Evolution of linguistic complexity in Hearthstone: a resource and an example in linguistic game studies

Jérôme Jacquin and Aris Xanthos
University of Lausanne

Abstract

The present paper uses corpus linguistics methods to study a born-digital cultural artifact, namely the collectible card game Hearthstone, released in 2014 and still under active development in 2020. As indicated by the developers' public statements, such a game, which is framed as an ongoing service rather than a final product, needs to strike a balance between two conflicting constraints: while the game must remain easily accessible at any stage of its development to increase the player base, the player's experience must be constantly renewed to keep the existing player base from getting bored and leaving the game. In practice, this translates into the regular addition of new card sets by the developer. Our hypothesis is that linguistic traces of these opposite forces can be found in the evolution of lexical diversity in the text of the game's rules, as opposed to its "flavour" text, which plays no role in the game's mechanics. To test this hypothesis, a corpus documenting the first two years of the game has been built and it is distributed in open access. The analysis of lexical diversity in these data shows that the rule text vocabulary is both highly controlled and subject to steady diversification. We view these features as reflecting stakes that characterise the ongoing management and development of a game-as-a-service like Hearthstone.

1. Introduction

In the current self-defining stage of digital humanities (DH), a number of scholars agree with the idea, notably discussed by Svensson (2009, 2010, 2012), that DH is not just a new name for humanities computing, but reflects a broadening of the scope of investigation, "a new focus or a different relation to traditional humanities computing work" (Svensson, 2009, p. 37). Arguably, most work that is subsumed under the DH umbrella nowadays can be predominantly related to one of two main research perspectives: the use of digital tools and methods for studying (typically digitised) cultural products–which corresponds to the historical meaning of humanities computing–, or the study of born-digital cultural artifacts with (mostly) traditional humanities and social sciences tools and methods. In this context, the present paper uses digital tools and methods to study a born-digital cultural artifact: the methodology is that of corpus and quantitative linguistics and the object of study is a

videogame. In particular, we will study the evolution of lexical diversity in the digital collectible card game Hearthstone.

Adopting a high-level view on the field of game studies, we find that contributions on linguistic aspects of videogames are relatively scarce, as compared to perspectives grounded in other disciplines such as computer science, pedagogy, medicine, or sociology (see Martin, 2018 for a large-scale review), and centered on questions related to artificial intelligence, educational practices (e.g. serious games), health issues and therapeutic uses, and gamers' social profiles for instance. The main topics of interest in the limited body of work conducted in a linguistic perspective appear to be language learning, game-mediated interaction between players, and videogame peritextuality (see e.g. Carrillo Masso, 2009; Losavio et al., 2014; Zagal et al., 2012; Zhu and Fang, 2015). In other words, game studies displaying an interest in language-related issues typically approach videogames as a learning environment, a communicative context, or a topic of discussion. In this context, the originality of the point of view adopted in this contribution is to envision *videogames as objects*–a perspective that still remains underrepresented in the field of game studies (linguistic or not), except maybe for typological purposes (i.e. classification of games into *genres*). Notable exceptions include early work by Thorne, Fischer and Lu (2012) on quest texts in World of Warcraft as well as some contributions in a book recently edited by Ensslin and Balteiro (2019).

Among the various aspects of a videogame's linguistic structure that could be considered (e.g. phonology, morphology, syntax), the present study focuses on the lexicon. Specifically, we will examine the lexicon of the texts of Hearthstone cards–texts which play a crucial part in the game's mechanics (see section 2.1)–in relation to the game's particular business model. Since its initial release in 2014, the basic version of Hearthstone is free but new cards are periodically released and players are encouraged to buy them. In this context, the game must strike a balance between two conflicting constraints: on the one hand, the game's mechanics must remain simple enough for new players to join as easily as possible at any time; on the other hand, the player's experience must be constantly renewed to keep the existing player base from getting bored and leaving the game. Our hypothesis is that traces of these opposite constraints can be found in the evolution of card texts over the game's lifetime. In particular, we consider the evolution of lexical diversity in these texts as an indicator of the evolution of the game's complexity.

This hypothesis echoes one of the major concerns of the game's developer, who has regularly and publicly discussed the need to strike a balance between simplification and diversification. For example, when being asked by Gamespot's journalist Rob Crossley about how the developer deals with the fact that "the game is going to get harder and harder to get into", Hearthstone executive producer Hamilton Chu answers:

> Yeah, that's a really good question, and it really hits on something we think about [...] There's a few different ways we try to manage it, and I think "manage"" is a good word for it. Something we think a lot about when adding new cards is making sure we don't add a lot of complexity. So even on individual cards, we count the number of words we use on a card, and if it's a high number, we actually try to find a way of phrasing it a little more elegantly.[1]

Reducing the number of words by rephrasing the text of a card is not the only linguistic strategy that Hearthstone developers use to control the increase in complexity. They also use short "keywords" that summarise information (see Yucel 2016 on the notion of "chunking" adapted to rule writing in game design). For example "Deathrattle" means that the death of a creature will activate a special effect. But as stated by Hearthstone Senior Game Producer Yong Woo, already in mid-2014:

> Keywords are double-edged swords. On the one hand, it really condenses information. I tell you *Deathrattle* and you know what it is right away. I tell somebody off the street *Deathrattle* and he is going to be like "What are you... selling me a snake? ". They will not know what I am talking about, so it is a balance we have to reach between condensing the information and making too hard a set of vocabulary for new players to understand. It is something we continue to work on.[2]

As part of an effort to study these issues in an empirical fashion, we have built a corpus containing the texts of all Hearthstone cards that have been released during the 20 first months of the game's development, a time period which corresponds to a particular regime in the developer's management of the game (see section 2.1). These texts have been linguistically analysed using a taylor-made, fine-grained annotation scheme, and the resulting corpus is made publicly available under an open access license. The next section of this paper is devoted to the presentation of this resource, including an introduction to the relevant features of the game. Section 3 presents the method, results, and discussion of our analysis of the evolution of linguistic complexity in these data. The final section summarises our findings and discusses further lines of research in the emerging field of linguistic game studies.

2. Resource: the Blizzword corpus

2.1 Context: Hearthstone

Developed and published by Blizzard Entertainment, Hearthstone is a free-to-play online digital collectible card game, available on desktop as well as mobile devices. It was released in March 2014 and is still under active development in 2020. Success has been massive since the beginning; while 10 million players were reported during the first month, they were 30 million in May 2015 and the last report issued in November 2018 stated that the 100 million players milestone had been reached.

A Hearthstone match is a turn-by-turn confrontation between two players or between a player and the computer. Each side uses a deck of 30 cards selected in the player's personal collection before the match. Winning the match depends on several factors, notably luck (since the player's cards are drawn in random order at each turn and the opponent's actions are only partially predictable) and skill (knowledge of game mechanics, cards, favoured strategies, deck building, etc.), but also–most importantly for this study–the particular set of cards available in the player's collection. About the latter, Hearthstone adopts a free-to-play

economic model, where players may develop their card collection using in-game currency ("gold") acquired by completing various quests and achievements. However, they are frequently encouraged to buy new packs of cards using real money and thereby increase the power level of their decks faster (see Maisenhölder, 2018 and Švelch, 2019 for discussions of the economic incentive to continuously update competitive decks in *Magic: the Gathering*, Hearthstone's oldest and most famous analog ancestor). The game thus operates as an ongoing service rather than a one-time released product (see e.g. Davidovici-Nora, 2013); for it to be economically profitable, the developer must at the same time retain previous adopters by regularly adding new cards and attract new players, who will discover the game from scratch, possibly years after the initial release.

In this study, we focus on the evolution of the game between its release in 2014 and the end of 2015. During this initial period, cards were only added to the game, whereas from early 2016 on, cards were also removed periodically–at least in the so-called "standard" mode, most popular among players due notably to its adoption for official competitions.[3] We chose the moment of this radical shift in the game's evolution as the end point of our data collection in order to focus on the strictly cumulative mode initially adopted by the developers. At the game's release, only two sets of cards were available: the "basic" set comprises 133 collectible cards, out of which 43 are immediately available to every player, the remaining 90 cards being obtained as rewards for reaching specific levels of experience.[4] The "classic" set is the largest set of the game, originally consisting of 245 collectible cards. Classic cards are mainly obtained by opening virtual card packs, each of which contains 5 random cards; every card type belongs to a specific "rarity" class, ranging from "common" to "legendary", so that acquiring some of them may require opening a large number of packs. Packs themselves can be obtained as rewards for certain quests and achievements, but they are most often purchased using either gold earned from other achievements or real money.

Besides the two core sets, five sets were added to the game until the end of 2015. Three of those were "adventures" ("Curse of Naxxramas", "Blackrock Mountain", and "League of Explorers"): cards in these smaller sets (30-45 collectible cards) could only be obtained by taking a number of single-player challenges, and their acquisition was practically guaranteed for every player having paid the adventure's fee (using gold or real money). Between these adventures, two larger "expansion" sets were released ("Goblins and Gnomes", 123 collectible cards, and "The Grand Tournament", 132 collectible cards). Cards in these sets were obtained in the same way as classic cards, i.e. by opening packs of the expansion in question, with the difference that these packs had to be purchased (they were not awarded for quests and achievements). Table 1 summarises the release of card sets during the initial, strictly incremental phase of the game development, in which we are specifically interested here.

Table 1: Summary of Hearthstone card set releases considered in this study

Code	Name	Release	Type	# Cards
BAS	*Basic*	March 11, 2014	Core	133

CLA	*Classic*	March 11, 2014	Core	245
NAX	*Curse of Naxxramas*	July 22, 2014	Adventure	30
GVG	*Goblins and Gnomes*	December 8, 2014	Expansion	123
BRM	*Blackrock Mountain*	April 2, 2015	Adventure	31
TGT	*The Grand Tournament*	August 24, 2015	Expansion	132
LOE	*League of Explorers*	November 12, 2015	Adventure	45
TOTAL				739

Cards are the most fundamental objects in Hearthstone. In each turn of a match at least one card is randomly taken from the player's 30 card deck and added to the player's hand. There are different types of cards, such as minions, spells, and weapons. Cards can have several numeric characteristics such as attack points and health points. Most of them also display textual information–rule text–, that is of crucial interest for our study. As an example, the rule text printed on the minion card named "Murloc Tidehunter"[5] is "Battlecry: Summon a 1/1 Murloc Scout". "Battlecry" is a keyword which means that a specific effect will be triggered when the card is played (in this case, a minion called *Murloc Scout* and having one attack point and one health point, will appear on the battlefield).

Each individual card is also associated with a "flavour" text, which is not visible during the match but which can be viewed in the player's card library, where all their cards are stored. The function of this text differs crucially from the rule text. Indeed, the flavour text does not give any information about the card's mechanics, nor does it have any impact on gameplay. It is used to relate the card to Hearthstone's cultural background, notably the developer's Warcraft franchise. In the case of "Murloc Tidehunter", the flavour text is "Death will rise, from the tides!", which is taken from a song performed by a band of non-playing characters in the very popular multiplayer game World of Warcraft.

2.2 Corpus constitution

The corpus we have built for this study contains the rule and flavour texts of all 739 collectible cards originally published in English between Hearthstone's public release in March 2014 and the end of 2015, i.e. shortly before the developer introduced the "standard" mode and started removing card sets from the game (see Table 1). The card data, which include not only their rule and flavour texts, but also a number of other indications such as

the set they are part of, their cost, attack, health, rarity, and so on, were initially retrieved from the HearthstoneJSON website[6] in November 2016.

After converting the original JSON format to XML and applying standard preprocessing steps, we used regular expressions to produce an initial segmentation of rule texts.[7] This enabled us to establish a distribution of about 400 form types, each of which was subsequently assigned a default lemma and part-of-speech (POS) tag–or marked as resulting from a tokenisation error. Based on this, we used Textable (Xanthos, 2014) to tag the tokenised texts, then proceeded to review them manually and craft more than a hundred ordered contextual rules belonging to one of two types: fusion or recategorisation. Fusion rules address tokenisation errors and replace a sequence of successive tokens with a single one, e.g. “if two successive tokens have lemmas *Murloc* and *Scout*, replace them with a single *Murloc Scout* token, that has lemma *Murloc Scout* and POS tag *n|sg* (singular noun)”[8]. Recategorisation rules address errors resulting from default lemma or POS tag assignment, e.g. “if a token that has lemma *draw* is preceded with a token that has lemma *you*, change its POS tag from *v|imp* (verb imperative) to *v|ind_pres* (verb indicative present)”. Using Textable again, we designed a workflow to convert these rules into regular expressions then apply them to rule texts, before manually double-checking each token and making the necessary corrections.

The workflow mentioned above could not be directly applied to flavour texts: indeed, not only did flavour texts contain about 40% more tokens than rule texts (when segmented with the same regular expression), but they also contained 7 times more form types. We therefore decided to use TreeTagger (Schmid, 1994) to produce an initial analysis of the flavour texts, then establish conversion rules between the TreeTagger categories and the scheme designed for rule texts, and use Textable to apply the conversions to the TreeTagger output. Finally, each token was manually reviewed twice, first by a student assistant, then by the authors. In the process, it became necessary to apply a number of changes to the rule text analysis, in cases where new ambiguities had arisen in flavour texts.

The result is a corpus of 755 cards, 743 of which are collectible. Table 2 summarises basic quantitative facts about the 720 cards which have a rule text and the 742 which have a flavour text.

Table 2 Global statistics concerning the corpus considered in this study

	# tokens	# types		
		forms	lemmas	POS tags
Rule texts	7926	446	355	62
Flavour texts	11281	3171	2593	78
Total	19207	3382	2727	81

Below is an example of the XML encoding of an individual card, namely “Light’s Champion”. Note that all characteristics of the card are stored as XML attributes, while rule

text “Battlecry: Silence a Demon” and flavour text “When there’s something strange (say, a gibbering demon) in your neighborhood, who are you going to call?” are enclosed in XML elements (*<rules>* and *<flavor>*):[9]

```
<card name="Light’s Champion" num="261" id="AT_106" set="6_TGT" cost="3"
attack="4" health="3" rarity="RARE" category="MINION" mechanics="BATTLECRY"
playRequirements="TARGET_IF_AVAILABLE:0 MINION_TARGET:0 TARGET_WITH_RACE:15"
targetingArrowText="&lt;b&gt;Silence&lt;/b&gt; a Demon.">
  <rules>
    <token lemma="Battlecry" pos-tag="n|sg" rend="bold">Battlecry</token>
    <token lemma=":" pos-tag="pun|colon" rend="bold">:</token>
    <token lemma="Silence" pos-tag="v|imp" rend="bold">Silence</token>
    <token lemma="a" pos-tag="det|indef">a</token>
    <token lemma="Demon" pos-tag="n|sg">Demon</token>
    <token lemma="." pos-tag="pun|period">.</token>
  </rules>
  <flavor>
    <token lemma="when|conj" pos-tag="conj|sub">When</token>
    <token lemma="there|pro" pos-tag="pro">there</token>
    <token lemma="be|v" pos-tag="v|ind_pres">’s</token>
    <token lemma="something" pos-tag="pro">something</token>
    <token lemma="strange" pos-tag="adj">strange</token>
    <token lemma="(|pun" pos-tag="pun|par">(</token>
    <token lemma="say" pos-tag="v|imp">say</token>
    <token lemma="," pos-tag="pun|comma">,</token>
    <token lemma="a" pos-tag="det|indef">a</token>
    <token lemma="gibber" pos-tag="v|par_pres">gibbering</token>
    <token lemma="Demon" pos-tag="n|sg">demon</token>
    <token lemma=")|pun" pos-tag="pun|par">)</token>
    <token lemma="in|prep" pos-tag="prep">in</token>
    <token lemma="your" pos-tag="det|poss">your</token>
    <token lemma="neighborhood" pos-tag="n|sg">neighborhood</token>
    <token lemma="," pos-tag="pun|comma">,</token>
    <token lemma="who|int" pos-tag="pro|int">who</token>
    <token lemma="be|vaux" pos-tag="vaux|ind_pres">are</token>
    <token lemma="you|subj" pos-tag="pro|subj">you</token>
    <token lemma="go|vaux" pos-tag="vaux|par_pres">going</token>
    <token lemma="to|mark" pos-tag="mark|inf">to</token>
    <token lemma="call|v" pos-tag="v|inf">call</token>
    <token lemma="?" pos-tag="pun|int">?</token>
  </flavor>
</card>
```

2.3 Availability

In the hope of fostering research in the emerging field of linguistic game studies, the corpus described above and studied in the next section is released publicly under the CC BY-SA

4.0 license. It can be freely downloaded in XML format from https://github.com/axanthos/blizzword.

3. Case study: evolution of linguistic complexity

The central research question of this study bears on the complexity of Hearthstone rule texts versus flavour texts, which are used for comparative purposes. The hypothesis is that, while these texts are part of the same game, they are differently affected by the ludic and economic constraints stated in the introduction. Phrased more precisely, our research question is the following: how and to what extent does the linguistic complexity of rule texts–in contrast to the flavour texts, which are not subjected to the same constraints–reflect the tension existing between the following contradictory objectives:

a. Retaining the existing player base by the regular addition of new cards, new keywords, and new rules, and therefore complexity; and
b. Keeping the game accessible for new players, who can join the Hearthstone community at any time during the game's ongoing development.

In the following sections, we present the method elaborated for addressing this question (3.1), summarise the results of our analysis (3.2), and discuss how they relate to the aforementioned game design objectives (3.3).

3.1 Method

We have chosen to assess the textual complexity of Hearthstone card texts in terms of lexical diversity–or more specifically, lexematic diversity[10]. To that effect, we have used an index that has been extensively discussed by McCarthy and Jarvis (2010, 2007) and which can be interpreted as the expected vocabulary size in subsamples of a fixed number of tokens randomly drawn from the sample(s) under consideration.[11] The main argument for using this particular measure of lexical diversity is that it is one of those that have been found to be most robust with regard to variations of sample size (McCarthy and Jarvis, 2010; Xanthos, 2013). In addition, the calculation of expected vocabulary size does not rely on particular hypotheses about the process by which the data have been generated, and the resulting measurement has a straightforward interpretation.

We are concerned in particular with the evolution of lexematic diversity over the course of card set releases (see table 1). That card sets are being used as context units in this analysis reflects the ideal player model adopted in this study: indeed, we make the simplifying assumption that whenever a card set is being released, the player is exposed to the rule and flavour texts of all cards of this set at once–an assumption that will be further discussed in Conclusion below. Although the "basic" and "classic" sets were published at the same time (at the game's initial release), they are also treated as successive data points, on the grounds that players are exposed to dozens of basic cards before they have the possibility to encounter or acquire their first classic cards.

Also, it is important to note that we are adopting a cumulative perspective on the data: while the first data point corresponds to the 133 cards of the basic set, the second one covers the 245 cards of the classic set *plus* the 133 basic cards, and so on. This is intended to approximate the experience of a player during the first two years of the game, when changes brought to the card pool by the developer were strictly additive. During that period, in effect, players needed to know an ever-increasing amount of cards and their texts to remain competitive. To make this methodological choice explicit, the label of each data point but the first will be prefixed with a rightwards arrow symbol ("→CLA", "→NAX", and so on) in the following data visualisations.

Based on these definitions, the least number of tokens associated with a data point is 918 (in the rule texts of the "basic" set), so the subsample size parameter involved in the calculation of expected vocabulary size has been set to 500 tokens. This ensures that even the lexematic diversity of the smallest data point can be estimated on the basis of a reasonably large amount of tokens (corresponding to the rule texts of about 50 cards).

The expected cumulated vocabulary size can be used for comparing the lexematic diversity of rule and flavour texts at each point in the timeline of our data (i.e. each card set release), however it can be less convenient for comparing the evolution of lexematic diversity in the two subcorpora. Indeed, if their overall lexematic diversity is substantially different, which our hypotheses predict, differences in their evolution over time can be harder to perceive due to scale differences. For this reason, we will consider not only the size of the expected cumulated vocabulary but also the increase of this size relative to its value at a reference point in the timeline. This makes it possible to compare the dynamics of lexematic diversity in the two subcorpora on a single, unified scale.

3.2 Results

Figures 1 and 2 respectively show the cumulated number of tokens and the evolution of raw cumulated vocabulary size in the rule and flavour text subcorpora.

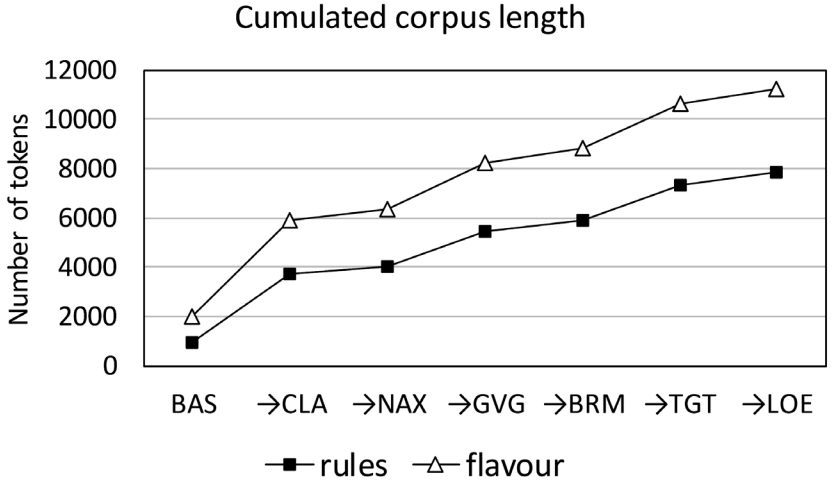

Fig. 1 Evolution of cumulated corpus length of rule and flavour texts

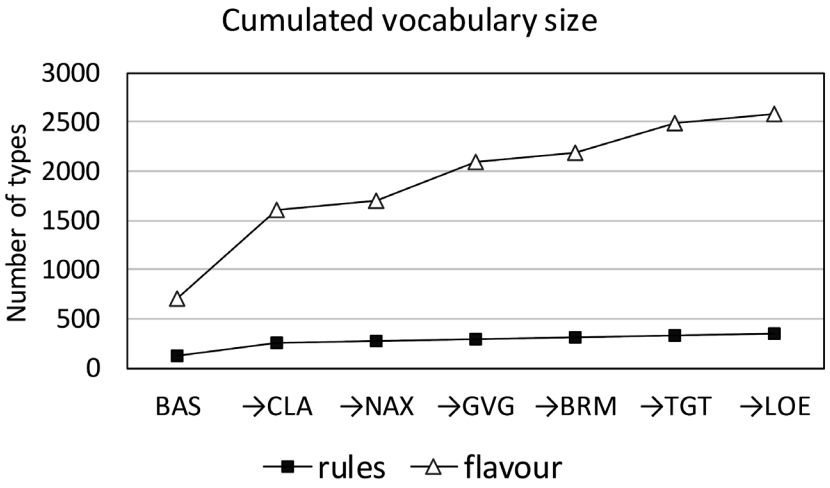

Fig. 2 Evolution of raw cumulated vocabulary size of rule and flavour texts

It must be stressed that the perceivable differences in raw numbers of lemma types do not lend themselves to a direct comparison, mainly for two reasons: (i) because the flavour text subcorpus is considerably larger than the rule text subcorpus (see table 2 and Figure 1) and (ii) because the cumulative approach adopted in this paper implies that the number of tokens is increasing over time within each individual subcorpus, which in turn entails that the raw number of types increases (all other things being equal). Figure 3 attempts to compensate for these shortcomings by displaying the expected cumulated vocabulary size in subsamples of 500 tokens.

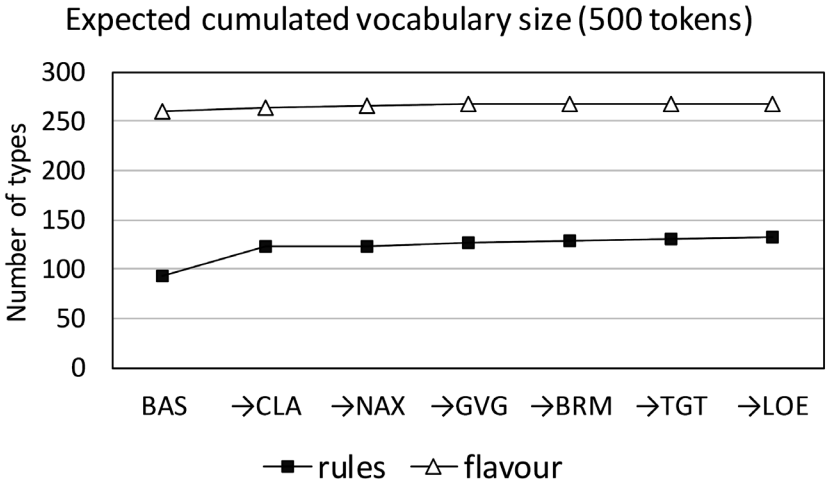

Fig. 3 Evolution of expected cumulated vocabulary size of rule and flavour texts (subsamples of 500 tokens)

Yet another picture arises when considering not the size of the expected cumulated vocabulary, but the increase of this size relative to the first release (BAS), as in Figure 4.

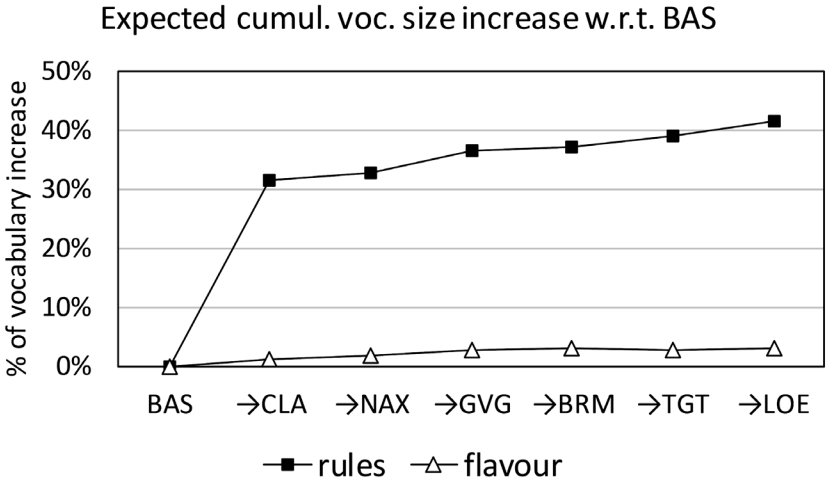

Fig. 4 Expected cumulated vocabulary size increase relative to the BAS card set (subsamples of 500 tokens)

Finally, Figure 5 shows how the resampled cumulated vocabulary increases relatively to the situation after the second core set release, thus focusing on what happens after the jump from BAS to CLA (which is imputable to the pedagogical function of BAS).

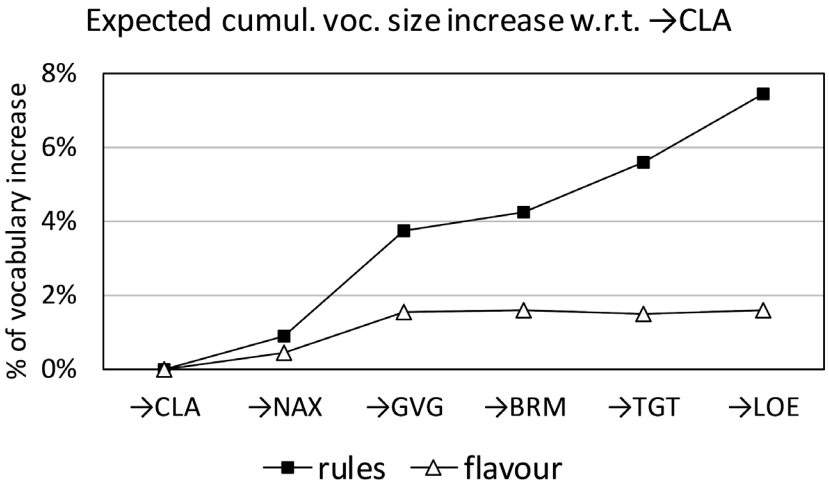

Fig. 5 Expected cumulated vocabulary size increase relative to the union of BAS and CLA card sets (subsamples of 500 tokens)

3.3 Discussion

As shown in Figure 1, while both the rule text curve and the flavour text curve exhibit a pattern that reflects the alternance between smaller sets (adventures NAX, BRM and LOE) and larger ones (core CLA, as well as expansions GVG and TGT), the number of tokens in flavour texts is consistently larger than in rule texts. On average, a card's flavour text is almost 40% larger than its rule text (15.2 vs 11 tokens). The differences in vocabulary size

between the two subcorpora are even larger than the differences in number of tokens, and they increase over time (see Figure 2). The gap between the number of lemma types in flavour and rule texts is already quite large in the BAS set (710 vs 120 lemma types), and the cumulated vocabulary is almost 8 times larger in flavour texts than in rule texts at the end of the observation period (2786 vs 350 lemma types).

As mentioned in the previous section, the expected cumulated vocabulary size in subsamples of fixed size (here 500 tokens) is a more robust way of comparing the two subcorpora than the raw number of lemma types. Examining this measure of lexical diversity in Figure 3 confirms that there is a large difference in lexical diversity (never lesser than a 200% factor) between rule and flavour texts. However, it also reveals that (i) there is a substantial increase of rule text vocabulary between core sets BAS and CLA–which, again, is imputable to the pedagogic role of the former–, and (ii) the difference in lexical diversity between flavour and rule texts does not increase over time as suggested by Figure 2. Rather, there appears to be an almost constant ratio of lexical diversity in flavour vs rule texts from the release of CLA to the end of the observation period (slightly more than 2:1 on average).

Rule texts are thus consistently characterised by a much more controlled vocabulary than flavour texts. Being one of the core mechanics of the game, rules are phrased in a consistent way, repetitively using a comparatively small set of lexical items (i.e. vocabulary). This reflects the objective for the game to remain easily accessible at any stage of its development: a new player joining the game experience one year after the initial release is not faced with a very different and much more complex game than the early adopters. Conversely, the flavour text vocabulary is far less controlled and shows no tendency to economy. This echoes the fact that flavour texts, which are used to ground the game in the fictional world previously built by the developer in one of its most famous franchises (i.e. Warcraft), have no impact on the gameplay.

Adopting a relative perspective on the same data (relative to the initial situation represented by the BAS card set), we find that the lexical diversity of rule texts of BAS is separated from that of CLA by a gap whose magnitude (+32%) is never matched in subsequent releases (see Figure 4). Most interestingly, while the increase is more gradual from this point on, it remains clearly visible until the end of the observation period, while the lexical diversity of flavour texts increases more slowly and reaches a plateau from GVG onwards. This is even more clearly shown in Figure 5, which focusses on what happens after the jump from BAS to CLA and confirms that the cumulated resampled vocabulary of rule texts keeps growing steadily with each new release, while the lexical diversity of flavour texts remains approximately constant in the last four releases of the observation period. In other words, while the vocabulary of rule texts is much more controlled than the vocabulary of flavour texts, it diversifies considerably faster. This reflects the progressive implementation, extension after extension, of new categories and new rules in the game, in order to renew–although still in a very controlled way–the experience for the existing player base.

4. Conclusion

The present paper started with a general observation: while videogames are increasingly becoming a topic of interest for research, the linguistic dimension of videogames envisioned as artifacts remains largely underexplored. Our study aims to illustrate the potentiality of linguistic game studies by providing an original resource and exploring it using digital tools and methods adapted to our hypothesis and to the characteristics of the considered game.

Our results demonstrate that the rule and flavour texts of cards in Hearthstone differ crucially when considered through the lens of lexical diversity. While flavour texts are much richer in terms of vocabulary, rule texts are characterised by a steady increase in lexical diversity. We view these differences as linguistic traces of two conflicting constraints that characterise the ongoing management and development of a game-as-a-service like Hearthstone: remaining accessible for new players by staying the same while retaining the actual players by enriching the gaming experience.

From here, a number of suggestions for further lines of research can be made, notably in relation with limitations of the present study. First, the corpus should be enriched by taking into account what happens after the last extension considered in this study, i.e. when the developer started to *remove* cards periodically, at least in the so-called “standard” mode. Could we observe linguistic traces of this shift in the way the evolution of the game is managed? To make the process of extending the corpus to the numerous extensions released after 2015 more efficient, the data previously analysed could be leveraged for training a POS-tagging system. The corpus could also be diversified by applying the same methodology to other (types of) games, including tabletop card games such as *Magic: the Gathering*, other digital card games such as *Gwent: The Witcher Card Game*, or even other types of evolutive videogames.

Regarding the language analyzed, Hearthstone was originally published in English but it now exists in a dozen of languages, including not only other Indo-European varieties such as German, French or Italian, but also translations into Chinese and Korean. This linguistic diversity opens up interesting cross-linguistic explorations, notably on the way lexical diversity is managed in other and possibly richer morphological systems than English.

Another line of research would consist in looking at other linguistic variables than lexical diversity. For example, a syntactic analysis using our POS-tagging scheme could provide relevant insights about another layer of complexity, that can be compared with our observations on the vocabulary of rule and flavour texts. Additionally, while the present study has been exclusively concerned with the “set” attribute of Hearthstone cards, a host of other attributes, such as “cost”, “rarity”, “health points” or “category”, are available and could also be readily exploited.

Multimodality should also be addressed, although tackling this dimension is both theoretically and methodologically challenging: in a videogame such as Hearthstone, information about gameplay is provided by other means than verbal phrasing (e.g. visual feedback such as change of positions, forms, colors; see Elias et al. 2012). Such information could be exploited in order to assess the extent to which multimodal traces of complexity reflect linguistic ones.

The last perspective we will discuss is probably the most ambitious. As stated above, our study reflects the trajectory of an ideal player by making the simplifying assumption that they are simultaneously exposed to the rule and flavour texts of all the cards released in each given set, one set after another. While this assumption seeks to adopt the developer's perspective on the game, examining more diverse trajectories through the game complexity, for example based on computer simulations or on ethnographic studies, would provide relevant and complementary insights on the complexity of the game considered as an actual practice. This line of research would also benefit from the massive Hearthstone gameplay data continuously crowdsourced and analysed by such organisations as HearthSim (https://hearthsim.info) and Vicious Syndicate (https://www.vicioussyndicate.com). These data provide the necessary empirical basis for operationalising the notion of "metagame" (cf. Garfield, 1995, Boluk and LeMieux, 2017) as applied to Hearthstone, for which the following explanation, written by Švelch (2019, p.6) about *Magic: the Gathering*, is equally relevant:

> any match can be considered a part of a series and thus should be approached with the information about previous matches in mind. In a tournament setting, this means that a player is not building the best deck in absolute terms (as this should be impossible if the game is supposed to be properly balanced), but a deck that is favored against the other popular decks, which constitute the current metagame.

As a consequence of the existence of an emerging metagame at any point of the game's history, the actual number of cards that a player interacts with on a regular basis is considerably lesser than the number of existing cards (even limited to those that constitute the standard format), a phenomenon which the aforementioned crowdsourced gameplay data would enable us to model to some extent.

In conclusion, by presenting a resource and a case study, our paper seeks to open a discussion on the relevance of linguistic game studies as an emergent field, where linguistics can provide insights about videogames mechanics and design while videogames' textuality appears as a particularly interesting locus to revisit traditional topics in linguistics, such as lexical diversity and complexity.

Acknowledgements

The authors are most grateful to the Faculty of Arts of the University of Lausanne for the financial support and to Laura Schött for her participation in the corpus annotation. They would also like to thank the two anonymous reviewers for their relevant comments on previous versions of this paper.

References

Boluk, S., LeMieux, P. 2017. *Metagaming: Playing, Competing, Spectating, Cheating, Trading, Making, and Breaking Videogames*. Minneapolis, MN: University of Minnesota Press.

Carrillo Masso, I., 2009. Developing a methodology for corpus-based computer game studies. *J. Gaming Virtual Worlds* 1, 143–169. https://doi.org/10.1386/jgvw.1.2.143/7

Davidovici-Nora, M., 2013. Innovation in business models in the video game industry: Free-To-Play or the gaming experience as a service. *Comput. Games J.* 2, 22–51. https://doi.org/10.1007/BF03392349

Elias, G. S., Garfield, R. and Gutschera K. R., 2012. *Characteristics of Games*. Cambridge, MA: MIT Press.

Ensslin, A., Balteiro, I. (Eds.), 2019. *Approaches to Videogame Discourse. Lexis, Interaction, Textuality*. Bloomsbury, London-Oxford.

Garfield, R., 1995. Lost in the shuffle: games within games. *The Duelist: The Official Deckmaster Magazine,* 5 June, 86–88.

HearthSim Developer Community, 2016. *HearthstoneJSON: JSON files for Hearthstone*, retrieved from https://hearthstonejson.com.

Losavio, A., Polyakova, S., Hayden, T., Losavio, M., 2014. Linguistic implementations in computer game and virtual world design, in: 2014 Computer Games: AI, Animation, Mobile, Multimedia, Educational and Serious Games (CGAMES). Presented at *the 2014 Computer Games: AI, Animation, Mobile, Multimedia, Educational and Serious Games (CGAMES)*, pp. 1–4. https://doi.org/10.1109/CGames.2014.6934142

Maisenhölder, P. 2018., Why Should I Play to Win If I Can Pay to Win? Economic Inequality and Its Influence on the Experience of Non-Digital Games. *Well Played* 7 (1), 60–83.

Martin, P., 2018. The Intellectual Structure of Game Research. *Game Stud.* 18.

McCarthy, P.M., Jarvis, S., 2010. MTLD, vocd-D, and HD-D: A validation study of sophisticated approaches to lexical diversity assessment. *Behav. Res. Methods* 42, 381–392. https://doi.org/10.3758/BRM.42.2.381

McCarthy, P.M., Jarvis, S., 2007. vocd: A theoretical and empirical evaluation. *Lang. Test.* 24, 459–488. https://doi.org/10.1177/0265532207080767

Schmid, H., 1994. Probabilistic Part-of-Speech Tagging Using Decision Trees, in: *Proceedings of International Conference on New Methods in Language Processing*. Manchester, UK.

Serant, D., 1988. A propos des modèles de raccourcissements de textes, in: Thoiron, P., Labbe, D., Serant, D. (Eds.), *Études sur la richesse et la structure lexicale*. Champion-Slatkine, Paris-Genève, pp. 43–66.

Švelch, J. 2019. Mediatization of a Card Game: Magic: The Gathering, Esports, and Streaming. *Media, Culture & Society*, October, Online First. https://doi.org/10.1177/0163443719876536.

Svensson, P., 2012. The digital humanities as a humanities project. *Arts Humanit. High. Educ.* 11, 42–60. https://doi.org/10.1177/1474022211427367

Svensson, P., 2010. The Landscape of Digital Humanities. *Digit. Humanit. Q.* 004.

Svensson, P., 2009. Humanities Computing as Digital Humanities. *Digit. Humanit. Q.* 003.

Thorne, S.L., Fischer, I., Lu, X., 2012. The semiotic ecology and linguistic complexity of an online game world. *ReCALL* 24, 279–301. https://doi.org/10.1017/S0958344012000158

Xanthos, A., 2014. Textable: programmation visuelle pour l'analyse de données textuelles, in: *Actes des 12èmes journées internationales d'analyse statistique des données textuelles (JADT 2014)*. pp. 691–703.
Xanthos, A., 2013. L'évaluation (de l'évaluation)+ de la diversité lexicale, in: Prikhodkine, A., Xanthos, A. (Eds.), *Mélanges offerts en hommage à Remi Jolivet*. Centre de linguistique et des sciences du langage, Lausanne, pp. 231–252.
Yucel, I., 2016. Rules for Writing Rules: How Instructional Design Impacts Good Game Design, in: Trammell, A., Torner, E. and Leigh Waldron, E. *Analog Game Studies: Volume I,* ETC Press, Pittsburgh, pp. 167-173.
Zagal, J.P., Tomuro, N., Shepitsen, A., 2012. Natural Language Processing in Game Studies Research: An Overview. *Simul. Gaming* 43, 356–373. https://doi.org/10.1177/1046878111422560
Zhu, M., Fang, X., 2015. A Lexical Approach to Study Computer Games and Game Play Experience via Online Reviews. *Int. J. Human–Computer Interact.* 31, 413–426. https://doi.org/10.1080/10447318.2015.1036228

[1] *Hearthstone Grand Tournament Interview:* "*I Don't Like Playing Aggro Decks Either*", Interview of Hamilton Chu by Rob Crossley (Gamespot), issued August 21st, 2015 (https://www.gamespot.com/articles/hearthstone-grand-tournament-interview-i-dont-like/1100-6429856, last accessed September 2nd, 2019).

[2] *GamesCom Interview: Hearthstone Observer Mode Coming Soon*, Interview of Senior Game Producer Yong Woo by Icy-Veins, issued August 25th, 2014 (https://www.icy-veins.com/forums/topic/6827-gamescom-interview-hearthstone-observer-mode-coming-soon/, last accessed September 2nd, 2019).

[3] Hearthstone's "standard" format was introduced in early 2016. At each point in time, this game mode is restricted to the card sets issued during the previous two calendar years, in addition to the two card sets originally shipped with the game ("basic" and "classic", see below). The so-called "rotation", i.e. the process of removing older sets from the standard format when the first new set of a given year is being issued, operates as a distinct mechanism contributing to the aforementioned balance between accessibility for new players and engagement of existing players.

[4] This set has an explicit pedagogical function, notably reflected in the fact that about 10% of its cards have no text, against less than 2% in the other card sets mentioned below. For a new player, moving from the basic to the next set is part of the "heuristics" of the game (Elias et al. 2012, pp. 29 and ss.), i.e. the process through which players implicitly or explicitly tackle the gameplay and game difficulty.

[5] Like many others, this card can be viewed on the game's official website (https://playhearthstone.com/en-us/cards/976-murloc-tidehunter?collectible=1&set=standard&textFilter=murloc%20scout).

[6] These data "are automatically converted from the game files" (HearthSim Developer Community, 2016). It is worth noting that the data we have been using are a snapshot of the game's textual contents as some point in time, and as such they abstract away from various changes that these contents have undergone before or after that particular moment (many such changes have notably occurred during the game's beta period). A more precise model of Hearthstone's linguistic evolution might consider the way in which individual card texts have evolved over the course of the game's development.

[7] A token was defined as a sequence of letters, a sequence of numbers, or any of the following characters: *+ - / % " . : , ; ' ( ).*

[8] The categorisation of capitalised words was a challenge regularly discussed among annotators. On the one hand, some tokens are easily analysed as proper nouns referring to specific and unique entities in the game's lore (e.g. *Antonidas* or *The League of Explorers*). On the other hand, many

words or expressions, like *Murloc Scout*, are capitalised without being actual proper nouns; they can be names for tribes (e.g. *Mechs*), races (e.g. *Night Elves*), hero classes (e.g. *Warrior*), or cards (e.g. *Molten Giant*). They were categorised as common nouns but lemmatised by keeping the capitalised letter(s).

[9] Comparing the rule and flavour text of this card immediately reveals a number of differences: in particular, the flavour text is much longer, it has a much more diverse vocabulary, and the presence of a vertical bar in many flavour text lemmas indicates that they have been disambiguated (e.g. to distinguish “when” as a conjunction from its adverbial and pronominal uses).

[10] i.e. the diversity of lemmas, such as the verb TO DEAL, rather than forms, such as *dealt* or *dealing*. Note that punctuation and other symbols whose interpretation is needed to understand the rules (e.g. slash in expressions such as “2/1”, referring to a minion with two attack points and one health point) are included in the count of lemmas.

[11] The calculation of this index, which is based on the hypergeometric law, has been described in French by Serant (1988).