

Predicting Human Card Selection in *Magic: The Gathering* with Contextual Preference Ranking

Timo Bertram
Dept. of Computer Science
Johannes-Kepler Universität
Linz, Austria
tbertram@faw.jku.at

Johannes Fürnkranz
Dept. of Computer Science
Johannes-Kepler Universität
Linz, Austria
juffi@faw.jku.at

Martin Müller
Dept. of Computing Science
University of Alberta
Edmonton, Canada
mmueller@ualberta.ca

Abstract—Drafting, i.e., the iterative, adversarial selection of a subset of items from a larger candidate set, is a key element of many games and related problems. It encompasses team formation in sports or e-sports, as well as deck selection in formats of many modern card games. The key difficulty of drafting is that it is typically not sufficient to simply evaluate each item in a vacuum and to select the best items. The evaluation of an item depends on the context of the set of items that were already selected earlier, as the value of a set is not just the sum of the values of its members - it must include a notion of how well items go together.

In this paper, we study drafting in the context of the card game *Magic: The Gathering*. We propose the use of the Contextual Preference Ranking framework, which learns to compare two possible extensions of a given deck of cards. We demonstrate that the resulting neural network is better able to better inform decisions in this game than previous attempts.

Index Terms—Preference Learning, Game-playing, Siamese Networks, Card Games, *Magic: The Gathering*

I. INTRODUCTION

Collectible card games have been around for decades and are among the most played tabletop games in existence. However, they are also among the most complex games [7]. 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. For the game of *Magic: The Gathering*, a lower boundary of the number of possible card configurations can be computed as follows. For one of their most restricted game modes, *Standard*, 1983 different cards are currently legal. Decks consist of at least 60 cards, of which usually about 37 are chosen from the aforementioned pool, which we will use as our lower bound. As each card can be put into a deck up to four times, this leads to $\binom{1983 \times 4}{37} > 10^{101}$ combinations of cards. Even with the assumption that a player will play a deck black and blue deck, which is a more reasonable assumption, this still results in $\binom{847 \times 4}{37} > 10^{87}$ possible decks.

As such numbers are vastly beyond the power of exhaustive computation, other methods must be developed to train agents to build decks. In this work, we study a specific game mode of MTG where the deckbuilding process is greatly simplified. To

train and evaluate our method, we use a dataset of expert draft selections, which provides information about which selections human experts preferred over others.

Our main technical contribution is to use Siamese networks in a way that has not been previously used. We train and employ them to decide between different choices by explicitly modeling the context of the decision. The general framework of this method of Contextual Preference Ranking is developed in Section V. Before that, we start with a brief description of the game (Section II) and a review of related work and Siamese networks (Sections III and IV). Our experiments and their results are presented in Sections VI and VII, followed by a discussion, our conclusions, and an outlook on open questions for future work.

II. MAGIC: THE GATHERING

Magic: The Gathering (MTG) is a collectible card game with several million players worldwide. We abstain from explaining the complex rules [19], as they are not necessary to understand the contribution of this work, but provide some background information in order to introduce the terminology used.

A. Drafting

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 a pool of cards, from which they later build their decks. The pool of cards is chosen from semi-random selections of cards, so-called *packs*. Each pack of MTG cards consists of 15 different cards of four different rarities: eleven *Common*, three *Uncommon*, and one *Rare* or *Mythic* card.¹ *Rare* cards appear more frequently than *Mythic* ones. Over the course of the whole draft, each player chooses a deck of 45 cards sequentially. 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. This

¹Some packs contain an extra sixteenth card. However, such packs did not occur in our dataset.

process continues until all 15 cards of the original packs are chosen. This process is repeated for an additional two packs, such that each player selects 45 cards in total. In each round, packs are passed around in the same direction.

Drafting differs from free deckbuilding since players can not choose any existing card. Still, the computational complexity of this problem is enormous, as a single draft leads to $(15 \cdot 14 \cdot 13 \dots)^3 > 2 \times 10^{36}$ possible decks for each individual player. As there are 8 players but 15 cards per pack, players will see most packs of cards twice. This gives players additional information, such as which cards have been selected by the opponents in the last round. Such information is disregarded in our current work. We evaluate each pick only in the context of our current selection of cards, without taking information about the opponents’ possible picks, future or past, into consideration.

B. Card Colors

An important property of a card is its *card colors*, which has a major impact on the composition of a good deck. Most cards are assigned a single color, but some cards have multiple colors, and a small subset of cards has no colors. While there are exceptions to this, most players’ final decks will only use cards of two different colors. This means that previously selected cards, especially their color, strongly influence the selection of subsequent cards, in order to build a consistent deck. This also means that colorless cards can be valuable, as they can be used in any deck. On the other hand, using a multicolored card however requires having *all* of its colors in the deck, which makes them much harder to incorporate.

C. Deck Building

Only about 23 of the 45 selected cards will be used in a player’s final deck, so almost half of all chosen cards do not participate in the play phase.² This opens strategic options such as making speculative picks and changing colors during the drafting phase. Players may also pick some strong cards without intending to play them, in order to deny the other players that card.

In actual games, the draft phase is followed by the play phase. In this work, instead of evaluating a drafting strategy directly by playing games with the resulting deck, we evaluate it by using a large database of human expert picks. This serves as the ground truth for which card is best in a specific situation. This has its limitations, as human choices are far from perfect and can be inconsistent, as well as having no information about the performance of the resulting decks. Still, this dataset is useful when trying to predict human decision-making in this context, and allows the study of draft picking independently of card play.

²The reason is that a legal deck only requires 40 cards, of which usually 17 are so-called *basic lands*, which are not part of the drafting process. For more information about this, visit <https://magic.wizards.com/en/articles/archive/lo/basics-mana-2014-08-18>.

III. RELATED WORK

Current work on selecting cards in the setting of collectible card games is limited due to available data. Most existing approaches either drastically reduce the complexity of the domain by choosing subsets of cards [1] or by using naive versions of games [10]. Evolutionary approaches are often used for deck building. However, computing the fitness of a deck is a difficult problem by itself. In practice, those approaches often use naive game-specific heuristics to play games [1], [6], which do not transfer to the context of real gameplay. A different way to circumvent the complexity of evaluating decks, which we follow here, is by training on expert decisions. *DraftSim* [18] is a large public domain simulator for human deck-building decisions, which provides an excellent basis for training. This dataset uses the eight-player drafting setting explained in Section II-A. In their work, they also proposed several card selection methods [18]. Their best performing method is a deep neural network, which was trained to directly pick the best-fitting card in each round. The input of the network is a feature-based encoding of the current set of cards, while the output is a vector of real-valued scores, which rank all possible card choices. A card with the maximum score within the current selection \mathcal{P} is chosen.

In our work, we train a Siamese network [3], [9] for the task of drafting. These networks are often used in one-shot learning for image recognition [9], [17] and process multiple inputs sequentially in the same network (see Section IV). Siamese networks have also been used in preference learning to compare two examples of a similar item [2]. This idea can also be viewed as an extension of Tesauro’s comparison training networks [14]. In his work, networks use pairwise comparisons without context, while we add the anchoring context sets. In our work, we use Siamese networks differently: we compare two items with a context by embedding both inputs as well as the context in a representation space with the help of the network. To the best of our knowledge, this is a novel approach.

IV. SIAMESE NETWORKS FOR PREFERENCE LEARNING

A key advantage of Siamese architectures over other, more traditional, neural networks is their independence of the order of inputs. Feeding both choices as one input into the network can lead to different outputs depending on the order, which is circumvented by having a separate forward-pass through the network for each input. The output for a given input is called the **embedding** of the input (see Figure 1).

To compare the different embeddings, Siamese networks often employ the distance between them to model similarities and preferences. The contrastive loss [11] and triplet loss are common loss functions.

$$L_{\text{triplet}}(\mathbf{a}, \mathbf{p}, \mathbf{n}) = \max(d(\mathbf{a}, \mathbf{p}) - d(\mathbf{a}, \mathbf{n}) + m, 0) \quad (1)$$

The *triplet loss* (Equation 1) uses an anchor (\mathbf{a}), a positive (\mathbf{p}) and a negative (\mathbf{n}) example. The anchor models the item to compare to, while the positive example \mathbf{p} is in some manner preferential to the negative example \mathbf{n} . As the loss decreases




Fig. 1. Training scheme for triplet loss using an anchor \mathbf{a} , a positive \mathbf{p} and a negative example \mathbf{n} . The loss function indicates whether \mathbf{a} is closer to \mathbf{p} or to \mathbf{n} . N is the network that maps an item into an embedding space.

with decreasing distance between \mathbf{a} and \mathbf{p} , and with increasing distance between \mathbf{a} and \mathbf{n} , this means that preferential items are embedded at closer positions in the embedding space than less preferential ones. While this choice is arbitrary, the Euclidian distance $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2$ is chosen as the distance metric for this work. 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 1. In preliminary experiments, the exact value of this parameter was not critical for the performance of the method.

Siamese networks are often used to model the similarity of items. For 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 of p and n is more likely to show the same individual as the anchor, therefore modeling similarity between items.

V. THE CONTEXTUAL PREFERENCE RANKING FRAMEWORK FOR SET ADDITION PROBLEMS

We use Siamese networks differently: instead of item similarity, we model preferences in a contextual, set-based setting, where \mathbf{p} and \mathbf{n} are possible additions to an existing anchor set \mathbf{a} . Formally, this *set addition problem* can be represented as follows: Given a set of items \mathcal{C} modeling the context, and a set of items \mathcal{P} that represent the current possible choices, select the item \mathbf{c}^* in \mathcal{P} , which fits the set \mathcal{C} best. Formally, if $u(\cdot)$ is an (unknown) utility function that returns an evaluation of a given set of items, then

$$\mathbf{c}^* = \arg \max_{\mathbf{c} \in \mathcal{P}} u(\mathcal{C} \cup \{\mathbf{c}\}) \quad (2)$$

The learning problem now is to learn the function $u(\cdot)$ from a set of example decisions. We propose to solve this problem by learning contextual preferences of the form

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

which means that item \mathbf{c}_j is a better addition to set \mathcal{C} than \mathbf{c}_k . In our application to drafting, all preferences are defined over one-element extensions $\{\mathbf{c}_i\}$. However, in principle, this framework can also be applied if the set \mathcal{C} can be extended by arbitrary larger sets of items \mathcal{C}_j and \mathcal{C}_k . For decisions without

a context, such as the first pick in a draft, $\mathcal{C} = \emptyset$. The distance to the empty set \emptyset can be used as a measure of the general utility of a card.

For training a network with such contextual preference decisions, we employ Siamese networks trained with the triplet loss. While such networks have been previously used for comparing the similarity of items (“anchor object \mathbf{a} is more similar to object \mathbf{p} than to object \mathbf{n} ”), we use them here in a slightly different setting. The anchor object \mathbf{a} is a set which needs to be extended with one of two candidate extensions \mathbf{p} or \mathbf{n} . The training information indicates that \mathbf{p} is a better extension than \mathbf{n} . This is very different from asking whether \mathbf{a} is more similar to \mathbf{p} or \mathbf{n} . For example in card drafting, we seek complementary cards that add to a deck, rather than endlessly duplicating the effect of similar cards picked earlier.

At testing time, pairwise comparisons are not needed, as we can directly evaluate each option in the context of their common anchor. This is possible because the resulting preferences are transitive w.r.t. 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.

This Contextual Preference Learning framework is the main contribution of this work, as it introduces a new way of thinking about the Siamese structure. Instead of comparing similar items, we train a preference of items based on a context. To our knowledge, Siamese networks have not previously been used in such a way. In addition, this contextual preference of comparing \mathbf{p} and \mathbf{n} with context \mathbf{a} also differs from comparing $\mathbf{a} + \mathbf{p}$ and $\mathbf{a} + \mathbf{n}$ as in *RankNet* [2]. We want to emphasize the generality of this framework - it is applicable to model any kind of preference learning problem with a context.

VI. EXPERIMENTAL SETUP

In this section, we evaluate the framework defined above in the domain of drafting cards in MTG.³ We define the context \mathcal{C} as the set of cards previously chosen by a player and train the networks with pairs of possible card choices p and n , where p was chosen by the player and n is another card that was available but not chosen. Therefore, we model that in the human expert’s opinion, p fits better into the current set than n . When using the network to make a pick decision in a game where we already hold cards \mathcal{C} , we compute the embedding $N(\mathcal{C})$ and the embeddings $N(\mathbf{c}_i)$ for all possible card choices \mathbf{c}_i , then choose the card \mathbf{c}^* with minimal distance to \mathcal{C} . Due to the nature of how we structure the training in Contextual Preference Ranking, it is important to emphasize that this distance does not model that this card is most similar to \mathcal{C} . Rather, the distance models how well cards fit into the context, with smaller distances equaling a better fit.

³The code used for all experiments can be found at <https://github.com/Tibert97/Predicting-Human-Card-Selection-in-Magic-The-Gathering-with-Contextual-Preference-Ranking>

A. Data preparation and exploration

The *DraftSim* dataset used in this research consists of 107,949 human drafts from the associated website [5]. Each draft consists of 24 packs of 15 cards distributed as explained in Section II-A. The dataset includes 2,590,776 separate packs and it contains a total of 265 different cards. It is important to note that those decisions are obtained from a simulator specifically created for drafting. Therefore, the dataset does not contain the playing phase of the game. In addition, the dataset is not tied to a larger *Magic: The Gathering* environment, which means that cards are not associated with a market where cards can be bought or sold. This is important, as otherwise, the physical or digital price of cards may influence decisions.

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 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 [18] to allow a direct comparison.

In order to better understand the characteristics of the dataset, we defined two metrics:

- (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 mainly depends on its individual card strength. In contrast, later card choices are heavily influenced by previously selected cards. In practice, players strongly prefer cards that match their collected colors, as those are most likely to be included in the final deck.

Figure 2 demonstrates that recognizing the first picked card is a much easier task than choosing cards later, since the human player's consensus is higher at that point. For the first pick decision, it is possible to simply consult a ranking of available cards [8], [12], [15]. However, even for this seemingly simple task, rankings are rarely completely 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. Therefore, the first-pick rate drops drastically as cards get weaker, as can be seen from the quick drop of the blue solid line in Figure 2. The lowest observed pick rate in the *DraftSim* set is 0.07, which is close to the theoretical minimum of $1/15 \approx 0.0667$ when a card is always chosen last. However,




Fig. 2. Pick rate of each individual card (higher pick rate equals better card)




Fig. 3. First-pick rate per rarity. Larger area of plot equals more cards of that pack rate. The cards with the highest first-pick rates are all *Rare* and *Mythic*.

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 first-pick rates are 0.9995 and 0.9987, showing clearly that in a vacuum, some cards are clearly regarded as the strongest. The pick-rate differs, as in those decisions the context of already chosen cards is important. There, the two highest pick rates are 0.98 and 0.77. This steep decline occurs, as the card with the highest rate is a colorless one and therefore playable in any deck. The second best however is a white card, which explains why a portion of decisions did not choose that card, as the player was likely already firmly drafting a deck of different colors than white.

Due to the properties of the game, a drafting system for this dataset does not need to be able to compare every card with all others, as a single pack of cards never includes multiple *Rare* or *Mythic* cards (Section II-A). Figure 3 visualizes the density of first-pick rates of cards separated by rarity. There, it is visible that all the very strongest cards in the set are in those two groups. For example, even one of the lesser picked *Rare* cards in the top cluster, with a first-pick rate of about 0.8, is picked more often than any *Common* or *Uncommon* card.




Fig. 4. Siamese network N architecture

However, just choosing *Rares* and *Mythics* whenever possible does not result in an appropriate heuristic. As also seen in Figure 3, a large number of these cards are also among the least-picked and therefore weakest cards. This is a result of *MTG* having multiple formats. Such cards can be strong within the context of very specific pre-constructed decks but are close to useless in the drafting format.

B. Network Architecture

This section outlines details of the architecture and training method used for the Siamese network in our experiments. The three different inputs \mathbf{a} (corresponding to the anchor card set \mathcal{C}), and \mathbf{p} and \mathbf{n} (corresponding to the picked card and one of the other cards) are sequentially processed by the network, as shown in Figure 1. Each forward-pass through the network encodes a set of input cards through multiple fully-connected network layers (Figure 4). Therefore, each training update consists of three sequential forward passes through the network, followed by the computation of the loss and a backward pass for updating the network parameters.

The embedding network takes a set of cards as input. The input space is 265-dimensional with one dimension representing 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 that 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 parameter, which we experimentally evaluate in Section VII-D. The output vector is the learned embedding of the input set.

Fully-connected layers are linked by exponential linear unit functions (ELU) [4]. 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. We do not use batch normalization as it did not provide benefits in our experiments but we use a dropout of 0.5. Most of those 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 [18]. Performance can likely be enhanced further with a guided search for the optimal parameters.

TABLE I
PERFORMANCE OF PROPOSED AGENT TO PREVIOUSLY SEEN:
¹ HEURISTIC AGENTS [18] ² TRAINED AGENTS [18] 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
SiameseBot ³ , D=2	53.69	0.98
SiameseBot ³ , D=256	83.78	0.2476

VII. RESULTS

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

A. Card Selection Accuracy

Our primary goal was to compare our Contextual Preference Ranking framework to the performance of the previous algorithms for this dataset. The best performing algorithm reported in [18] uses a traditional deep neural network to learn a ranking over all possible cards for a given context. It was trained by directly mapping an 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 Magic draft agent SIAMESEBOT instead learns on pairwise comparisons between the picked card and any other card in the candidate pack \mathcal{P} and therefore generates 2 to 14 training examples from a single pick, depending on the size of \mathcal{P} .

This additional constant factor in the training complexity is to some extent compensated by the fact that we were able to train our network with a much smaller number of training epochs. Due to the large size of the dataset, we split it into 220 sub-datasets. For the results in Figure 7, only a single epoch of 50 of those datasets was used. We, therefore, used less than $1/4$ of an epoch of the whole dataset, in contrast to 20 epochs of training on the complete dataset in [18].

Following [18] we report 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)* is how far away the correct pick is from the chosen card when ranking all possible choices. In both of those metrics, with embedding dimension $D = 256$, we achieve substantially improved results on the dataset, as can be seen from Table I.

This strong increase in performance suggests that our Contextual Preference Ranking approach works well for this domain. Furthermore, the proposed approach is completely domain-agnostic. Apart from having a fixed one-hot encoding of each card, we do not provide the network any other information about the game or the cards. This leads us to speculate that our method will likely work well for other contextual decision-making problems.




Fig. 5. Accuracy of pick-prediction per number of already chosen cards. SIAMESEBOT performs much better than all other methods and remains a more stable accuracy in the middle of the packs.

B. Draft Analysis

We also compare the performance of our proposed network 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 our SIAMESEBOT and those reported in [18] over the three consecutive picking rounds with 15 cards each. We clearly see that our method generally provides substantially more accurate decisions. Interestingly, the accuracy of picks does not show the same performance curve as previous methods. Those methods have U-shaped curves and are more accurate at the start and end of each pack. Our method remains a relatively stable quality throughout most of the pack. The worst accuracy for our SIAMESEBOT occurs at pick number 2, which is an interesting observation. A possible reason for this may be, that the embedding of the context with a single card, and the embedding of a possible choice card itself are the same, which could lead to problems there.

C. Visualization

Finally, we can use the resulting embedding of cards to visualize the decision process of the network. Since embeddings for cards are constant, the card selection decision is solely determined by the embedded representation of the anchor.

Visualizing the embedding of single cards aids in understanding of the decision process of the network. As the embedding is 256-dimensional, we use t-SNE [16], a stochastic algorithm to reduce the dimensionality of data points, to plot the embedding in two dimensions. We graph all cards in their respective color. Cards of exactly two colors are shown with one color as their border and the other as the filling. Purple is used for colorless cards, and gold for cards of more than two colors. The empty set is shown as the anchor in the middle, which corresponds to visualizing the first pick.

Although the network did not receive any information about the color of cards, Figure 6 shows clear clusters with cards




Fig. 6. 2D visualization of embedded cards with the anchor as the empty card set and colors matching the card colors in the game. Clear clusters of equal-colored cards are visible, although the network did not receive this information.

of the same color. Within each cluster, cards seem to be roughly linearly ordered, where cards closer to the empty set are stronger in a vacuum. This leads to a star-like structure. Between clusters, single multicolored cards are visible, which correspond to multicolored cards of the two adjacent colors.

However, some clusters outside the star structure are visible. As the dimensionality of the embedding was reduced drastically, giving an accurate explanation of those is far from trivial. Firstly, as t-SNE is stochastic, the resulting 2-dimensional representation changes in different iterations. Therefore, some clusters sometimes seem to be connected, while they are disconnected in other iterations. Figure 6 shows two distinct black clusters, but in other iterations with the same parameters, the structure of those is more connected. We also use k-means to cluster the data in its original space. There, the far-seeming white points are sometimes clustered together, which further strengthens the explanation that some structures are artifacts of t-SNE. This is even more drastic when taking different hyperparameters into account.

Finally, we tested how the observable distances to the anchor in Figure 6 correlate to the real distances in the 265-dimensional embedding space. Those two measures achieve a Kendall's Tau correlation of 0.6243, which means they are strongly correlated, showing that Figure 6 still gives good intuition about the decision-process of the network. When investigating the correlation further, the loss in correlation mainly comes from singular points and clusters which have drastically different distances in the two embedding spaces, while clusters themselves seem to achieve a similar distance.

D. Sensitivity to Embedding Dimension

The embedding dimension D is the most important hyperparameter of the proposed method. Figure 7 shows the learning curves for different choices of D . Increasing D leads to strong improvements in network accuracy up to about $D = 32$. After this, diminishing returns set in and the performance stops improving at $D = 128$. However, even with




Fig. 7. Influence of D on performance. Increasing D leads to higher accuracies until $D \approx 128$.

$D = 2$, the Siamese network achieves an accuracy of 53.69%, which is higher than the 48.67% of the best previous method NNETBOT (see Table I).

VIII. DISCUSSION

With the proposed approach, we achieve much higher accuracies than previously reported, while using a much smaller amount of training. Besides, the embedding of cards and decks provides valuable information about the dataset without any added computational effort. The fact that cards of the same color are clustered together is intuitive and further confirms the validity of our approach. Embedding a different dataset would likely look vastly different, for example, a set where specific colors are more likely to be drafted together. One surprising finding was that the embedding was not intuitively perfect. A few outside clusters of cards (Figure 6) seem to be rated drastically worse than the main structure around the anchor. A possible explanation is that those clusters contain weak cards which are chosen near the end. This would still cluster them together, as they are only chosen in accordance with the color of the deck while being exceptionally far away from the empty set since they are never chosen as the first pick. It however could also just be an artifact of the dimensionality reduction of the data. Another reassuring finding is that two-colored cards lie between their two colors, as those cards are equally relevant for both.

We can use the resulting embedding to construct a rating of all cards by computing their distances to the empty set. Interestingly, this differs from expert rankings. We compare this resulting ranking to two expert opinions in Table II. The last column ranks the cards based on how often they were first picked in the dataset (compare Figure 2). The rarity of each card is encoded behind its name as either *Uncommon* (U), *Rare* (R) or *Mythic* (M).

From this, a few stark differences are immediately obvious. While *Ajani, Adversary of Tyrants*, *Djinn of Wishes* and *Leonin Warleader* are rated similarly to the experts, the extraordinarily high rating of *Goblin Trashmaster* is surprising. Note that the rates for the top cards are very similar, e.g. 99.94% for *Ajani*,

TABLE II
RANKINGS OF FIRSTPICKS OF THE PROPOSED METHOD COMPARED TO EXPERT OPINIONS. FPR = FIRSTPICKRATE

Card	Siamese	Expert 1 [15]	Expert 2 [8]	FPR
Spit Flame (R)	1	18	22	17
Leonin Warleader (R)	2	15	4	8
Goblin Trashmaster (R)	3	51	112	32
Ajani, Adversary of Tyrants (M)	4	7	5	1
Djinn of Wishes (R)	5	14	6	14
Tezzeret, Artifice Master (M)	20	1	2	3
Resplendent Angel (M)	30	9	1	2
Murder (U)	12	21	9	39




Fig. 8. Correlation between first-pick rate and distance. Cards with a higher first-pick rate are embedded closer to the empty set. Kendall's Tau = 0.74

Adversary of Tyrants and 99.87% for *Resplendent Angel*, and those *Rare* and *Mythic* cards are never in direct competition due to the composition of the decks. This can make a correct ranking very hard for the network. We can also observe from the Siamese ranking that, surprisingly, four of the top five cards are *Rare*. We speculate that this is due to the fact that *Rare* cards occur more frequently in the dataset and the training sees more positive examples of these. It is possible to combat this by oversampling mythic examples, or by adding features to the cards, but this would stand in contrast to the domain-agnostic approach chosen. While the high ranking of *Goblin Trashmaster* is unusual, the network has, however, made a precise estimation about a hard-to-rate card, *Murder*. This is by far the best *Uncommon* card in the set. It is ordered at rank 9 and 21 by CFB and DraftSim respectively. Our Siamese network ranks it at 12, although its first-pick rate is only 39.

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 8, 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 more than just the first picks, the embedding distance is a smoother measure of how strong the card is according to the network.

Finally, we can use the embedding to extract meta-information about this dataset. For example, the Siamese network seems to strongly favor the colors red and white, as it rates four white and five red cards higher than the best green one.

IX. CONCLUSION

We showed that the proposed method of using a Siamese network to model preferences in the context of drafting cards in *Magic: The Gathering* worked well and vastly outperformed previous results. Compared to [18], we report an increase in accuracy by more than 56%, while also decreasing the pick distance by more than 83%. Even when our network makes an incorrect choice, the network ranks the correct choice very high. In addition to this performance, we show that the resulting embedding makes intuitive sense. It can be used to learn further from the dataset, apart from only using it for draft predictions. For this dataset, we were able to create absolute rankings of cards and could speculate which colors SIAMESEBOT prefers.

With this first implementation of a contextual preference ranking framework, we showed that Siamese networks work well for adding items to an existing set. We want to reemphasize that while this is the first practical test of this framework, there is no reason to believe that the success is limited to this particular setting. We did not incorporate any domain information into SIAMESEBOT beyond the ID of cards used to encode the input. Therefore, we speculate that our proposed framework will work well for other problems where preference has to be modeled in a context.

X. FUTURE WORK

In order to further test the generality of this approach in other domains, more work with other datasets is required. One possible area for future work is sequential team-building in a MOBA game. It could also be possible to extend this approach beyond sequential decision-making. An example is a game where decks are played against each other, and the context is the intersection of both decks, with positive and negative examples taken from the remaining cards of the winning and losing deck respectively. This may introduce a lot of noise into the training, as winning or losing with a deck is subject to a multitude of factors besides the chosen cards, but may extend the method to a larger variety of domains.

There is potential to use this method not only for pre-game decision-making but for game playing as well. Given a dataset of expert moves in a game, we can model the anchor as the current game state, and the chosen and one not chosen move as positive and negative examples. A concern with all of those ideas however is the fact that we are solely training on human expert examples, which provides an upper limit on how well this can perform in a general context. To circumvent this, one could also generate datasets on self-play games as part of an agent training loop, as in AlphaZero [13] and similar approaches. For further improving performance within MTG, we could build refined architectures that use

meta-information and a history, which allows inferences about opponent strategies and color choices which are used by strong human players. We could also try to train separate networks for specific numbers of already chosen cards, especially for the case of 1 chosen card where performance is worst overall.

Acknowledgements We thank the authors of [18] for making the data publicly available and for sharing their experimental data, and Johannes-Kepler Universität Linz for supporting Müller’s sabbatical stay through their Research Fellowship program.

REFERENCES

- [1] Bhatt, A., Lee, S., De Mesentie Silva, F., Watson, C. W., Togelius, J., and Hoover, A. K. (2018). Exploring the Hearthstone deck space. *Proceedings of the 13th International Conference on the Foundations of Digital Games (FDG)*, Malmö, Sweden. ACM.
- [2] Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. (2005). Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning* (pp. 89-96).
- [3] Chicco, D. (2021). Siamese neural networks: An overview. In Hugh M. Cartwright (ed.) *Artificial Neural Networks*, 3rd edition. Springer.
- [4] Clevert, D. A., Unterthiner, T., and Hochreiter, S. (2016). Fast and accurate deep network learning by exponential linear units (ELUs). In *4th International Conference on Learning Representations (ICLR), Conference Track Proceedings*, 1–14.
- [5] Draftsim dataset, <https://draftsim.com/draft-data/>
- [6] García-Sánchez, P., Tonda, A. P., Squillero, G., García A. M., Merelo Guervós J. J. (2016). Evolutionary deckbuilding in Hearthstone. *Proceedings of the IEEE Conference on Computational Intelligence and Games (CIG)*.
- [7] Hoover, A. K., Togelius, J., Lee, S., and de Mesentier Silva, F. (2020). The Many AI Challenges of Hearthstone. *KI – Künstliche Intelligenz* 34(1):33–43.
- [8] Karsten, F (2018). An Early Pick Order List for Core Set 2019. Retrieved April 02, 2021 from <https://strategy.channelfireball.com/all-strategy/mtg/channelmagic-articles/an-early-pick-order-list-for-core-set-2019/>
- [9] Koch, G., Zemel, R., and Salakhutdinov, R. (2015, July). Siamese neural networks for one-shot image recognition. In *Proceedings of the ICML’15 Deep Learning Workshop*.
- [10] Kowalski, J., and Miernik, R. (2020). Evolutionary approach to collectible card game arena deckbuilding using active genes. *Proceedings of the IEEE Congress on Evolutionary Computation (CEC)*, Glasgow, United Kingdom.
- [11] Lian, Z., Li, Y., Tao, J., and Huang, J. (2018). Speech emotion recognition via contrastive loss under Siamese networks. In *Proceedings of the Joint Workshop of the 4th Workshop on Affective Social Multimedia Computing and 1st Multi-Modal Affective Computing of Large-Scale Multimedia Data (ASMMC-MMAC)*, pp. 21–26.
- [12] Scott-Vargas, L. (2018). Core Set 2019 Limited Set Review: White. Retrieved April 02, 2021 from <https://strategy.channelfireball.com/all-strategy/mtg/channelmagic-articles/core-set-2019-limited-set-review-white/>
- [13] Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., ... and Hassabis, D. (2017). Mastering the game of go without human knowledge. *nature*, 550(7676), 354-359.
- [14] Tesauro, G. (1988). Connectionist Learning of Expert Preferences by Comparison Training. *Advances in Neural Information Processing 1 (NIPS)*, pp. 99–106.
- [15] Troha, D. (2018). Draftsim’s Pick Order List for Core Set 2019. Retrieved April 02, 2021 from <https://draftsim.com/M19-pick-order.php>
- [16] Van der Maaten, L., and Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research* 9:2579–2605.
- [17] Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., and Wierstra, D. (2016). Matching networks for one shot learning. *Advances in Neural Information Processing 29 (NIPS)*, pp. 3630–3638.
- [18] Ward, H. N., Brooks, D. J., Troha, D., Khakhalin, A. S., and Mills, B. (2020). AI solutions for drafting in Magic: The Gathering. *arXiv preprint 2009.00655*.
- [19] Wizards of the Coast (2021). Magic: The Gathering Comprehensive Rules. Retrieved April 07, 2021 from <https://media.wizards.com/2021/downloads/MagicCompRules%2020210224.pdf>