Exceptional Gestalt Mining: Combining Magic Cards to Make Complex Coalitions Thrive

Wouter Duivesteijn[1] (✉) and Thomas C. van Dijk[2]

[1] Technische Universiteit Eindhoven, Eindhoven, the Netherlands
w.duivesteijn@tue.nl
[2] Ruhr-Universität Bochum, Bochum, Germany
thomas.vandijk@ruhr-uni-bochum.de

Abstract. We propose Exceptional Gestalt Mining (EGM), a variant of Exceptional Model Mining that seeks subgroups of the dataset where a coalition becomes more than the sum of its parts. Suppose a dataset of games in which several roles exist within a team; the team can combine forces from any subset of roles, to achieve a common goal. EGM seeks subgroups for which games played employing a large role set have a higher win rate than games played employing any strict subset of that role set. We illustrate the knowledge EGM can uncover by deploying it on a dataset detailing *Magic: The Gathering* games: we find combinations of cards that jointly work better in multicolor decks than in decks employing fewer colors. We argue that EGM can be deployed on datasets from sports where several roles exist that directly interact in play, such as ice hockey.

Keywords: Local Pattern Mining · Magic: The Gathering · Team Composition · Gestalt · Subgroup Discovery · Exceptional Model Mining.

1 Introduction

Magic: The Gathering (MTG, or *Magic* for short) is a turn-based collectible card game, where prior to playing the game proper, one builds a synergistic deck from a set of available cards; in the actual game of Magic, each player uses their own deck (which is typically constructed from the same set, but not necessarily the same sample from the set: some ways of deck construction involve chance or a drafting mechanic). The game was first released in 1993 as physical playing cards and has taken quite a flight since then: the initial stock of 10 million cards sold out within three months; a "Pro Tour" and several independent tournament circuits have existed since at least 1996. The game was patented in 1997 [13] and the game was acquired by Hasbro in 1999 for 325 million dollars. After almost two decades and multiple versions of the much-maligned *Magic Online*, Magic only took off as a proper e-sport in 2018 with the release of *Magic Arena*.

Fig. 1: Two example Magic cards. On the left hand, a basic *land* that provides *mana* of a specific *color* (in this case: green). On the right hand, a *spell* that can be played when the mana cost is paid, some of which must be in specified colors (in this case: two mana of any color plus one green mana). In this example, the Forest can contribute to paying the cost for Jorn, God of Winter.

1.1 Core Game Mechanics of Magic: The Gathering

A crucial feature of Magic's gameplay is the "mana system". Some of the cards are *lands* that produce mana of a certain *color*[3] (white, blue, black, red, or green). This mana is a resource used to play other cards called *spells*, which actually produce game effects as indicated on the card. This typically advances the game state in a manner that brings the player closer to their end goal. A game of Magic can be won in various ways, most prominently by reducing the opponent's life points from the starting total of 20 down to zero (or below) and by depleting the opponent's deck; other win conditions also exist, as specified by individual cards.

Of particular relevance to our analysis are the five *basic lands*, which are the main ways to acquire each color of mana. Most spells specify that part of their cost must be paid with mana of a specific color (or set of colors); a spell is said to be a certain color if it requires mana of that color. Figure 1 gives an example of a basic land and a spell that share a color; hence, the former can contribute to paying the cost required for playing the latter.

Beginning players often limit themselves to building a deck of cards using only a single color (for example: only forests and green spells). This is at least somewhat synergistic, since the cards of each color tend to lean into particular playing styles (with strengths and weaknesses) and the deck needs only one color of mana – and therefore only one type of basic land. Playing more than one color in the same deck allows the player to balance weaknesses with strengths

[3] Most of our descriptions of the game are not completely accurate. With over 20 000 distinct cards there is an exception to almost any generalization. In this case: some lands can produce mana of multiple colors, or produce no mana at all. Our descriptions only serve to illustrate the context of the dataset.

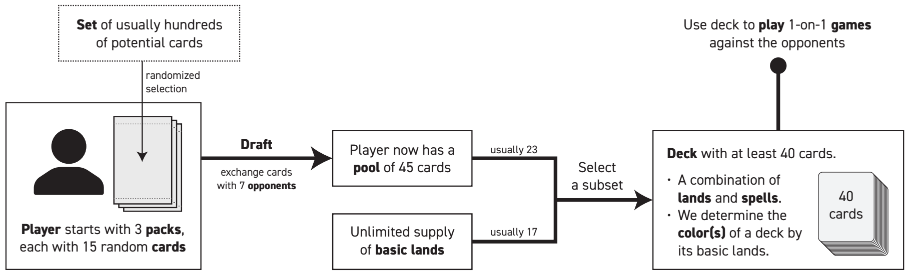

Fig. 2: Flowchart detailing the process of a Magic draft.

and to combine powerful cards that happen to be in different colors. However, this introduces the risk of drawing mismatched lands and spells. The strategy behind color choice and selecting the optimal mix of lands (the "mana base") is deep and extensively discussed in the community. Here, we leave it at the observation that many decks play two colors. Even among professional players, a typical rule of thumb is that one is best served by limiting colors in a deck to two or three; exceptions of course exist.

In this paper, we analyze data from *draft* Magic, where the players open a booster pack containing a random selection of 15 cards from a particular set, select a card from the pack for their own pool, and pass the remainder of the pack to the next player. This process of selecting a card and passing on the others continues in a cycle until no cards remain, after which new packs are opened. In total, three packs are opened per player, and these three packs typically contain an identical distribution stemming from a certain collection of cards (called a *set*) such that gameplay can be expected to be thematically coherent. All players in a draft of Magic will see only a small subset of cards in the game: they will not have access to the same sample of cards and each will need to craft a coherent deck from whichever cards they managed to acquire through the drafting process. Independently of the draft packs, they may add any number of each of the five basic lands. See Figure 2 for an overview of this process.

1.2 Gathered Data

For players who want to analyze their performance in online draft Magic, the people at `17lands.com` – so-named after the conventional wisdom that playing 17 lands in a 40-card draft deck is optimal (although, again, exceptions apply) – have created an application that monitors the Magic Arena log files to track the state of the drafts performed and games played. The goal is not only to help players understand their own performance, but also to enable analytical insights from aggregated data. Some of these datasets are publicly available [20].

We analyze the Kaldheim (KHM) Traditional Draft dataset [19]. This is a single csv file containing game-level data. Every row describes a single game: its outcome (win/loss), some metadata on the players, the list of cards drafted by the player, and the list of cards in the player's deck. Existing analyses of this dataset focus on finding individual cards that affect the win rate: if it is in the deck, how is the win rate affected? Or: does it come at an increased or decreased win rate for the top-10% or bottom-10% players?

1.3 Main Contribution

In this paper we investigate the concept of *gestalt*: a total that is bigger than the sum of its parts. We seek combinations of cards, that work better when they are in decks containing a bigger coalition of colors: we will deem the combination of cards interesting, if its performance in the multicolor combination is substantially better than its performance in subcoalitions of fewer colors.

Our main contribution is Exceptional Gestalt Mining, an instance of Exceptional Model Mining [12] that finds such exceptional combinations. This is also valuable for some major sports leagues, especially team sports where players have distinct roles or styles but do interact synergistically. Association Football and Ice Hockey are examples of such sports; American Football is not.

2 Related Entities

2.1 Related Work on Magic: The Gathering

Magic has already received some interest from the academic community, though not necessarily from an analytics perspective. Magic as a rule system is very complex – indeed, following the rules is Turing complete [6][4] and even checking whether particular game moves are legal is coNP-complete [4]. As a game with partial hidden information, its strategy is also challenging [5]. More tangentially, Magic and its community have been the subject of experimental studies on auction design [22] and marketing [23].

Outside of academia there is a large and active corpus of online strategy advice; as mentioned in the introduction, this ranges from folklore and anecdote to fairly advanced probability theory and more recently: data mining. Central to a lot of the discussion is the concept of *win percentage*: what fraction of games played against a (theoretical) field of opponents, including the distribution of player skill and their deck selection, do you expect to win? This is precisely the kind of data that is sampled by `17lands.com`. In informal blog publications, they discuss wide-ranging themes revolving around this data, including the effect on win rate of a changed mulligan rule [8], simulations of draft strategies [9], winning deck archetypes in the Kaldheim set [24], and the limitations of win rate as a measure of success [25].

[4] Cards can react to in-game events with new effects of their own and the game rules define that if an unbreakable loop occurs, then the game is a draw; simulating a Turing machine with the available cards is nontrivial, but was finally achieved in 2019 after multiple steps of partial progress.

2.2 Related Work on Local Pattern Mining

Exceptional Gestalt Mining is a form of Local Pattern Mining (LPM), a subfield of data mining exploring local structures in a dataset. The best-known LPM method is Frequent Itemset Mining and its cousin Association Rule Mining [1]. Here, all data is binary, and the focus lies on finding sets of products that co-occur frequently in transactions. Associations between these sets are sought; the process is unsupervised. Supervised Descriptive versions of LPM [18] exist. Subgroup Discovery (SD) [16,26,15] is such a version, seeking subgroups of the dataset where a designated target variable displays an unusual distribution. Emerging Pattern Mining [10] and Contrast Set Mining [2] employ different task formulations, which have been shown [18] to be unifiable with SD.

Exceptional Model Mining (EMM) [21,12] can be seen as a multi-target generalization of SD. Rather than a single target, EMM typically designates several columns as targets. Whereas SD assesses the interestingness of subgroups in terms of a difference in target distribution, EMM typically seeks a behavioral difference in some sort of interaction between the targets. For instance, one could seek an unusual correlation between two targets [21], allowing the identification of subgroups of the housing market where the lot size and sales price of a house are uncorrelated, while this correlation is positive on the entire dataset.

This last example illustrates an interesting design choice in EMM: in order to find subgroups displaying exceptional behavior, one must choose when behavior is exceptional; to do so, one must characterize what constitutes normal behavior. Naive approaches [12, Section 3.2.2] take behavior on the full dataset or the complement of the subgroup as a proxy for normality. This assumes that behavior on the entire dataset can be treated as a monolithic whole, which may not be realistic. A recent alternative [11] employs the Chinese Restaurant Process to model multiple kinds of normal behavior in the dataset; a subgroup is exceptional if its behavior matches none of the normal behaviors. Yet another alternative only compares a subgroup's behavior locally, with the behavior of a peer group [17,14]. We take yet another approach, evaluating subgroups by seeking an *internal* behavioral bifurcation: a subgroup is interesting if it performs much better in a specific coalition of colors than it does in all subcoalitions; this evaluation is agnostic to behavior outside the subgroup.

2.3 Related Sports

So far we talked a lot about a card game and the potential for mining due to the availability of data. This paper, however, appears in the Proceedings of the 8th Workshop on Machine Learning and Data Mining for Sports Analytics. Besides arguing that Magic is a legitimate e-sport with professional players, we believe that there are fruitful parallels particularly to *team* sports. One way to make the analogy is as follows.

A deck consisting of Magic cards can be seen as the total set of players that a sports organization[5] has available to compete in matches in a season. Not all of these players participate in every individual team sports game, and similarly, not all cards in the deck are drawn during a particular game of Magic. A draft, like the one described for Magic, happens in all the big North American sports leagues (NBA, NFL, MLB, NHL), where organizations take turns to select new players from a limited pool; specific draft mechanics will be different across sports league from the ones in Magic, but the turn-based limited resources concept is shared. The colors of Magic – with distinct strengths and weaknesses that a deck can emphasize or counteract through combination – can be likened to strategic or tactical roles that players can have within the team. Some of these roles might have a tendency to synergize, such as “patient defense” with “counter attack” in Association Football, or the combination of red and white cards in so-called aggro decks; this does not mean that every red card works well with every white card, or that every patient defender plays well with every fast runner. Some players or combinations of players might perform especially well in teams that include a particular role other than their own (most interestingly when this not a general synergy between the roles, but in some way an exceptional synergy). A specific hockey center might score more shorthanded goals from assists of certain defenders on a specific penalty kill unit, than even-strength goals from open play on regular lines; a cheap green card that enables “mana fixing” might do particularly well when combined with red and white, even if green cards in general might not.

The concept we strive to find in this paper, a combination of Magic cards (or: players in a sports team) that works better in a particular multicolor coalition than in smaller subcoalitions, maps better to some sporting contexts than others. The analogy does not work well for the NFL, for example. In an American Football team, the offensive and defensive lines both work towards a common goal, but they never interact: these lines are never on the field at the same time. There may be some overlap between the offense or defense and some of the special teams, but by and large, gestalt cannot be expected between these lines. The analogy works much better for the NHL as seen by the example in the previous paragraph. In an Ice Hockey team, offensive and defensive lines each have their own major roles: offense should largely score goals, defense should largely avoid conceding goals. But these lines are on the field simultaneously, and the team cannot function well by focusing on a single strategy to the neglect of all others[6].

[5] Depending on what a specific sports league allows, this may include the current main squad, youth players, minor league affiliate team players, loan players, and players acquired in mid-season transfers.

[6] In fact, one standard hockey player performance metric, called +/–, acknowledges that good defenders enable a strong offense and good offensive lines contribute to a strong defense. The metric records the number of goals the team scores while you are on the ice, minus the number of goals the team concedes while you are on the ice. Hence, top-scoring centers or wingers who neglect their defensive duties can be found out by comparing their performance in terms of goals and assists with their performance in terms of +/–.

In some top NHL teams, offensive lines can be clustered into those that excel in offense itself, and those that excel in defensive support. Here, the synergy between offensive and defensive lines becomes of paramount importance: if an offensive line gels really well with a defensive line, their combination becomes much more than the sum of their parts. This is where exceptional gestalt can be found.

3 Exceptional Gestalt Mining

We assume a dataset Ω, which is a bag of N *games* of the form $g^j = (a_1^j, \ldots, a_k^j, R^j, \ell^j)$; the superscript index j refers to a specific game, and is omitted when unnecessary. Here, $a_1, \ldots, a_k$ are the *entities* in the data: binary columns indicating Magic cards, or sports players, that can form a part of an exceptional combination. The final column, ℓ, is the *outcome* of the game, which is binary: the game is either won or lost. Finally, we assume the existence of a set $\mathcal{R}$ of *roles* within the game: this can be the set of colors in an Magic game or the set of roles players take on a hockey team[7]. In every game g^j, a set $R^j \subseteq \mathcal{R}$ of roles is *deployed* during g^j.

Exceptional Gestalt Mining (EGM) falls under the framework [12] of Exceptional Model Mining (EMM), which seeks interpretable subsets of the dataset that behave exceptionally. A subgroup is defined as a conjunction of conditions on the entities (for instance: $a_7 = \text{false} \ \wedge \ a_{112} = \text{true}$), and this logical expression selects a subset of games from Ω. Informally, we call the conjunction the *description*, and the selected subset the corresponding *subgroup*, although we will conflate the two concepts if no confusion can result. Formally, we let a description D be a function $D : (a_1, \ldots, a_k) \to \{0, 1\}$, and the subgroup G_D corresponding to description D will be $G_D = \{ \ g^j \in \Omega \mid D(a_1^j, \ldots, a_k^j) = 1 \ \}$.

The beam search algorithm for EMM [12, Algorithm 1] traverses the space of candidate subgroups, evaluating their exceptionality along the way. That evaluation requires a *quality measure* φ, taking a description (or subgroup) as input and returning a value in $\mathbb{R}$, where conventionally higher is better. Philosophically, the quality measure must reflect how exceptional within-subgroup behavior is, when compared with behavior on a well-chosen reference group. This behavior can be a simple correlation coefficient or slope of a regression model, and the reference group could be the entire dataset. For EGM, however, we make different choices. The *win rate* of a subset $G \subseteq \Omega$ is $\varphi_{\text{winrate}}(G) = \frac{1}{|G|} \sum_{g^j \in G} \ell^j$. We could use this quality measure to find (combinations of) cards where the win rate is exceptionally high or low, but this would merely tell us what the best or worst cards in the pool are. Instead, we look at how win rate and roles interact.

[7] There may be multiple granularities on which roles make sense. For instance, in a hockey team, one can specify $\mathcal{R} = \{\text{defender, goalkeeper, forward}\}$ or $\mathcal{R} = \{\text{center, defender, goalkeeper, left winger, right winger}\}$. In Exceptional Gestalt Mining, we need to pick one set of roles, and stick with it.

3.1 Measuring Exceptional Gestalt

It is quite likely that a given combination of cards, when used as a description, defines a subset of games featuring more than a single distinct role set. We explore this spectrum of role sets, and the distribution of win rates across that spectrum. Given a subgroup G_D, define its *Role Set Set* (RSS) to be

$$RSS(G_D) = \left\{ R^j \;\middle|\; g^j \in G_D \right\}$$

and its *Conditional Win Rate* (CWR) *given role* R as:

$$\varphi_{\text{cwr}}(G_D, R) = \frac{\sum\limits_{g^j \in G_D \text{ s.t. } R^j = R} \ell^j}{\left| \left\{ g^j \in \Omega \;\middle|\; D(a_1^j, \ldots, a_k^j) = 1 \,\wedge\, R^j = R \right\} \right|}$$

The *gestalt* quality measure can now be defined:

$$\varphi_{\text{gestalt}}(G_D) = \max_{R_i \in RSS(G_D)} \left(\varphi_{\text{cwr}}(G_D, R_i) - \max_{\substack{R_j \in RSS(G_D) \\ R_j \subset R_i}} \varphi_{\text{cwr}}(G_D, R_j) \right) \qquad (1)$$

Hence, we seek subgroups for which the CWR given a role set R_i is larger than its CWR given the best possible strict role subset $R_j \subset R_i$. The subgroup for which this distance is largest, is the subgroup with the most exceptional gestalt.

3.2 Why Does This Make Sense, Intuitively Speaking?

If a card r requires mana of a specific color, say, red, it will only function in decks that can produce red mana. If another card b requires black mana, most decks featuring r and b will contain both Mountains and Swamps (the red and black basic lands). For this combination of two cards, we expect the RSS to consist of supersets of {red, black}. It is possible that the combination of r and b works even better when combined with playing styles that are typically associated with a third color, e.g. blue. In this case, there is not necessarily a *specific* blue card u that interacts well with r and b (though a positive association with most good blue cards is likely present). However, one can expect that the combination of r and b has a higher CWR in decks with role set {red, black, blue} than its CWR in decks with role set {red, black}. In fact, even more colors may be necessary to unleash the full potential of the card combination. This sort of added value is the gestalt for which EGM is designed.

4 Experimental Setup

The Kaldheim Traditional Game Dataset [19] encompasses 182 401 rows, each detailing a single game. The dataset logs match information from the point of view of one of the two players in the game, namely the player who has installed the 17-Lands plugin. Hence, in each row there is a player and an opponent, which the dataset does not treat symmetrically. We use the 321 columns detailing which of the available cards are in the deck of the player. (The opposing deck is unknown.) In the original dataset – and the game – a card can occur more than once in a deck. We convert these columns to binary conditions: does it appear in the deck at all? These columns form the *entities* $a_1, \ldots, a_{321}$ in EGM. The binary column indicating if the player won the game is the outcome ℓ in EGM.

To perform EGM, we need to know what colors are present in the player's deck. We infer this by aggregating over the card information columns representing that color's basic lands (including snow-covered variants): if the deck contains *Forest*, we say the color *green* plays a role in this deck. Applying this for all basic lands and colors, we obtain the set of deployed roles R for this game.

Our EGM implementation is built on an existing EMM implementation [11]. Source code can be found at our companion website[8]. We parameterize the underlying beam search algorithm for EMM ([12, Algorithm 1]) by setting the beam width $w = 100$, the search depth $d = 3$, and the number of reported subgroups $q = 100$; these settings are in line with existing work. In order to prevent the discovery of spurious subgroups, we skip any terms R_i and R_j in either maximum of Equation (1) if the support for that G_D combined with that role set is below 200.

5 Experimental Results

First and foremost, notice that $\varphi_{\text{winrate}}(\Omega) = 0.54$. Since every MTG game has a winner and a loser, the win rate in the general population must necessarily be 0.5. However, the players in our dataset opted to install the 17 Lands plugin to track their statistics, and it stands to reason that this sample of players skews away from the most casual players, which in turn likely skews the win rate upwards.

For each experiment that follows, we report the top-five distinct subgroups found with Exceptional Gestalt Mining. More results can be found for all these experiments on our companion website[8]. Here we report distinct subgroups; if a subgroup of the form A $\wedge$ B performs best, the subgroup B $\wedge$ A typically ranks second. Such duplicates are filtered out of the results reported here.

5.1 Main Results

Table 1 lists the top-five distinct subgroups found with Exceptional Gestalt Mining. The top subgroup combines two cards. Raise the Draugr is a black card that resurrects creatures that died earlier in battle. Glittering Frost is a green card that allows a land to generate additional mana of any color. Decks in which these cards jointly appear almost surely have black and green in their role set; the result in the table indicates that these cards perform best in an extended role set with white and blue (the latter is abbreviated as U, since

[8] `http://wwwis.win.tue.nl/~wouter/Gestalt/`

Table 1: Top-5 distinct subgroups found with Exceptional Gestalt Mining.

#	Description	φ_{gestalt}	Best Role Set R_i (CWR(R_i))	Second Best R_j (CWR(R_j))
1.	Raise the Draugr ∧ Glittering Frost	0.2002	WUBG (64.5%)	WBG (43.5%)
2.	Jorn, God of Winter ∧ Snow-Covered Forest ∧ Sculptor of Winter	0.1902	WUBRG (73.7%)	UBRG (54.6%)
3.	Jorn, God of Winter ∧ Glittering Frost	0.1872	WUBRG (72.7%)	UBG (54.0%)
4.	Sulfurous Mire ∧ Masked Vandal	0.1821	WUBRG (63.9%)	WBRG (45.7%)
5.	Inga, Rune Eyes ∧ Behold the Multiverse ∧ Berg Strider	0.1800	UBR (65.8%)	UB (47.8%)

B is black); the conditional win rate of this card duo in WUBG decks is 20 percentage points higher than its CWR in any strict subset of this color set. It is not obvious how these two cards synergize specifically with each other, but a more general interpretation is possible. Glittering Frost makes any color of mana, thereby enabling large role sets. However, due to the specifics of Magic gameplay, Glittering Frost is best suited for long, drawn-out games – which is exactly the kind of game where Raise the Draugr is effective.

The second subgroup combines three cards that do combine explicitly. Jorn, God of Winter is a green snow creature card; when it attacks, all snow permanents are untapped (i.e.: become available for further use in the same turn). Snow-Covered Forest is a source of green mana that would be untapped by Jorn attacking. Sculptor of Winter is a green snow creature that can untap a snow land, which could be the Snow-Covered Forest; since it is a snow creature itself, it gets untapped by Jorn attacking, which allows it to untap a further snow land. All these cards are green, which does not make it immediately apparent why they would be prime candidates for the all-color role set that EGM finds for them. However, all three cards interact with snow, which is a recurring theme throughout cards of all colors in the Kaldheim set, and they make snow-related mana sources and cards more flexible in use. We postulate that these cards make it feasible to unleash the full potential of decks encompassing powerful cards of all colors: without these cards, it would be difficult to juggle mana sources of all colors; with these cards, this is less of a problem. As a consequence, the CWR of this trio of cards in all-color decks is 19.02 percentage points higher than its CWR for any strict subset of colors.

We observe a theme in the top results. They do not necessarily involve cards that display multicolor synergies among themselves. Instead, they are often sets

Table 2: Top-5 distinct subgroups found with Exceptional Gestalt Mining while limiting evaluation to role sets involving a maximal number of roles.

(a) Evaluation limited to roles sets with at most four roles.

# Description	φ_{gestalt}	R_i
1. Raise the Draugr ∧ Glittering Frost	0.2002	WUBG
2. Inga, Rune Eyes ∧ Berg Strider ∧ Behold the Multiverse	0.1800	UBR
3. Disdainful Stroke ∧ Bind the Monster	0.1754	WUBG
4. Shimmerdrift Vale ∧ Narfi, Betrayer King ∧ Snow-Covered Island	0.1705	UBR
5. Shimmerdrift Vale ∧ Jarl of the Forsaken ∧ Snow-Covered Swamp	0.1701	UBGR

(b) Evaluation limited to roles sets with at most three roles.

# Description	φ_{gestalt}	R_i
1. Inga, Rune Eyes ∧ Berg Strider ∧ Behold the Multiverse	0.1800	UBR
2. Shimmerdrift Vale ∧ Narfi, Betrayer King ∧ Snow-Covered Island	0.1705	UBR
3. Behold the Multiverse ∧ Bind the Monster ∧ Ice Tunnel	0.1664	UBR
4. Inga, Rune Eyes ∧ Behold the Multiverse ∧ Snow-Covered Island	0.1648	UBR
5. Disdainful Stroke ∧ Augury Raven ∧ Snow-Covered Island	0.1645	UBR

(c) Evaluation limited to roles sets with at most two roles.

# Description	φ_{gestalt}	R_i
1. Tuskeri Firewalker ∧ Axgard Cavalry	0.0483	UR
2. Tuskeri Firewalker	0.0420	UR
3. Breakneck Berserker ∧ Axgard Cavalry	0.0417	WR
4. Battlefield Raptor	0.0339	WG
5. Run Amok	0.0326	WR

of coherently-behaving cards that allow for more color flexibility (either explicitly by making mana of different colors or by stalling for time, so more lands may be drawn), thus allowing a deck to make the most of multicolor synergies in the wider card set.

5.2 Results when Limiting the Number of Roles

In order to perhaps find more specific results, we change our experiments as follows: in the evaluation of exceptionality with φ_{gestalt}, we only allow role sets R_i to take part that consist of at most r roles. Games with bigger role sets still contribute to the support of a subgroup, but they take no part in CWR computations. Results when limiting the evaluation in this manner can be found in Table 2, for $r \in \{4, 3, 2\}$.

We observe that all top-5 subgroups for $r = 3$ (and indeed, 69 of the top-100) have UBR (blue-black-red) as the superior role set. For all those subgroups, UB is the second-best role set and has a losing record, even though UB in general (outside the subgroup) has good win percentage. That is, of the decks playing these specific cards, the ones with role set UBR have high winrate, but the ones with role set UB have winrate under 0.5 (and other role sets are even worse or have insufficient support).

6 Conclusions

We introduce Exceptional Gestalt Mining (EGM) as a form of Exceptional Model Mining, seeking subgroups whose combination delivers more than just the sum of their parts. EGM is applicable on datasets of games that are either won or lost, where a set of entities can play part in the game or not, and where a set of *roles* exist in the dataset of which a subset is deployed during the game. In such a dataset, EGM finds subgroups of entities whose win rate when a larger coalition of roles is deployed, is substantially higher than its win rate when each strict subset of that role coalition is deployed. Hence, these subgroups display an exceptional level of gestalt.

On a dataset of the Kaldheim Traditional Draft, an online game setting of the collectible card game *Magic: The Gathering*, EGM finds card combinations having a higher win rate in large multicolor decks than in decks of subsets of that color coalition. For instance, in Table 1 we see that the combination of *Jorn, God of Winter*, *Snow-Covered Forest*, and *Sculptor of Winter* has a win rate 19 percentage points higher in decks of all five colors than in decks of fewer colors. All these cards are green; they could function fine in a monocolor green deck. However, their combination makes the deck more versatile, which enables easier combination of forces of multiple colors in a single deck. This is the sort of gestalt that EGM can detect in a dataset.

We argue that EGM is directly deployable on data from major sports leagues, such as the NHL. In its current form, EGM requires a game to have a binary outcome, and a clearly defined subset of the team entities that contributed to the win or loss. A complete NHL game seems incompatible with this setting: players are subbed on or off the rink at will, and a game has a score as outcome (which can be converted to a win/loss binary outcome, but is richer than that); it is not immediately apparent which players contributed to the win or loss. However, the gap can be bridged by decomposing an NHL game into individual scoring plays; a goal scored corresponds to a win for the players currently on the ice, while a goal conceded corresponds to a loss. Hence, Exceptional Gestalt Mining can find synergy between NHL players on specific lines, measured in terms that are related to the standard hockey $+/-$ player performance metric. In future work, we plan to deploy EGM as is on such data.

Additionally, we intend to make intelligent use of more information that is present in the Kaldheim dataset (or can be derived from it), but which currently goes untouched. For instance, we know for each game which roles were observed from the deck of the opponent (possibly a subset of all colors in their deck); it stands to reason that certain coalitions have stronger gestalt when opposed by certain colors than when opposed by others. We also know how many turns each game took (gestalt in faster/slower decks[9]), and in what win rate bracket the player resides (gestalt for experienced/novice players); incorporating such information in EGM may uncover further interesting subgroups.

References

1. Rakesh Agrawal, Ramakrishnan Srikant: Fast Algorithms for Mining Association Rules in Large Databases. Proc. VLDB, pp. 487–499, 1994
2. Stephen D. Bay, Michael J. Pazzani: Detecting Group Differences: Mining Contrast Sets. Data Min. Knowl. Discov. 5(3):213–246, 2001
3. Robert A. Bosch: Optimal Card-Collecting Strategies for Magic: The Gathering. College Mathematics Journal 31(1):15-–21, 2000
4. Krishnendu Chatterjee and Rasmus Ibsen-Jensen: The complexity of deciding legality of a single step of Magic: The Gathering. Proc. European Conf. on AI, 2016
5. P. I. Cowling, C. D. Ward and E. J. Powley: Ensemble Determinization in Monte Carlo Tree Search for the Imperfect Information Card Game Magic: The Gathering, IEEE Trans. on Comp. Int. and AI in Games, 4(4):241–257, 2012
6. Alex Churchill, Stella Biderman, Austin Herrick: Magic: The Gathering is Turing Complete. arXiv preprint arXiv:1904.09828 [cs.AI], 2019
7. R. Conroy: Do the Bots Send Signals?. Available online at `https://www.17lands.com/blog/bot_signals`, last accessed June 16, 2021, published July 01, 2019
8. R. Conroy: Impact of the London Mulligan. Available online at `https://www.17lands.com/blog/london_mulligan`, last accessed June 16, 2021, published October 23, 2019
9. R. Conroy: Simulated Draft Strategies. Available online at `https://www.17lands.com/blog/simulating_draft_strategies`, last accessed June 16, 2021, published July 28, 2020
10. Guozhu Dong, Jinyan Li: Efficient Mining of Emerging Patterns: Discovering Trends and Differences. Proc. KDD, pp. 43–52, 1999
11. Xin Du, Yulong Pei, Wouter Duivesteijn, Mykola Pechenizkiy: Exceptional spatio-temporal behavior mining through Bayesian non-parametric modeling. Data Min. Knowl. Discov. 34(5):1267–1290, 2020
12. Wouter Duivesteijn, Ad Feelders, Arno J. Knobbe: Exceptional Model Mining - Supervised descriptive local pattern mining with complex target concepts. Data Min. Knowl. Discov. 30(1):47–98, 2016
13. Richard C. Garfield: Trading card game method of play. United States Patent US5662332A, United States Patent and Trademark Office, Sept. 02, 1997
14. Arnaud Giacometti, Arnaud Soulet: Dense Neighborhood Pattern Sampling in Numerical Data. Proc. SDM, pp. 756–764, 2018
15. Francisco Herrera, Cristóbal J. Carmona, Pedro González, María José del Jesus: An overview on subgroup discovery: foundations and applications. Knowl. Inf. Syst. 29(3):495–525, 2011
16. Willi Klösgen: Explora: A Multipattern and Multistrategy Discovery Assistant. Advances in Knowledge Discovery and Data Mining, pp. 249–271, 1996
17. Rob M. Konijn, Wouter Duivesteijn, Wojtek Kowalczyk, Arno J. Knobbe: Discovering Local Subgroups, with an Application to Fraud Detection. Proc. PAKDD (1):1–12, 2013

[9] Compare the discussion of the Raise the Draugr/Glittering Frost in Section 5.1.

18. Petra Kralj Novak, Nada Lavrac, Geoffrey I. Webb: Supervised Descriptive Rule Discovery: A Unifying Survey of Contrast Set, Emerging Pattern and Subgroup Mining. J. Mach. Learn. Res. 10:377–403, 2009
19. 17 Lands: KHM Traditional Draft Game Data. Available online at `https://17lands-public.s3.amazonaws.com/analysis_data/game-data.KHM.TradDraft.tar.gz`, last accessed June 17, 2021, published April 04, 2021
20. 17 Lands: Public Datasets. Available online at `https://www.17lands.com/public_datasets`, last accessed June 17, 2021
21. Dennis Leman, Ad Feelders, Arno J. Knobbe: Exceptional Model Mining. Proc. ECML/PKDD (2):1–16, 2008
22. David Lucking-Reiley: Using Field Experiments to Test Equivalence between Auction Formats: Magic on the Internet. American Economic Review 89(5):1063—1080, 1999
23. Brett A.S. Martin: Using the Imagination: Consumer Evoking and Thematizing of the Fantastic Imaginary. Journal of Consumer Research 31(June):136–149, 2004
24. Sierkovitz: 5000 Words for Snow. Available online at `https://www.17lands.com/blog/khm_snow_decks`, last accessed June 16, 2021, published February 18, 2021
25. Sierkovitz: Servant, not Master. Available online at `https://www.17lands.com/blog/using_wr_data`, last accessed June 16, 2021, published April 27, 2021
26. Stefan Wrobel: An Algorithm for Multi-relational Discovery of Subgroups. Proc. PKDD, pp. 78–87, 1997