Bots trained to play like a human are more fun.
- [20] SUN, Q., LIU, Y., CHUA, T.-S., AND SCHIELE, B. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 403–412.
- [21] VACI, N., EDELSBRUNNER, P., STERN, E., NEUBAUER, A., BILALIĆ, M., AND GRABNER, R. H. The joint influence of intelligence and practice on skill development throughout the life span. *Proceedings of the National Academy of Sciences* 116, 37 (2019), 18363–18369.
- [22] WANG, D., WEISZ, J. D., MULLER, M., RAM, P., GEYER, W., DUGAN, C., TAUSCZIK, Y., SAMULOWITZ, H., AND GRAY, A. Human-ai collaboration in data science: Exploring data scientists’ perceptions of automated ai. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.
- [23] YOSINSKI, J., CLUNE, J., BENGIO, Y., AND LIPSON, H. How transferable are features in deep neural networks? *Advances in neural information processing systems* 27 (2014).
- [24] ZHANG, R., MCNEESE, N. J., FREEMAN, G., AND MUSICK, G. “An ideal human” expectations of AI teammates in human-AI teaming. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW3 (2021), 1–25.