Contrastive Learning for Game Artificial Intelligence

On the Example of Magic: The Gathering

Timo Bertram, Johannes Fürnkranz, Martin Müller
Johannes Kepler Universität Linz

Introduction

Contrastive Learning is a powerful tool to train neural networks with comparisons rather than explicit targets. In our work, we investigate how we can use contrastive learning for games, e.g. for comparing actions or game states. Here, we show how to compare cards in a collectable card game.

Instead of comparing items in a vacuum, we focus on comparisons with a context, which is often more applicable than trying to compare with no basis. For this showcase, we apply our method to the domain of drafting, i.e. building a deck of cards to play Magic: The Gathering with, From human data, we learn sample triplets of decisions, consisting of a card preferred over another card given a deck of cards.

$$(\mathbf{c}_j \succ \mathbf{c}_k \mid \mathcal{C}) \, .$$

Method

We term our method **contextual preference ranking** and model it with a Siamese Neural Network. Such networks are often used in image recognition for one-shot learning, which results in a learned embedding space with similar images in close proximity. In our case, we model decisions, i.e. one item is better than another, rather than similarity. We embed the deck, the chosen card and the unchosen card and train the Siamese Network to adjust the embedding distances according to the preferences.

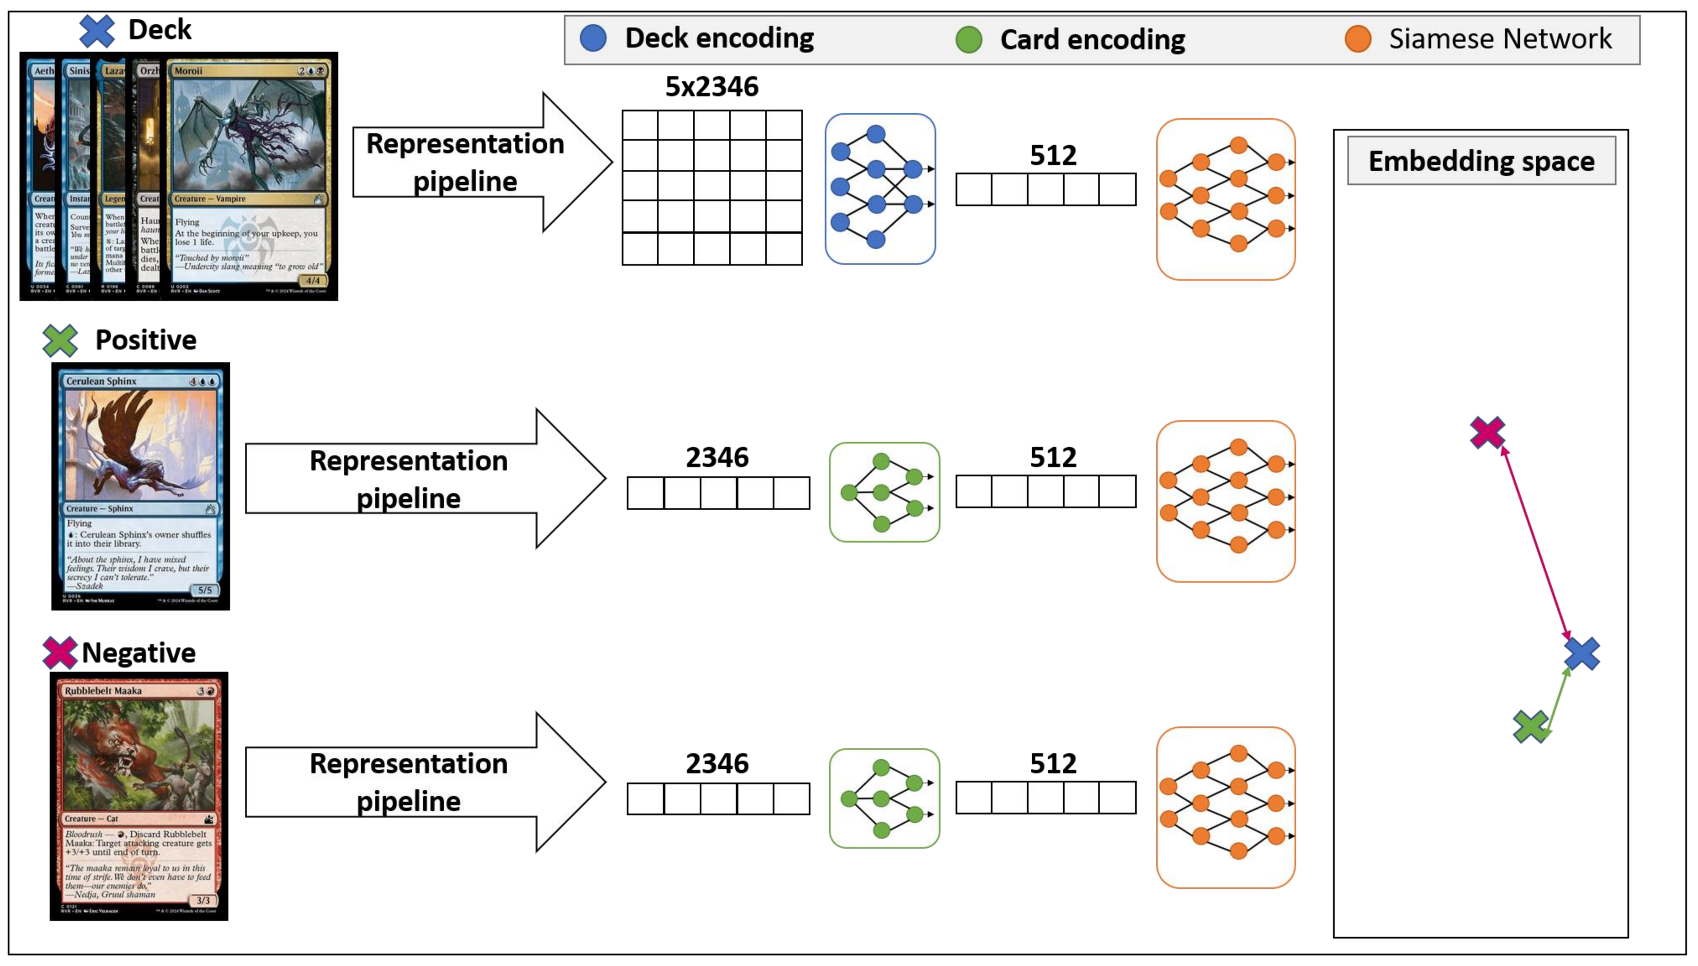

This research is concerned with two aspects:

- **How accurately can we predict human decisions?**
- **How can we meaningfully represent cards such that we can generalize to unseen ones?**

Results

We test several different ways to represent cards, including one-hot encodings, random encodings, image encodings, hand-crafted feature encodings and combinations of them. In this experiment, we train on one set of cards (250 unique ones) and test on held-out decisions of that set as well as completely unknown cards. We report the accuracy of picking the correct card out of a larger set, varying between 2 and 14.
When testing on known cards, i.e. ones used in training, we find little difference between the representations. Thus, the model remembers cards rather than generalizing across features. However, when regarding unknown cards, i.e. requiring an understanding of their characteristics, the specific encoding used is crucial for success.

TOP-1 TESTING ACCURACY OF MODELS WITH DIFFERENTLY ENCODED INPUTS. ALL MODELS ARE TRAINED SOLELY ON *NEO* AND TESTED ON A HELD-OUT *NEO* TEST-SET AND UNSEEN CARDS. ACCURACY IS AVERAGED ACROSS ALL UNSEEN SETS.

Model	Input size	*NEO* test	unseen sets
One-hot	302	67.80%	NaN
Random	1024	67.87%	23.79%
Image *32*	32	65.93%	28.69%
Image *1024*	1024	68.09%	31.10%
Meta	16	64.73%	42.14%
Features	1306	67.76%	33.57%
Features + Meta	1322	68.07%	34.74%
Features + Image *1024*	2330	67.81%	35.59%
Features + Meta + Image *1024*	2346	68.00%	42.87%

To create a generalized drafting model, being able to model preferences between arbitrary cards is important. This includes existing, but unseen, cards, as well as not yet released ones. For this, the best-performing encoding uses a combination of features, a learned image representation and card-specific meta-information.

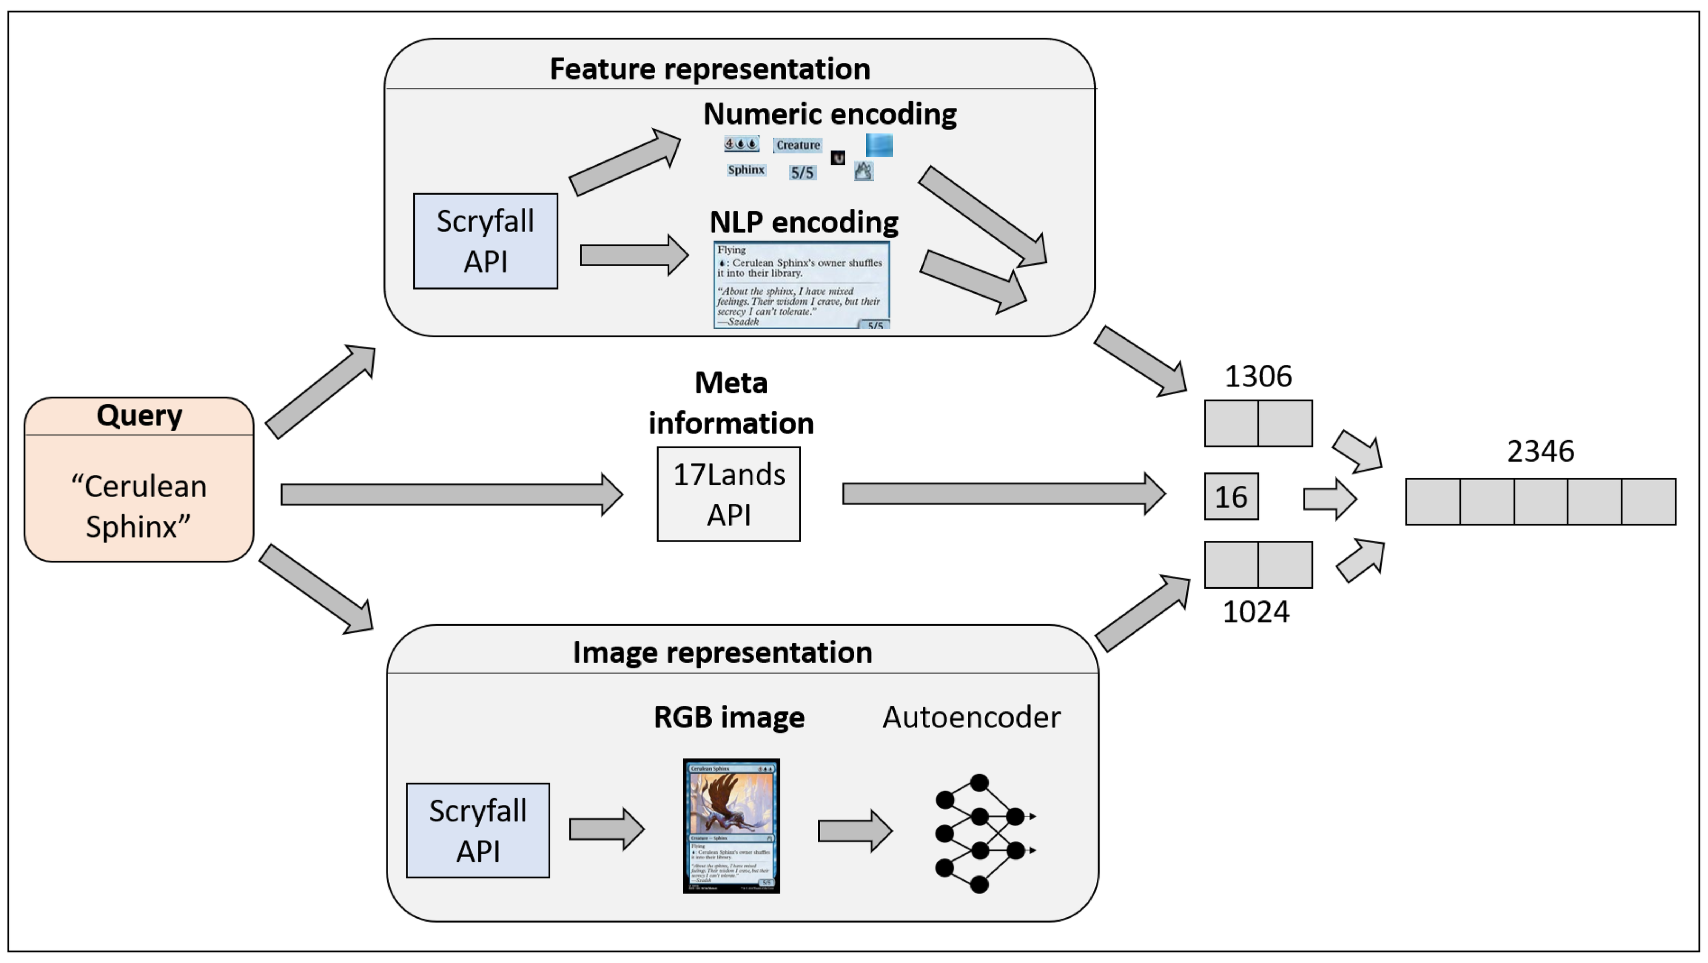

Conclusion

For this domain, we achieved good success with the **contextual preference ranking** framework.
In addition to strong predictive capabilities, we receive an intuitive embedding space with similar items positioned in close proximity. While we so far solely use the distances for predictions, one could further investigate the embedding and learn more about the underlying relations.

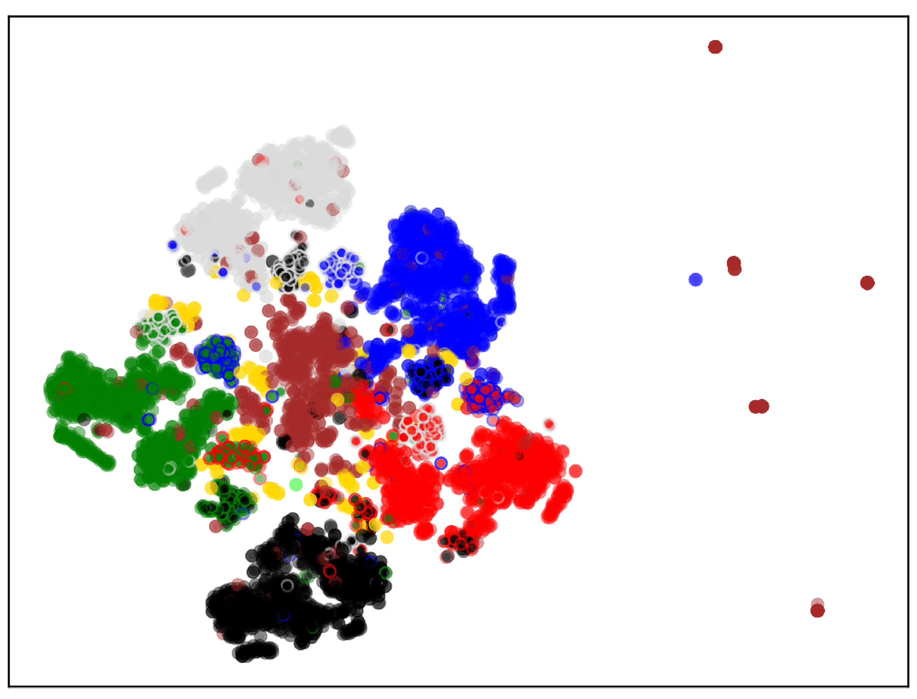

We believe that with this research, we constructed the currently strongest model for drafting in Magic. The Gathering. The result could be used as an opponent or as a player-aid for less experienced drafters.

Other Applications

In addition, the underlying contextual preference ranking framework is a general tool applicable to any kind of preference within a context. For example, we conducted further research in the area of *Reconnaissance Blind Chess*, where we learn which (hidden) game board is most likely given the player's information.