
A Comparison of Contextual and Non-Contextual Preference Ranking for Set Addition Problems

Timo Bertram¹ Johannes Fürnkranz¹ Martin Müller²

Abstract

In this paper, we study the problem of evaluating the addition of elements to a set. This problem is difficult, because it can, in the general case, not be reduced to unconditional preferences between the choices. Therefore, we model preferences based on the context of the decision. We discuss and compare two different Siamese network architectures for this task: a twin network that compares the two sets resulting after the addition, and a triplet network that models the contribution of each candidate to the existing set. We evaluate the two settings on a real-world task; learning human card preferences for deck building in the collectible card game “Magic: The Gathering”. We show that the triplet approach achieves a better result than the twin network and that both outperform previous results on this task.

1. Introduction

Set addition problems can be commonly found in many applications. The problem is to evaluate which of several possible candidates is the best addition to an existing set, such that the resulting set achieves a high evaluation according to a latent set evaluation function. Examples include adding cards to a player’s deck, adding players to a football team, or buying stocks to complement an existing portfolio. Naturally, the evaluation of such items depends on the elements that are already in the set. For example, a mediocre goalkeeper may be the better addition to a team of excellent field players than the best striker, and the evaluation of a promising high-risk share may depend on the risk profile of the stocks you are already holding.

In this paper, we look at this problem in a preference learning context (Fürnkranz & Hüllermeier, 2010): we assume

¹Dept. of Computer Science, Johannes-Kepler Universität, Linz, Austria ²Dept. of Computing Science, University of Alberta, Edmonton, Canada. Correspondence to: Timo Bertram <tbertram@faw.jku.at>.

that we are given training information that specifies which of two possible choices is preferential over the other. Based on this problem formulation, we study and compare two different learning schemes based on Siamese neural networks. The first is a classical preference learning setting, where the learner is trained to predict which of the two sets resulting from the addition is preferable. As a second variant, we consider a setting that can directly model which of the two additions fits the context better. As such, the former implicitly models the context of the decision by comparing the two resulting sets, while the latter models the context explicitly as a separate input to the learner.

We formally define contextual preferences and the set addition problem in Section 2, show two Siamese neural network approaches to this problem in Section 3, and evaluate and compare them for the task of deck-building in the collectible card game *Magic: The Gathering (MTG)*. To train and evaluate our method, we use a dataset of sequential expert deck-building decisions, which provides information about which selections the human experts preferred over others and allows us to compare the two preference-based methods on this data. These experiments and their results are presented in Sections 4 and 5, followed by some conclusions. This work builds open our previous research for the Contextual Preference Ranking framework (Bertram et al., 2021).

2. Contextual Preferences and the Set Addition Problem

As described above, we are concerned with problems where we have to evaluate the addition of an item to an existing set of items, such as adding players to a sports team, buying stocks for a portfolio, putting products in a shopping basket, or adding cards to a player’s deck. In the general case, this problem is hard, because the value of the items that are added changes depending on the set it is part of. In many cases, items may have different, hidden properties, and the evaluation of the set depends on properties that are covered by the individual items. While each item a value of its own, and some are better than others, the overall value of a set does not equal the sum of values of items. It crucially depends on the overall composition.

Formally, this *set addition problem* can be represented as follows: Given a set of items \mathcal{C} as the *context*, and a set of items \mathcal{P} that represent the current possible choices, select the item \mathbf{c}^* in \mathcal{P} which is the best addition to \mathcal{C} . Let us assume an (unknown) utility function $u(\cdot)$ which returns an evaluation for a given set of items, then

$$\mathbf{c}^* = \arg \max_{\mathbf{c} \in \mathcal{P}} u(\mathcal{C} \cup \{\mathbf{c}\}) \quad (1)$$

The learning problem is to learn the function $u(\cdot)$ from a set of example decisions. The training information can be given in various ways. In this paper, we assume a preference-based formulation: we do not have access to a direct evaluation of a set, but are given pairwise comparisons between them. In particular, we assume that we have access to a set of *contextual preferences* of the form

$$(\mathbf{c}_j \succ \mathbf{c}_k \mid \mathcal{C}) \quad (2)$$

which means that item \mathbf{c}_j is a better addition to the context \mathcal{C} than item \mathbf{c}_k . In our set-based setting, this is equivalent to unconditional preferences of the form

$$(\mathcal{C} \cup \{\mathbf{c}_j\}) \succ (\mathcal{C} \cup \{\mathbf{c}_k\}). \quad (3)$$

The main contribution of this paper is a case study that translates these two formulations into corresponding neural network architectures and compares them on learning human preferences in a real-world card game. For decisions without a context, when the first item is added to the set, $\mathcal{C} = \emptyset$. This special case of an empty context may be viewed as a comparison of the general, unconditional utility of two items \mathbf{c}_j and \mathbf{c}_k .

While we will not tackle this in the current paper, this framework can in principle easily be generalized to using arbitrarily large sets of items \mathcal{C}_j and \mathcal{C}_k for the context, i.e., for dealing with contextual preferences of the form

$$(\mathcal{C}_j \succ \mathcal{C}_k \mid \mathcal{C}_i). \quad (4)$$

3. Learning Contextual Preferences with Siamese Networks

In the following, we briefly describe how to use Siamese networks for preference learning from sets. While they are typically used on multiple examples of the same type, e.g. images, we employ them to allow comparisons of two items with a context by embedding both inputs as well as the context in a uniform representation space. To the best of our knowledge, this is a novel approach.

3.1. Siamese networks for preference learning

Siamese networks implement the idea that the same network is used multiple times in order to encode multiple items. The




Figure 1. Training scheme for cross-entropy loss using a positive \mathbf{p} and a negative example \mathbf{n} . The loss function aims to increase the evaluation of \mathbf{p} over the one of \mathbf{n} . N is the network that maps the two objects to an evaluation.

encodings are then compared and trained by a supervision signal (Chicco, 2021). A prototypical application of such networks is one-shot learning for image recognition (Vinyals et al., 2016; Koch et al., 2015). Going back to *comparison training* (Tesauro, 1989), similar symmetrical architectures have also been used in preference learning, where the task is not to encode the similarity of objects but the preference between them.

A more traditional neural network approach to preference learning would compare two items by having both as a concatenated input to one network, which then outputs a single signal to model a preference. One problem with this is that reordering the inputs can lead to different results. While this can in practice be combated by training the network with random orderings, there is no guarantee that this fully eliminates the error. An order-dependent output is problematic and should not occur in practice. Siamese networks circumvent this problem by processing multiple inputs sequentially by the same network. This leads to a separate output for each input, called the **embedding** of the input.

3.2. Unconditional Preferences

One way to model the preference of which item to add to a set is to model the preferences over the resulting unions. To use preferences of this type, one branch of the Siamese network encodes the preferred object \mathbf{p} , and the other branch the losing object \mathbf{n} . The two encodings are then compared by using their difference, as in Figure 1.

For sets, we define $\mathbf{p} = \mathcal{C} \cup \{\mathbf{c}_j\}$ and $\mathbf{n} = \mathcal{C} \cup \{\mathbf{c}_k\}$, as shown in (3). The output of the network is a single real value, which can be regarded as an evaluation $u(\mathcal{C} \cup \{\mathbf{c}_j\})$ of the set \mathbf{p} . The preference of one set over another is modeled by a higher evaluation $u(\cdot)$.

This setting corresponds to comparison training, which has




Figure 2. Training scheme for triplet loss using an anchor \mathbf{a} , a positive \mathbf{p} and a negative example \mathbf{n} . N is the network that maps each of the three objects into an embedding space. The loss function indicates whether \mathbf{a} is closer to \mathbf{p} or to \mathbf{n} .

been proposed by Tesauro (1989) in a game-playing context. For comparisons between two arbitrary items, the RankNet approach (Burges et al., 2005) uses a cross-entropy loss function and the sigmoid of the difference between the two evaluations. We directly follow this method in our first set-based approach and will refer to it as *RankNet*.

3.3. Triplet Siamese Networks for Contextual Preference Ranking

For directly using the contextual preferences of the form (2), we employ triplet Siamese networks, as shown in Figure 4. The key idea of this approach is to use an anchor (\mathbf{a}), a positive (\mathbf{p}) and a negative (\mathbf{n}) example. The anchor models the context, and the positive example \mathbf{p} is preferred to the negative example \mathbf{n} in this context, ($\mathbf{p} \succ \mathbf{n} \mid \mathbf{a}$).

Such networks are trained with a triplet loss

$$L_{\text{triplet}}(\mathbf{a}, \mathbf{p}, \mathbf{n}) = \max(d(\mathbf{a}, \mathbf{p}) - d(\mathbf{a}, \mathbf{n}) + m, 0). \quad (5)$$

The loss decreases with decreasing distance between \mathbf{a} and \mathbf{p} , and with increasing distance between \mathbf{a} and \mathbf{n} . This moves the embeddings of the anchor and the positive example closer together, and pushes the embedding between the anchor and the negative example apart. For this work, the Euclidian distance $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2$ is chosen as the distance metric between embeddings. The *margin* m is a parameter of the loss function and controls how far embeddings are pushed away from each other. We used a margin of $m = 1$. In preliminary experiments, the exact value of this parameter was not critical for the performance of the method. As an example, Siamese architectures can compare pictures of individuals and be trained to recognize whether two different images show the same person. In that case, the preference indicates which picture is more likely to show the same individual as the anchor.

We use them here in a slightly different, set-based setting. In our case, the anchor object \mathbf{a} is the context set \mathcal{C} , which needs to be extended with one of two candidate extensions $\mathbf{p} = \mathbf{c}_j$ or $\mathbf{n} = \mathbf{c}_k$. The training information indicates that \mathbf{p}

is a better extension than \mathbf{n} . This is very different to asking whether \mathbf{a} is more similar to \mathbf{p} or \mathbf{n} . For example, card selection tasks seek cards that complement the deck, rather than duplicating the effect of similar cards picked earlier.

3.4. Discussion

At testing time, we do not need to query all possible pairwise comparisons of options, but can directly evaluate each option to formulate an overall ranking. In the case of Contextual Preference Ranking, this becomes possible because the resulting preferences are transitive w.r.t. to the given anchor set, i.e.,

$$(\mathbf{c}_1 \succ \mathbf{c}_2 \mid \mathcal{C}) \wedge (\mathbf{c}_2 \succ \mathbf{c}_3 \mid \mathcal{C}) \Rightarrow (\mathbf{c}_1 \succ \mathbf{c}_3 \mid \mathcal{C})$$

The reason for this is that all objects are embedded with the same embedding network N , which always outputs the same signal for the same input, regardless of the position of the item in the comparison. The same principle applies to unconditional preferences.

We view the adaptation of triplet Siamese networks to a set-based Contextual Preference Ranking framework as the main contribution of this work, as it introduces a new way of thinking about the Siamese triplet structure. Instead of comparing similar items, we train a preference of items based on a context. To our knowledge, we are the first to use triplet Siamese networks in such a way (Bertram et al., 2021). This contextual preference of comparing \mathbf{p} and \mathbf{n} with context \mathbf{a} also differs from trying to model the unconditional preferences $\mathbf{a} + \mathbf{p}$ and $\mathbf{a} + \mathbf{n}$. We want to emphasize the generality of this framework; it is applicable to model any kind of preference learning problem with a context.

4. Experimental setup

The goal of our experimental evaluation is to compare the two different solutions for contextual preference problems described in Sections 3.2 and 3.3. As a domain, we choose the problem of *drafting*, or selecting cards, in the collectible card game *Magic: The Gathering (MTG)*. We define the context \mathcal{C} as the set of previously chosen cards of a player and train the networks with pairs of cards \mathbf{p} and \mathbf{n} , where \mathbf{p} was chosen by the player and \mathbf{n} is another card that was available but not chosen. For Contextual Preference Ranking, we model that in the human expert’s opinion, \mathbf{p} fits better into the current set \mathcal{C} than \mathbf{n} . For RankNet, we model that the set $\mathcal{C} \cup \{\mathbf{p}\}$ should receive a higher evaluation than $\mathcal{C} \cup \{\mathbf{n}\}$. Both approaches rank choices; CPR by distance to the anchor and RankNet by the evaluation of resulting sets.

In the following, we briefly describe the game setting, the dataset, and the used network architectures.

4.1. Drafting in Magic: The Gathering

Collectible card games have been around for decades and are among the most played tabletop games. However, they are also among the most complex games (Hoover et al., 2020). Of course, a good player needs to be able to play the game itself, which requires an understanding and knowledge of thousands of cards. Furthermore, *deck-building*, choosing a suitable set of cards to play with, is a gigantic challenge in itself and is vastly beyond the power of exhaustive computation. We abstain from explaining the complex rules (Wizards of the Coast, 2021), as they are not necessary to understand the contribution of this work. Instead, we provide some background information about the way cards are chosen in the used dataset.

MTG is played in a variety of different styles. For this work, we consider the format of *drafting* in a game with eight players. In contrast to formats where decks are constructed separately from playing, drafting features a first game phase in which players form their decks from a selection of cards, so-called *packs*. Over the course of the whole draft, each player chooses a pool of 45 cards sequentially, from which their deck is built afterward. Players get their cards by choosing from many packs as follows: Each of the eight players in a draft starts with a full pack of 15 cards, selects a single card from it, and passes the remaining 14 cards on to the next player. In the following rounds, players select from 14, 13, . . . cards until the packs are emptied. This process is repeated two more times with new decks, such that in the end, each player has selected $3 \times 15 = 45$ cards in total.

4.2. Data preparation and exploration

The *DraftSim* dataset used in this research has been collected by Ward et al. (2020) and contains 107,949 human drafts simulated on the Web.¹ Each draft consists of 24 packs of 15 cards distributed as explained in Section 4.1. The dataset includes 2,590,776 separate packs, which consist of a total of 265 different cards.

We train the network on pairs of possible cards in the context of the set of cards that are already held by the player. For each decision to choose the best card from a pack of k cards, $k - 1$ training examples are generated, for pairing the human-selected card with each of the $k - 1$ other cards in the pack. The *DraftSim* dataset contains 217,624,680 such training examples. These examples are split 80/20 into training and test data, using the same split as in (Ward et al., 2020) to allow a direct comparison.

To better understand the characteristics of the *DraftSim* dataset, we defined two metrics:

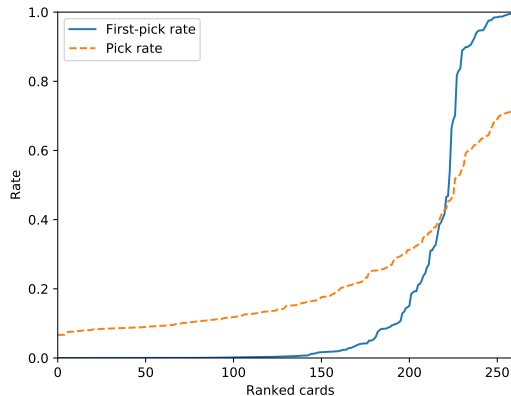


Figure 3. Pick rate per card

- (i) The *pick rate* of each individual card c captures how often the card was selected when being offered.

$$p_{\text{pick}}(c) = \frac{\text{number of times } c \text{ chosen}}{\text{number of times } c \text{ offered}}$$

- (ii) The *first-pick rate* captures how often a card c was selected on the very first pick.

$$p_{\text{firstPick}}(c) = \frac{\text{number of times } c \text{ chosen first}}{\text{number of times } c \text{ offered first}}$$

The former metric defines how likely a card is to be chosen over the whole range of the draft, while the second only considers the very first pick. Whether a card is selected first depends mainly on its individual card strength. In contrast, later card choices are heavily influenced by previously selected cards.

Figure 3 demonstrates that recognizing the first pick is a much easier task than choosing cards later since the consensus is higher at that point. For the first pick decision, it is possible to simply consult a ranking of available cards (Karsten, 2018; Scott-Vargas, 2018; Troha, 2018). However, even for this seemingly simple task, rankings are rarely unanimous, which underlines the complexity of the domain.

Over the whole draft, all cards will be chosen at some point. For the first pick, the number of reasonable choices is relatively small, as can be seen from the quick drop of the blue solid line in Figure 3. In addition, the **lowest** observed pick rate in the *DraftSim* set is 0.07, which is close to the theoretical minimum of $1/15 \approx 0.0667$. However, the lowest first-pick rate in the data set is 0.00001, which can safely be regarded as a misclick or otherwise unexplainable decision. In contrast, the two **highest** pick-rates 0.98 and 0.77. The best card is colorless and therefore playable in any deck. However, the second-best card is white, which explains why a portion of decisions did not choose that card, as the player was likely already firmly drafting a deck that did not include white cards. This again confirms the importance of context.

¹<https://draftsim.com/draft-data/>

4.3. Network Architecture

This section shows details of the architecture and training method for the Siamese networks used in the experiments. The Siamese network encodes a set of input cards through multiple fully-connected network layers (Figure 4). Therefore, each training update consists of two or three sequential forward passes through the network, followed by the computation of the loss and a backward pass for updating the network parameters. The way this network is used for the twin network (RankNet) and the triplet network (CPR) is shown in Figures 1 and 2 respectively.

The network takes a set of cards as input. The input space is 265-dimensional, with one dimension for each possible card. For \mathbf{p} and \mathbf{n} , the input is a one-hot encoding, while the anchor \mathbf{a} uses an encoding in which each dimension encodes the number of already chosen cards of each type. The output of the network is a D -dimensional vector of real numbers in the range $[-1, 1]$, where $D \geq 1$ is a hyperparameter ($D = 1$ for RankNet). This output vector is the learned embedding of the input set.

Fully-connected layers are linked by exponential linear unit functions (ELU) (Clevert et al., 2016). In preliminary experiments, this led to quicker training than rectified linear (RELU) and leaky RELU activations. We use a learning rate of 0.0001 and the Adam optimizer with a batch size of 128. For the output layer, the tanh function was chosen. Batch normalization was not used as it did not seem to help, but we do use a dropout of 0.5. Most of the parameters, such as the learning rate, the size of the network, and the optimizer, were not optimized, as reaching the absolute highest performance was not the priority of this work. Rather, we used intuitive parameters, which were comparable to the ones used in previous research (Ward et al., 2020).




Figure 4. Siamese network N architecture

5. Results

In this section, we discuss the performance of our networks for the card selection task and visualize the obtained card embeddings.

5.1. Card Selection Accuracy

Our primary goal was to compare Contextual Preference Ranking with RankNet and with previous methods for this task (Ward et al., 2020). The best performing algorithm of that study used a traditional, single-branch deep neural network to learn a ranking over all possible cards for a given context. It was trained by directly mapping a feature-based encoding of the current set of cards \mathcal{C} to a one-hot encoded vector that represents the selected card. Thus, it generated exactly one training example per card pick. Our two agents, CPRBOT and RANKNETBOT, instead learn on pairwise comparisons between the picked card and any other card in the candidate pack \mathcal{P} , and therefore generate 2 to 14 training examples from a single pick. However, this additional constant factor in the training complexity is to some extent compensated by the fact that we were able to train our networks with a much smaller amount of training epochs.

Following Ward et al. (2020), Table 1 reports two measures: the *mean testing top-one accuracy (MTTA)* is the percentage of cases in which the network chooses the correct card in the pack. The *mean testing pick distance (MTPD)* shows how far away the correct pick is from the chosen card when ranking all possible choices. For CPRBOT, we report the performance for two different agents, which only differ in the output dimension of the neural network. RANKNETBOT is able to achieve higher accuracies than the previously proposed agents, but the 256-dimensional CPRBOT was able to achieve the best performance by a large margin. This strong increase suggests that the Contextual Preference Ranking approach works well for this domain, and outperforms the direct comparison between two sets by RankNet.

Table 1. Performance of proposed agent to previously seen.

¹ Heuristic agents (Ward et al., 2020) ² Trained agents (Ward et al., 2020) ³ This work

Agent	MTTA (%)	MTPD
RandomBot ¹	22.15	NA
RaredraftBot ¹	30.53	2.62
DraftsimBot ¹	44.54	1.62
BayesBot ²	43.35	1.74
NNetBot ²	48.67	1.48
RankNetBot ³	69.09	0.80
CPRBot ³ , D=2	53.69	0.98
CPRBot ³ , D=256	83.78	0.25

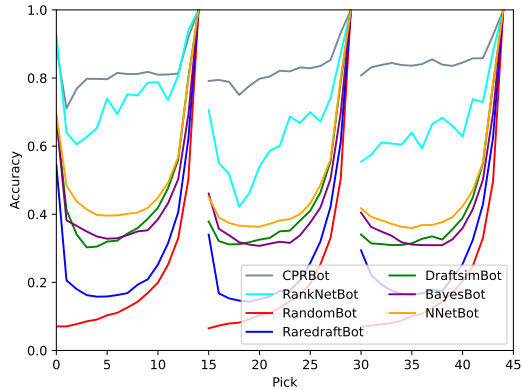


Figure 5. Accuracy of pick-prediction for size of context. Separate lines model the three separate packs.

5.2. Draft Analysis

We compare the performance of both methods over the course of the whole draft. Since already chosen cards strongly influence the current decision, we explore whether a growing set of chosen cards influences the accuracy of picks. Figure 5 shows the accuracy of all agents over the three consecutive picking rounds. Clearly, both preference-based algorithms achieve higher performance than the others. Interestingly, the accuracy of picks does not show the same performance curve for CPRBOT as for the other methods. Those methods have U-shaped curves and fail to make good decisions in the middle of packs. CPRBOT remains stable throughout the picks while increasing in the end due to a smaller number of choices. CPRBOT’s performance has an outlier at the second pick. This may be because there is no difference between the embeddings of the anchors and the embedding of the card choice. For the second pick only, the anchor set is modeled the same as a single card.

5.3. Discussion

Both preference-based agents achieved higher accuracy than previously reported. This is especially well pronounced in the middle of the pack, where weaker cards have to be compared against each other. To further visualize correlations between the network predictions and the underlying data, we plot the first-pick rate of cards against the distance to the empty set in Figure 6, showing a strong correlation with a Kendall rank correlation coefficient of 0.74. The main difference between these two statistics is that the distance is much smoother than the first-pick rate, which decreases rapidly for weaker cards. The first-pick rate is only subject to binary choices, i.e., $c_1 \succ c_2$ without giving any weight to how close the decision between those cards was. Due to the training with decisions between mediocre cards, the embedding distance is a smoother measure of how strong the card is according to the network.




Figure 6. Correlation between FPR and distance to empty anchor. Kendall’s tau = 0.74.

6. Conclusions and Future Work

We showed that the Siamese-based agents which model preferences worked well in the context of drafting cards in *Magic: The Gathering* and vastly outperformed previous results. Compared to (Ward et al., 2020), we report an increase in accuracy by more than 56%, while also decreasing the pick distance by more than 83% with our CPRBOT. Even when the Siamese network makes an incorrect choice, it typically ranks the correct choice very high. In addition, we also showed that the CPR-based triplet network architecture outperformed a more conventional twin network. We, therefore, speculate that the former is better suited for modeling set-addition problems with Contextual Preference Ranking.

We want to reemphasize that while these are early tests, there is no reason to believe that the success is limited to this particular setting. We did not incorporate any domain information beyond the ID of cards used to encode the input into our networks. Therefore, we expect that our proposed framework may work well for other problems where preference has to be modeled in a context. In order to further test the generality of this approach in other domains, more work with other datasets is required. We could also aim to evaluate whole sets of extensions at once instead of only single items. In addition, there is potential to use this method not only for pre-game decision-making but also for game playing for games that can be represented as sets.

One concern with our work so far is that we have only trained on human expert examples, which limits performance in a general context. For domains where an automatic evaluation of the chosen sets is available, we could generate datasets using self-play as part of an agent training loop. This is the approach used in Tesauro’s TDGammon (Tesauro, 1995), in DeepMind’s AlphaZero (Silver et al., 2018), and in many similar approaches.

References

- Bertram, T., Johannes, F., and Martin, M. Predicting human card selection in Magic: The Gathering with contextual preference ranking. In *Proceedings of the 3rd IEEE Conference on Games*, 2021.
- Burges, C. J. C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. N. Learning to rank using gradient descent. In Raedt, L. D. and Wrobel, S. (eds.), *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, pp. 89–96, Bonn, Germany, 2005. ACM. doi: 10.1145/1102351.1102363.
- Chicco, D. Siamese neural networks: An overview. In Cartwright, H. M. (ed.), *Artificial Neural Networks*, volume 2190 of *Methods in Molecular Biology*, pp. 73–94. Springer, 3rd edition, 2021. doi: 10.1007/978-1-0716-0826-5_3.
- Clevert, D., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (ELUs). In Bengio, Y. and LeCun, Y. (eds.), *Proceedings of the 4th International Conference on Learning Representations (ICLR)*, San Juan, Puerto Rico, 2016.
- Fürnkranz, J. and Hüllermeier, E. (eds.). *Preference Learning*. Springer-Verlag, 2010. ISBN 978-3642141249.
- Hoover, A. K., Togelius, J., Lee, S., and de Mesentier Silva, F. The many AI challenges of hearthstone. *Künstliche Intelligenz*, 34(1):33–43, 2020. doi: 10.1007/s13218-019-00615-z.
- Karsten, F. An early pick order list for core set 2019, 2018. URL <https://strategy.channelfireball.com/all-strategy/mtg/channelmagic-articles/an-early-pick-order-list-for-core-set-2019/>. Retrieved April 02, 2021.
- Koch, G., Zemel, R., and Salakhutdinov, R. Siamese neural networks for one-shot image recognition. In *Proceedings of the ICML’15 Deep Learning Workshop*, 2015.
- Scott-Vargas, L. Core set 2019 limited set review: White, 2018. URL <https://strategy.channelfireball.com/all-strategy/mtg/channelmagic-articles/core-set-2019-limited-set-review-white/>. Retrieved April 02, 2021.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- Tesauro, G. Connectionist learning of expert preferences by comparison training. In Touretzky, D. (ed.), *Advances in Neural Information Processing Systems 1 (NIPS-88)*, pp. 99–106. Morgan Kaufmann, 1989.
- Tesauro, G. Temporal difference learning and TD-Gammon. *Communications of the ACM*, 38(3):58–68, March 1995.
- Troha, D. Draftsim’s pick order list for core set 2019, 2018. URL <https://draftsim.com/M19-pick-order.php>. Retrieved April 02, 2021.
- Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., and Wierstra, D. Matching networks for one shot learning. In Lee, D. D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29: Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, pp. 3630–3638, Barcelona, Spain, 2016.
- Ward, H. N., Brooks, D. J., Troha, D., Mills, B., and Khakhalin, A. S. AI solutions for drafting in magic: the gathering. *arXiv preprint*, abs/2009.00655, 2020.
- Wizards of the Coast. Magic: The Gathering – comprehensive rules, 2021. URL <https://media.wizards.com/2021/downloads/MagicCompRules%2020210224.pdf>. Retrieved April 07, 2021.