

Unravelling Names of Fictional Characters

Katerina Papantoniou

Institute of Computer
Science, FORTH
Heraklion, Greece
papanton@ics.forth.gr

Stasinou Konstantopoulos

Institute of Informatics &
Telecommunications, NCSR ‘Demokritos’
Ag. Paraskevi 153 10, Athens, Greece
konstant@iit.demokritos.gr

Abstract

In this paper we explore the correlation between the sound of words and their meaning, by testing if the polarity (‘good guy’ or ‘bad guy’) of a character’s role in a work of fiction can be predicted by the name of the character in the absence of any other context. Our approach is based on phonological and other features proposed in prior theoretical studies of fictional names. These features are used to construct a predictive model over a manually annotated corpus of characters from motion pictures. By experimenting with different mixtures of features, we identify phonological features as being the most discriminative by comparison to social and other types of features, and we delve into a discussion of specific phonological and phonotactic indicators of a character’s role’s polarity.

1 Introduction

Could it be possible for fictional characters’ names such as ‘Dr. No’ and ‘Hannibal Lecter’ to be attributed to positive characters whereas names such as ‘Jane Eyre’ and ‘Mary Poppins’ to negative ones? Could someone guess who is the hero and who is the competitor based only on the name of the character and what would be the factors that contribute to such intuition? Literary theory suggests that it should be possible, because fictional character names function as expressions of experience, ethos, teleology, values, culture, ideology, and attitudes of the character.

However, work in literary theory, psychology, linguistics and philosophy has studied fictional names by analysing individual works or small clusters of closely related works, such as those of a

particular author. By contrast, we apply tools from computational linguistics at a larger scale aiming to identify more general patterns that are not tied to any specific creator’s idiosyncrasies and preferences; in the hope that extracting such patterns can provide valuable insights about how the sound of names and, more generally, words correlates with their meaning.

At the core of our approach is the idea that the names of fictional characters follow (possibly subconsciously) a perception of what a positive or a negative name ought to sound like that is shared between the creator and the audience. Naturally the personal preferences or experiences of the creator might add noise, but fictional characters’ names will at least not suffer (or suffer less) from the systematic cultural bias bound to exist in real persons’ names.

In the remainder of this paper, we first present the relevant background, including both theoretical work and computational work relevant to peoples’ names (Section 2). Based on this theoretical work, we then proceed to formulate a set of features that can be computationally extracted from names, and which we hypothesise to be discriminative enough to allow for the construction of a model that accurately predicts whether a character plays a positive or negative role in a work of fiction (Section 3). In order to test this hypothesis, we constructed a corpus of characters from popular English-language motion pictures. After describing corpus construction and presenting results (Section 4), we proceed to discuss these results (Section 5) and conclude (Section 6).

2 Background

2.1 Onomastics

The procedure of naming an individual, a location or an object is of particular importance and serves

purposes beyond the obvious purpose of referring to distinct entities. Characteristics such as place of origin, gender, and socioeconomic status can often be guessed from the name or nickname that has been attributed to an individual. *Onomastics*, the study of the origin, history, and use of proper names has attracted scholarly attention as early as antiquity and Plato's 'Cratylus' (Hajdú, 1980).

In fiction and art, in particular, names are chosen or invented without having to follow the naming conventions that are common in many cultures. This allows creators to apply other criteria in selecting a name for their characters, one of which being the intuitions and preconceptions about the character that the name alone implies to the audience. Black and Wilcox (2011) note that writers take informed and careful decisions when attributing names to their characters. Specifically, while care is taken to have names that are easily identifiable and phonologically attractive, or that are important for personal reasons, these are not the only considerations: names are chosen so that they match the personality, the past, and the cultural background of a character.

According to Algeo (2010) behind each name lies a story while Ashley (2003) suggests that a literary name must be treated as a small poem with all the wealth of information that implies. Markey (1982) and Nicolaisen (2008) raised concerns on whether onomastics can be applied to names in art given the different functional roles of names as well as their intrinsic characteristics, namely sensitivity and creativity. 'Redende namen' (significant names) is a widespread theory that seeks the relationship between name and form (Rudnyckyj, 1959). According to this theory, there is a close relationship between the form of a name and its role. This consideration is still prevalent to date as shown by Chen (2008) in her analysis of names in comic books, where names transparently convey the intentions of the creator for the role of each character. Another concern is whether the study of literary names should be examined individually for each creative work or if generalizations can be made (Butler, 2013). However, the scope of most studies is limited to individual projects or creators, creating an opportunity for computational methods that can identify generalizations and patterns across larger bodies of literary work than what is manually feasible.

2.2 Related Work

Although serving radically different purposes and applications than our investigation, various methods for the computational analysis of proper nouns have been developed in natural language processing. Without a doubt, some of the oldest and most mature technologies that exploit the properties of proper nouns are those addressing *named entity recognition and categorization (NERC)*. In this direction, there is a recently ongoing effort for the extension of NERC tools so that they cover the needs of literary texts (Borin et al., 2007; Volk et al., 2009; Kokkinakis and Malm, 2011).

Moving beyond recognition, effort has been made to explore characteristics and relationships of literary characters (Nastase et al., 2007). Typically, however, these efforts take advantage of the context, and very little work tries to extract characteristics of literary characters from their names alone. One example is the application of language identification methods in order to extract the cultural background of proper names (Konstantopoulos, 2007; Bhargava and Kondrak, 2010; Florou and Konstantopoulos, 2011). This work showed that people's names in isolation are more amenable to language identification than common nouns. Konstantopoulos (2007), in particular, reports inconclusive results at pinpointing the discriminative features that are present in people's names but not in other words.

Another relatively recent and related research direction that does not focus on proper nouns investigates elements of euphony mostly by examining phonetic devices. The focus is to identify how the sound of words can foster its effectiveness in terms of persuasion (Guerini et al., 2015) or memorability (Danescu-Niculescu-Mizil et al., 2012).

3 Approach

These earlier attempts relied on the examination of predictive models of n-grams in order to identify the n-grams that are the best discriminants. The aim was that by inspecting these most discriminative n-grams, meaningful patterns would emerge and serve as the vehicle for formulating hypotheses about the correlation between what names sound like and the cultural background of the persons bearing them.

This approaches largely ignored the background in onomastics and literary research. By contrast, we exploit this prior body of theoretical work

ID	Feature	Category	Type
1	words count	phonological	numeric
2	vowels count	phonological	numeric
3	consonants count	phonological	numeric
4	plosives count	phonological	numeric
5	fricatives count	phonological	numeric
6	affricates count	phonological	numeric
7	nasals count	phonological	numeric
8	vowel start	phonological	categorical
9	vowel end	phonological	categorical
10	voice start	phonological	categorical
11	subsequent letters count	phonological	categorical
12	low vowel	phonological	categorical
13	high vowel	phonological	categorical
14	definite article	lexical form	categorical
15	consonance	poetic	numeric
16	assonance	poetic	numeric
17	alliteration	poetic	numeric
18	name and title resemblance	domain	numeric
19	credit index	domain	numeric
20	genre	domain	categorical
21	sentiment soundex wordnet	emotions	numeric
22	sentiment leven-shtein wordnet	emotions	numeric
23	gender	social	categorical
24	foreign suffix	social	categorical
25	first name frequency	social	numeric
26	last name frequency	social	numeric
27	full name frequency	social	numeric
28	honor	social	categorical

Table 1: List of features

to define more sophisticated features that directly correspond to theoretical hypotheses. Our empirical experiments are now aimed at identifying the features (and thus hypotheses) that are the most discriminative, rather than at hoping that a coherent hypothesis can be formulated by observing patterns in n-gram features.

In the remainder of this section, we will present these hypotheses and the machine-extracted features that reflect them. The features are also collected in Table 1.

3.1 Emotions

Hypothesis 1 *The (positive or negative) polarity of the sentiment that a character’s name evokes is associated with the polarity of the character’s role.*

The understanding of how the language transmits emotions has attracted significant research attention in the field of Computational Linguistics. Most of the relevant literature is directed towards calculating sentiment for units at the document or

sentence level. These works are usually boosted by semantic dictionaries that provide information about the emotional hue of concepts such as the Linguistic Inquiry and Word Count (LIWC) (Pennebaker et al., 2001), the Harvard General Inquirer (Stone et al., 1966), the WordNet Affect (Strappavara and Valitutti, 2004) and SentiWordNet (Esuli and Sebastiani, 2006). In our task, the absence of context and the inherent arbitrariness in naming (even in fictional names) increases the difficulty in conveying emotional quality to names. More specifically, the intriguing part was to associate fictional names with concepts from a semantic sentiment resource in order to approximate a sentiment value. To achieve this we used SentiWordNet: a linguistic resource that has derived from the annotation of WordNet synsets according to the estimated degree of positive, negative or neutral hue. The overall valence for a given name is calculated as the sum of the valence of its elements (first name, surname). The valence of each name element is the average valence of all SentiWordNet concepts that are associated with it. To associate a name element and a SentiWordNet concept we used the Soundex phonetic distance and the Levenshtein lexicographic distance (Levenshtein, 1966). A heuristic threshold is used to decide whether a name and a SentiWordNet concept are associated.

More formally, the valence $\text{val}(n)$ of a name n comprising name elements e_i is calculated as follows:

$$\text{val}(e_i) = \frac{\sum_{u \in \text{ass}_S(e_i)} \text{swn}(u) + \sum_{v \in \text{ass}_L(e_i)} \text{swn}(v)}{|\text{ass}_S(e_i)| + |\text{ass}_L(e_i)|}$$

$$\text{val}(n) = \sum_i \text{val}(e_i)$$

where $\text{ass}_S(\cdot)$ is the set of SentiWordNet concepts that are Soundex-associated with the given name element, $\text{ass}_L(\cdot)$ the set of SentiWordNet concepts that are Levenshtein-associated with the given name element, and $\text{swn}(\cdot)$ the valence assigned to the given concept by SentiWordNet.

3.2 Stylistic and poetic features

Hypothesis 2 *Assuming Ashley’s (2003) and Butler’s (2013) position that ‘a name can be a whole “poem” in as little as a single word’ we assume that stylistic features usually found in poems can be extracted from the names of fictional charac-*

ters, and that such features correlate with the polarity of their roles.

The first quantitative analysis efforts of the poetic style can be found in the 1940s and in the study of the poet and literary critic Josephine Miles (1946; 1967) where she studied the features of poems over time. Despite the great contribution of this work and others that followed, the creation of a framework for quantitative poetic style analysis remained limited to a small number of poems and much of the work was done manually. The work of Kaplan and Blei (2007) is an attempt to automate and analyze large volumes of poems exploring phonological, spelling and syntax features. For our work, we identified the following poetic devices that can be applied to isolated names:

- **Alliteration:** a stylistic literary device identified by the repeated sound of the first consonant in a series of multiple words, or the repetition of the same sounds of the same kinds of sounds at the beginning of words or in stressed syllables of a phrase.
Examples: Peter Parker, Peter Pan
- **Consonance:** a poetic device characterized by the repetition of the same consonant two or more times in short succession.
Examples: Lillian Hellman, Freddy Krueger, Hannibal Lecter, Kristen Parker
- **Assonance:** a poetic device characterized by the repetition of the same vowel two or more times in short succession.
Examples: Bobbie Ritchie

3.3 Phonological features

Hypothesis 3 *The presence of specific phonological features can reveal evidence of the role of a character in an artistic artifact.*

Linguistic theory widely adopts the concept of arbitrary relationship between the *signifier* and the *signified* (de Saussure, 1916 1983; Jakobson, 1965). However, an increasing volume of works in various fields investigates the existence of non-arbitrary relations between phonological representation and semantics, a phenomenon known as *phonological iconicity*. Standing from the side of Computational Linguistics and with the intuition that in fictional names the correlation between a word's form and the emotion it expresses will be stronger, we examined a wide range of phonology-related features, shown in Table 1. It should be

noted that these features are extracted from the phonetic representation of names derived by applying the spelling-to-phoneme module of the *espeak* speech synthesizer.¹

3.4 Sociolinguistic features

Hypothesis 4 *We hypothesize that social aspects of names — such as frequency of use or use of foreign names in a given environment — can relate to role of a fictional character. For instance, a ‘girl next door’ role is more likely to be assigned a very popular female name than a name that sounds hostile or foreign.*

The frequency of names in U.S.A was calculated based on the Social Security Death Index (SSDI), a publicly available database that records deaths of U.S.A citizens since 1936.² The same dataset was also used to build a model for recognizing foreign-looking names. More specifically, we trained n-gram language models of order 2–5 against the dataset for both orthographic and phonetic representation using the *berkeleylm* library (Pauls and Klein, 2011). We then heuristically defined a threshold that correlates well with foreign-looking suffixes. Analogously with the name frequency we extract the gender of each name using a baby names dataset that includes gender information.³ For unisex names the prevalent gender was picked. Finally, honorific titles (e.g. Professor, Phd, Mr, Mrs etc.) were also extracted from names. Honorific titles are intriguing due to their ambiguous meaning since they can express respect and irony in different contexts.

3.5 Domain features

Hypothesis 5 *We pursued indications to check if domain-related features such as the appearance time of a character in a movie, the movie title or the movie genre is associated (correlates) with the problem under study.*

In this category lies the *feature_{sameastitle}* since anyone with a quick glance in a list of films would notice that a fictional name often consists of, or is the part of, the movie title, as in, There's Something about Mary, Hannibal, Thelma & Louise, Rocky, etc. On IMDB character names are presented in the form of a list in

¹Please cf. <http://espeak.sourceforge.net>

²Please cf. <https://archive.org/details/DeathMasterFile>

³Specifically, we used <https://www.ssa.gov/oact/babynames/state/index.html>

descending order based on screen credits. In the $feature_{creditindex}$ we want to check if the naming process is more assiduous for the roles of protagonists based on this list. In the same direction, we examine the $feature_{genre}$ for a possible correlation between the role of a character and the genre of a film.

4 Experiments and Results

4.1 Data Collection and Annotation

In order to validate our approach, we first need a corpus of names of fictional characters, annotated with the polarity of their role. As such a resource does not exist to the best of our knowledge, we have created it for the purposes of the work described here.

Our decision to use motion pictures rather than other fictional work is motivated by the relative ease of finding annotators familiar with the plot of these works, so that we could get reliable annotations of the polarity of the leading roles. We compiled a list of 409 movies based on the following criteria:

- That they are widely known films, covering all genres of film production. We automatically crosschecked if the candidate movies are included in DBPedia⁴ and YAGO⁵, as these are indicators that the films are known to the general public.
- That they have received some award or are positively evaluated by users (i.e., have an IMDB rating of 5.0 or higher). The underlying assumption is that this criterion selects major productions where care has been given to even the most minute detail, including the names of the major characters and what these names connote to the audience.
- That they are recent productions, so that annotators can easily recall the plot and the characters.

We then asked volunteers to select any movie from the list that they were very familiar with, and to assign one of *positive*, *negative* or *neutral* to the top-most characters in the credits list, working only as far down the credits list as they felt confident to. The three categories were defined as follows in the annotation guidelines:

⁴<http://wiki.dbpedia.org>

⁵<http://www.mpi-inf.mpg.de/yago>



Figure 1: Character annotation tool

- **Positive:** when the role of the character in the plot left a positive impression on you when you saw the movie.
- **Negative:** when the role of a character left a negative impression on you when you saw the movie.
- **Neutral:** when the role of the character is important for the plot, but you are in doubt or cannot recall whether it was a positive or a negative role.

Neutral tags are ignored in our experiments. They were foreseen only to allow annotators to skip characters and still have a sense of accomplishment, so that they only make choices that they are confident with.

We used the Hypothes.is⁶ open source annotation application. The annotation was carried out by having volunteers install the Hypothes.is Web browser extension and then visit the IMDB⁷ page of any of the movies on our list (direct links were provided to them in the guidelines). IMDB was chosen due to its popularity, so that annotators would already be familiar with the online environment. The annotators tagged the character names directly on the IMDB page and the annotations were collected for us by Hypothes.is (Figure 1).

Eight annotators participated in the procedure and provided 1102 positive and 434 negative tags for characters of 202 movies, out of the 409 movies in the original list. Table 2 gives the annotation distribution per movie genre.

The reliability of the annotated collection by means of inter-rater agreement was also measured. For this purpose, various standard agreement measures (Meyer et al., 2014) were calculated, all showing very high agreement among the annotators (Table 3). This demonstrates that the annota-

⁶<https://hypothes.is>

⁷<http://www.imdb.com>

	Original		Resampled	
	Pos	Neg	Pos	neg
Action	262	107	244	102
Adventure	126	63	133	62
Animation	73	22	63	27
Biography	28	6	39	8
Comedy	78	25	23	21
Crime	68	25	81	18
Drama	81	40	76	32
Horror	16	12	28	13
Musical	0	0	0	0
Mystery	20	13	26	17
Sci-Fi	2	0	2	0
Thriller	0	0	0	0
Western	1	2	3	2
Sum	755/315	768/302		

Table 2: Number of annotations per genre before and after resampling

Measure	Value
Percentage Agreement	0.963
Hubert Kappa Agreement	0.980
Fleiss Kappa Agreement	0.973
Krippendorff Alpha Agreement	0.979

Table 3: Inter-annotator agreement

tion task is well-formulated, but does not guarantee that our classification task is consistent, since the latter will use different information than that used by the annotators. That is to say, the annotators had access to their understanding of the movies’ plot to carry out the task, whereas our classification task will be performed over the characters’ names alone.

The collection is publicly available, including the guidelines and instructions to the annotators, the source code for the annotation tool, and the source code for the tool that compiles Weka ARFF files from the JSON output of the annotation tool.⁸

4.2 Experimental Design

The experimental design consisted of an iterated approach performing experiments with different sets of features. This process was driven by a preliminary chi-squared analysis in order to exploit feature significance. The algorithms that are used for the experiments are Naive Bayes and J48

⁸<https://bitbucket.org/dataengineering/fictionalnames>

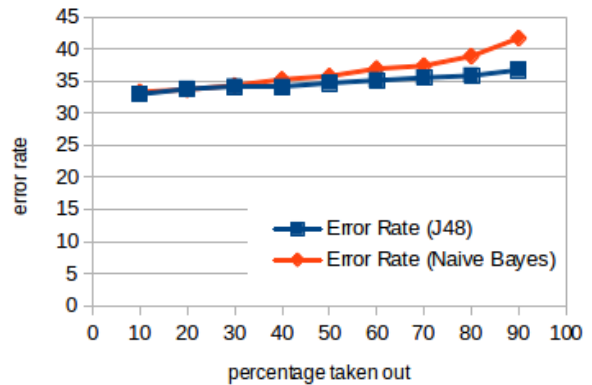


Figure 2: Learning curve for the number of instances

	All		Without domain	
	NB	J48	NB	J48
Recall	0.723	0.824	0.718	0.803
Prec.	0.731	0.822	0.515	0.801
F-score	0.618	0.823	0.6	0.802

Table 4: Comparison of Naive Bayes and J48

(Salzberg, 1994) decision trees. Each experiment is done using a 10-fold cross validation on the available data, using a confidence factor of 0.25 for post-pruning. For all the experiments we used the Weka toolkit (Hall et al., 2009). Due to the imbalance of our dataset in favor of positive classes (see Table 2), we sub-sampled the dataset maintaining the initial genre distribution. We also applied *principal component analysis (PCA)* in order to guarantee the independence of the classification features, as required by the Naive Bayes algorithm. To explore the behavior of the algorithms to the change of trained data we generated the learning curves shown in Figure 2. In both cases the learning curves are well-behaved since the error rate grows monotonically as the training set shrinks. However, the precision, recall, and F-scores achieved by J48 are significantly better than those of Naive Bayes (Table 4).

This preliminary experiment led us to use J48 for the main experiment, where we try different features in order to understand which are the most discriminative ones. These results are collected in Table 5 and discussed immediately below.

5 Discussion of Results

A first observation that can be easily made is that the domain features are good discriminants. As these features exploit information such as credit-

	Rec.	Prec.	F-score
Without domain features	0.803	0.801	0.802
Only domain features	0.725	0.699	0.667
Only phonological features	0.790	0.786	0.787
Without poetic features	0.836	0.832	0.833
Without consonance feature	0.823	0.820	0.821
Without emotions features	0.814	0.810	0.811
Without phonological features	0.798	0.792	0.793
Without social features	0.807	0.803	0.804
All features	0.824	0.822	0.823

Table 5: Performance of J48 for different feature settings

Most frequent in positive characters	
Phoneme	Examples
n-gram	
/lɪ/	Ned Alleyn (Shakespeare in Love)
/an/	Anouk Rocher (Chocolat)
/aɪ/	Eliza Doolittle (My Fair Lady)
/nɪ/	Linguini (Ratatouille)
/ɪst/	Kevin McCallister (Home Alone)
/ɪəʊ/	Frodo (The Lord of the Rings)
/and/	Dylan Sanders (Charlie’s Angels)
/stə/	C.C. Baxter (The Apartment)
Most frequent in negative characters	
Phoneme	Examples
n-gram	
/əɪ/	Tom Buchanan (The Great Gatsby)
/əʊ/	Iago (Aladdin)
/tə/	Norrington (Pirates of the Caribbean)
/ɪ/	Tom Ripley (The Talented Mr. Ripley)
/mən/	Norman Bates (Psycho)
/mɪs/	Mystique (X-Men)
/ktə/	Hannibal Lecter (Hannibal)

Table 6: Frequent phoneme {2,3}-grams

ing order that is outside the scope of our hypotheses, there were expected to be good discriminants and are included for comparison only.

By comparing the performance of all features (F = 82%), domain-only features (F = 68%), and all-except-domain features (F = 80%), we can immediately understand that our name-intrinsic features are better discriminants than domain features; in fact, name-intrinsic features not just better than domain features, they are by themselves almost as good as domain and name-intrinsic features combined. This is a significant finding, as it validates our core hypothesis that there is a correla-

tion between what fictional character names look and sound like and the role they play in the plot of the fictional work they appear in.

We will now proceed to look in more detail into the different categories of features used, in order to gain further insights about specific discriminants.

5.1 Phonological Features

The phonological features are important separation criteria as evidenced by the drop in performance when they are excluded from the experimental setup (Table 5). Specifically, using all features except phonological features is equivalent to using phonological features alone (about F = 79% in both cases) and slightly worse than using all name-intrinsic features (about F = 80%). By comparison, removing any other category *increases* performance, leading us to believe that all other features are actually adding noise (rather than discriminatory power) to the feature space.

In order to delve more into this category of features, we proceeded with an n-gram analysis (of order 1 through 4) to look for correlations between phonemes. The results clearly demonstrated the positive effect of the number of vowels (normalized by the length of the utterance) to the positive category. As far as the consonants are concerned, voiced (e.g. /ʌ/, /g/, /d/, /w/) seem to relate more to the negative class. Table 7 summarizes a more fine-grained analysis for the consonants based on their categorization.

The environment plays an important role, with specific combinations showing tendencies that are not observed with isolated phonemes. For example, diphoneme /an/ relates to positive class while /əɪ/ to negative. Table 6 lists some frequent phoneme 2- and 3-gram examples. The position of each phoneme also seems to play an crucial role

Phonemes	Class
/p/, /b/ (bilabial plosive)	P
/l/ (alveolar lateral)	P
/f/, /v/ (labiodental affricative)	N
/k/, /g/ (velar plosive)	N
/t/, /d/ (alveolar plosive)	N
/dʒ/, /tʃ/ (affricate)	N
/m/, /n/ (nasal)	N
/ɹ/ (alveolar retroflex)	N

Table 7: Consonants behavior

in the classification task. Specifically, we note that starting with a vowel or a consonant are among the most discriminating features. These observations are consistent to a great extent with work in psychology and literary theory that studied phonological iconicity for common words (Nastase et al., 2007; Auracher et al., 2011; Schmidtke et al., 2014).

Some contradictory conclusions in these works are attributed by researchers to the methodologies applied, while at the same time concerns are raised whether such methodologies can inductively lead to cross-language and general conclusions (Auracher et al., 2011). Table 8 summarizes some of the outcomes of these works.

5.2 Emotion and Affect

The analysis showed that the features that calculate the emotional load of fictional names based on SentiWordNet contribute to the classification task. However, we believe that there is still room for improvement for the performance of this feature mainly towards the optimization of the selection threshold in order to reduce the degree of false positive matches as well as the addition of more lexical resources for example WordNet Affect or LIWC.

5.3 Social Features

The annual publication *It's a Man's (Celluloid) World* examines the representation of female characters every year. According to its 2015 results (Lauzen, 2015), gender stereotypes were abundant with female characters being younger than their male counterparts and more likely to have pro-social goals including supporting and helping others. This bias makes the gender feature discriminative, but in a way that is not linguistically interesting: female characters are simply related to the

Reference	Description
Taylor and Taylor (1965)	evidence that pleasantness relations are language specific
Fonagy (1961)	sonorants (e.g., /l/,/m/) more common in tender poems, plosives (e.g., /k/,/t/) in aggressive ones
Miall (2001)	Passages about Hell from Miltons "Paradise Lost" were found to contain significantly more front vowels and hard consonants than passages about Eden while the latter contained more medium back vowels
Whissell (1999)	plosives correlate with unpleasant words
Auracher et al. (2011)	nasals (e.g., /m/) relate to sadness, plosives (e.g., /p/) to happiness, parallels across remote languages
Zajonc et al. (1989)	umlaut /y/ causes negative affective states

Table 8: Phonological iconicity studies

positive class.

A somewhat surprising result was that the foreign suffix feature is not discriminative. The hypothesis that the concept of the 'other' is stereotyped negatively does not seem to be true in our dataset. A closer investigation might identify genres where this hypothesis holds (e.g., war movies), but this would be implicit pragmatic information about the context of the film rather than a linguistically interesting finding.

5.4 Poetic and Stylistic Features

The experimental findings show that literary devices can actually be identified in fictional characters names, but the same findings also indicate that they do not contribute significantly to the classification task. More specifically, consonance is the only stylistic/poetic feature that affects classification.

6 Conclusions and Future Work

In this paper we test the hypothesis that the sound and the form of fictional characters' names correlates with meaning, in our particular case with

the respective characters' role in the work of fiction. We restricted our study to fictional characters since they are not tied to cultural conventions of naming, such as names that run in a family, so that we are able to look for patterns that are perceived as positive or negative by the audience and used as such (consciously or not) by the creator.

Our experiments have verified that features intrinsic to the names and without any reference to the plot or, in general, any other context are discriminative. Furthermore, we have discovered that the most discriminative features are of phonological nature, rather than features that hint at pragmatic information such as the gender or origin of the character. A further contribution of our work is that we ran an annotation campaign and created an annotated corpus of fictional movie characters and their corresponding polarity. This corpus is offered publicly, and can serve experimentation in the digital humanities beyond the scope of the experiments presented here.

Our future research will test the correlation between the polarity and the name of a fictional character beyond the movie domain. It would, for example, be interesting to seek differences between spoken names (as in films) and names that are only meant to be read (as in literature). In addition, using written literature will allow us to compare texts from different periods, pushing earlier than the relatively young age of motion pictures. Character polarity annotations in written literature could be created by, for example, applying sentiment analysis to the full text of the work.

References

- [Algeo2010] John Algeo. 2010. Is a theory of names possible? *Names*, 58(2):90–96.
- [Ashley2003] Leonard R. N. Ashley. 2003. *Names in Literature*. Bloomington, IN: Authorhouse (formerly 1st Books).
- [Auracher et al.2011] Jan Auracher, Sabine Albers, Yuhui Zhai, Gulnara Gareeva, and Tetyana Stavniychuk. 2011. P is for happiness, N is for sadness: Universals in sound iconicity to detect emotions in poetry. *Discourse Processes*, 48(1):1–25.
- [Bhargava and Kondrak2010] Aditya Bhargava and Grzegorz Kondrak. 2010. Language identification of names with SVMs. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 693–696, Los Angeles, California, June. Association for Computational Linguistics.
- [Black and Wilcox2011] Sharon Black and Brad Wilcox. 2011. 188 unexplainable names: Book of Mormon names no fiction writer would choose. *Religious Educator*, 12(2).
- [Borin et al.2007] Lars Borin, Dimitrios Kokkinakis, and Leif-Jöran Olsson. 2007. Naming the past: Named entity and animacy recognition in 19th century Swedish literature. In *ACL 2007 Workshop on Language Technology for Cultural Heritage Data (LaTeCH 2007)*, pages 1–8.
- [Butler2013] James Odelle Butler. 2013. *Name, Place, and Emotional Space: Themed Semantics in Literary Onomastic Research*. Ph.D. thesis, University of Glasgow.
- [Chen2008] Lindsey N. Chen. 2008. Ethnic marked names as a reflection of United States isolationist attitudes in Uncle Scrooge comic books. *Names*, 56(1):19–22.
- [Danescu-Niculescu-Mizil et al.2012] Cristian Danescu-Niculescu-Mizil, Justin Cheng, Jon M. Kleinberg, and Lillian Lee. 2012. You had me at hello: How phrasing affects memorability. *CoRR*, abs/1203.6360.
- [de Saussure1916 1983] Ferdinand de Saussure. [1916] 1983. *Course in General Linguistics*. Duckworth, London. (translation Roy Harris).
- [Esuli and Sebastiani2006] Andrea Esuli and Fabrizio Sebastiani. 2006. SENTIWORDNET: A publicly available lexical resource for opinion mining. In *Proceedings of the 5th Conference on Language Resources and Evaluation (LREC'06)*, pages 417–422.
- [Florou and Konstantopoulos2011] Eirini Florou and Stasinos Konstantopoulos. 2011. A quantitative and qualitative analysis of Nordic surnames. In *Proceedings of the 18th Nordic Conference of Computational Linguistics (NODALIDA 2011), May 11-13, 2011, Riga, Latvia*, volume 11 of *NEALT Proceedings Series*.
- [Fonagy1961] Ivan Fonagy. 1961. *Communication in Poetry*. William Clowes.
- [Guerini et al.2015] Marco Guerini, Gözde Özbal, and Carlo Strapparava. 2015. Echoes of Persuasion: The Effect of Euphony in Persuasive Communication. *CoRR*, abs/1508.05817.
- [Hajdú1980] Mihály Hajdú. 1980. The history of Onomastics. *Onomastica Uralica*, 2:7–45.
- [Hall et al.2009] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA data mining software: An update. *SIGKDD Explor. Newsl.*, 11(1):10–18, November.

- [Jakobson1965] Roman Jakobson. 1965. Quest for the Essence of Language. *Diogenes*, 13(51):21–37.
- [Kaplan and Blei2007] David M. Kaplan and David M. Blei. 2007. A computational approach to style in American poetry. In *Proceedings of the 7th IEEE International Conference on Data Mining (ICDM 2007)*, pages 553–558, October.
- [Kokkinakis and Malm2011] Dimitrios Kokkinakis and Mats Malm. 2011. Character profiling in 19th century fiction. In *Workshop: Language Technologies for Digital Humanities and Cultural Heritage in conjunction with the Recent Advances in Natural Language Processing (RANLP). Hissar, Bulgaria*.
- [Konstantopoulos2007] Stasinou Konstantopoulos. 2007. What’s in a name? In Petya Osenova, Erhard Hinrichs, and John Nerbonne, editors, *Proceedings of Computational Phonology Workshop, International Conf. on Recent Advances in NLP, (RANLP), Borovets, Bulgaria, September 2007*.
- [Lauzen2015] Martha Lauzen. 2015. It’s a man’s (celluloid) world: On-screen representations of female characters in the top 100 films of 2011. Technical report, San Diego State University Center for the Study of Women in Television and Film, School of Theatre, Television and Film, San Diego State University, San Diego, CA.
- [Levenshtein1966] Vladimir I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady*, 10:707, feb.
- [Markey1982] T. L. Markey. 1982. Crisis and cognition in onomastics. *Names*, 30(3):129–142.
- [Meyer et al.2014] Christian M. Meyer, Margot Mieskes, Christian Stab, and Iryna Gurevych. 2014. DKPro Agreement: An Open-Source Java Library for Measuring Inter-Rater Agreement. In *Proceedings of the 25th International Conference on Computational Linguistics: System Demonstrations (COLING)*, pages 105–109, Dublin, Ireland, August.
- [Miall2001] David Miall. 2001. Sounds of contrast: An empirical approach to phonemic iconicity. *Poetics*, 29(1):55–70.
- [Miles1946] Josephine Miles. 1946. Major adjectives in English poetry: From Wyatt to Auden. *University of California Publications in English*, 12.
- [Miles1967] Josephine Miles. 1967. *Style and Proportion: The Language of Prose and Poetry*. Little, Brown and Co., Boston.
- [Nastase et al.2007] Vivi Nastase, Marina Sokolova, and Jelber Sayyad Shirabad. 2007. Do happy words sound happy? A study of the relation between form and meaning for English words expressing emotions. In *Recent Advances in Natural Language Processing (RANLP 2007)*.
- [Nicolaisen2008] William F. H. Nicolaisen. 2008. On names in literature. *Nomina*, 31:89–98.
- [Pauls and Klein2011] Adam Pauls and Dan Klein. 2011. Faster and smaller n-gram language models. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*.
- [Pennebaker et al.2001] JW Pennebaker, ME Francis, and RJ Booth. 2001. *Linguistic inquiry and word count [computer software]*. Mahwah, NJ: Erlbaum Publishers.
- [Rudnycky1959] Jaroslav B. Rudnycky. 1959. Function of proper names in literary works. *Internationalen Vereinigung für moderne Sprachen und Literaturen*, 61:378–383.
- [Salzberg1994] Steven L. Salzberg. 1994. C4.5: Programs for machine learning. *Machine Learning*, 16(3):235–240.
- [Schmidtke et al.2014] David S. Schmidtke, Markus Conrad, and Jacobs Arthur M. 2014. Phonological iconicity. *Frontiers in Psychology*, 12.
- [Stone et al.1966] Philip J. Stone, Dexter C. Dunphy, Marshall S. Smith, and Daniel M. Ogilvie. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, Cambridge, MA.
- [Strapparava and Valitutti2004] Carlo Strapparava and Alessandro Valitutti. 2004. WordNet-Affect: An affective extension of WordNet. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, pages 1083–1086. ELRA.
- [Taylor and Taylor1965] I. K. Taylor and M. M. Taylor. 1965. Another look at phonetic symbolism. *Psychological Bulletin*, 65.
- [Volk et al.2009] Martin Volk, Noah Bubenhofer, Adrian Althaus, and Maya Bangerter. 2009. Classifying named entities in an Alpine heritage corpus. *Künstliche Intelligenz*, pages 40–43.
- [Whissell1999] Cynthia Whissell. 1999. Phonosymbolism and the emotional nature of sounds: Evidence of the preferential use of particular phonemes in texts of differing emotional tone. *Perceptual and Motor Skills*, 89(1):19–48, August.
- [Zajonc et al.1989] R. B. Zajonc, Sheila T. Murphy, and Marita Inglehart. 1989. Feeling and facial efference: Implications of the vascular theory of emotion. *Psychological Review*, 96(3):395–416, July.