Chapter 20

Looking for meaning in names

Stasinos Konstantopoulos
National Centre of Scientific Research "Demokritos"

This article discusses the concept that proper names are not semantically empty denotations, but characterize in many, often subliminal, ways their denotee. The discussion is driven by computational experiments on using just a name to guess the linguistic and cultural background of a person and the positive or negative polarity of a fictional character's role.

1 Introduction

This article discusses the concept that proper names are not semantically empty denotations, but they characterize, sometimes subliminally, their denotee. To give a first, straightforward example, the reader will intuitively agree that one can guess (or at least considerably narrow down) a person's linguistic and cultural background from that person's name alone. But there might also be less intuitive correlations between names and their bearers' properties.

To qualify such a broad and general claim into a more concrete object of investigation, we will narrow its scope to correlations between how a name 'sounds', its phonetic properties, and properties of the name's bearer. In our language guessing example, we expect that guesses are based to large extent on knowledge of a language's orthographic conventions, characteristic morphological markers, and characteristic lexical units. The question we discuss here is if *phonetic* characteristics are also relevant to our intuition about where somebody comes from, and also what other properties can be intuited besides origin.

However, in the case of actual people's names it is hard to imagine that they would be associated with anything but pleasant and positive meanings, if anything at all. But if we turn our attention to *fictional* characters' names, we expect their creators to have (consciously or not) named them to sound like they behave, so that there will be both positively and negatively sounding names. Ponder, for example, whether it would sound natural if *Hannibal Lecter* was a positive character and *Frodo* was a negative one. If it wouldn't, and we had never heard of these names before, that would

imply a prior on what a positive or negative character should sound like, essentially a semantics that is linked to the sound of the name.

2 Names and computational methods

In fiction and art names serve a more complex functional role than denotation alone and should be treated differently than other proper names (Markey 1982; Nicolaisen 2008). Since they are chosen or invented to satisfy a different role, creators apply different criteria in selecting a name for their characters than the conventions or aesthetics used to select people's names. One of these criteria is the intuitions and preconceptions about the character that the name alone implies to the audience (Rudnyckyj 1959; Algeo 2010). In fact, Ashley (2003) suggests that a literary name must be treated as a 'small poem' with all the wealth of information that such a statement implies.

This more general observations can sometimes be framed into specific intuitions and preconceptions in the context of each given work. Chen (2008), for instance, argues that the ethnic-marked names created by Carl Barks for the *Uncle Scrooge* comic books in the 1950s and 1960s contributed to the books' success by feeding into the isolationist feelings of post-war US. However, there is legitimate concern regarding the validity of generalizations made by studying individual creative works (Butler 2013).

This creates an opportunity for computational methods that can extract patterns from larger bodies of literary work than what is manually feasible. Naturally, more relevant to us are methods that concern the meaning of isolated words rather than the grammatical structure that combines them into text. Coming back to guessing linguistic and cultural background from a name, the relevant literature originates in *speech synthesis* where *language identification* is used to select different pronunciation models for foreign names. Starting from hand-crafted rules (Spiegel 1985), the speech synthesis community has since moved to machine-learned language identification models (Font Llitjós & Black 2001). The same general methodology has been applied to automatic transliteration of named entities for the purposes of *multi-lingual information extraction* (Virga & Khudanpur 2003) and *machine translation* (Huang 2005).

In these applications, the language identification pre-processing was reported to improve accuracy on the overall task, but the accuracy of the language identification step itself was not reported. The first results on how indicative a name is of its linguistic background comes from experiments on applying n-gram modelling to a corpus of European names and nationalities, compiled by harvesting from the Web information about football players and their national squad eligibility (Konstantopoulos 2007). This work and subsequent analyses of the same data (Konstantopoulos 2010; Florou & Konstantopoulos 2011) aimed at comparing language identification of people's surnames versus common words and identifying the features that make the former more characteristic of their linguistic background than the latter. But they have also yielded lateral results more directly related this discussion, and specifically the

analysis of the roots and of the derivational morphology used in surname formation.

3 Word form and meaning

In another strain of related word, form-meaning systematicity has been investigated in the context of exploring the idea that the mental lexicon would be maximally easy to organize if there were a transparent, structure-preserving relationship where words sound similar to the extent that they mean similar things (Shillcock et al. 2001). This idea has been tested in English (Shillcock et al. 2001; Monaghan, Shillcock & Christiansen 2014; Gutiérrez, Levy & Bergen 2016) and Spanish (Tamariz 2008), typically restricted to mono-morphemic words to avoid confusing simultaneous form/meaning similarity due to derivation for form-meaning systematicity.

This body of work provides significant evidence of form-meaning systematicity. It should be noted, however, that this work assumes a *distributional semantics* to interpret 'meaning' and is, accordingly, evaluated on the task of using form features to predict what other words the given word co-occurs with. As Gutiérrez, Levy & Bergen (2016: Section 6) also note, these experiments tell us very little about human intuition regarding the meaning of an unknown word heard outside of any context.

Naturally, such questions offer themselves to psycholinguistics research such as the experiments conducted by Ramachandran & Edward (2001) to identify cross-lingual and cross-cultural correlations between nonce words and shapes. But, and remaining on the more familiar ground of computational linguistics, they also offer themselves to computational experiments on predicting the meaning of words from their form, including non-distributional interpretations of 'meaning'.

4 Names of fictional characters

This brings us back to names in fiction: the names of fictional characters reflect (possibly subconsciously) a perception shared between the creator and the audience of what a character's name ought to sound like. The personal preferences or experiences of the creator might add noise, but given a large enough corpus fictional characters' names can uncover analogies and familiarities within a given linguistic and cultural background.

Motivated by this idea, Papantoniou & Konstantopoulos (2016) created a corpus of names of fictional characters in motion pictures, annotated with the *polarity* of their role in the plot. Eight annotators identified 1102 positive and 434 negative characters in 202 movies. The annotation guidelines stressed that only clear-cut cases should be annotated and the overall setup made it easy for the annotators to avoid making a commitment and move on to the next character or movie, yielding a high degree of agreement (Table 1).

Obviously, the fact that the annotators made their decision by looking at the cast credits cannot be construed as anything deeper than their knowledge of the movie's plot. That is to say, regardless of whether the character's name was as common

as *Jane Smith* or as unique as *Darth Vader*, it was already familiar and the results of the annotation task are *not* meant to be used to extract any conclusions about predicting polarity from nonce word forms. They can, however, be used to create computational models that predict polarity where we can experiment by controlling what background knowledge we make available to the machine learning algorithm and observing the resulting prediction accuracy. By identifying the background features that are the best predictors, we can become informed about the inventory of phonological characteristics, semantic and pragmatic analogies, and other devices that creators use to share their perception of the character with their audience.

As noted earlier, experimenting with very difficult language identification tasks (Florou & Konstantopoulos 2011: Nordic surnames) has shown that even closely related, and otherwise difficult to separate backgrounds (e.g., Norwegian and Danish) exhibit different patterns regarding what derivational morphology is applied to roots from different semantic classes (such as in occupational, locative, or patronymic surnames). In the Nordic surname experiments, n-gram modelling was unable to identify such fine patterns which were only discovered by human observation and verified by encoding them as a DCG that was evaluated on the corpus.

This led Papantoniou & Konstantopoulos (2016) to define more sophisticated features that encode prior theoretical work as well as more informed features that incorporate lexical knowledge:

- the literary analysis of poems is a natural place to look for theoretical insights regarding how words sound. Swooshing past the staggering volume of relevant work through the centuries, we assume the framework developed by Kaplan & Blei (2007) for the computational analysis of the phonemic, orthographic, and syntactic features of English-language poems. Of these features, *alliteration* (as in *Peter Pan*), *consonance* (*Freddy Krueger*), and *assonance* (*Frodo*) can be directly applied to names outside any context. Such features encode dependencies longer than what can be discovered by n-gram modelling, so that the prior knowledge that alliteration, consonance, and assonance might be relevant needs to be encoded in pre-computed features.

- An increasing volume of work investigates *phonological iconicity*, the existence of non-arbitrary relations between phonetic representation and semantics. The

Table 1: Inter-annotator agreement on the role polarity task.

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

findings are often be based on individual works and not generalizable, for example Miall (2001) notes that passages about Hell from Milton's *Paradise Lost* contain significantly more front vowels and hard consonants than passages about Eden, which contain more medium back vowels. When some generalizations have been made, these can be contradictory. Auracher et al. (2011), for instance, found that across different languages (including remote ones), nasals relate to sadness and plosives to happiness, parallels across remote languages, which might be consistent with the earlier finding that sonorants (including nasal /m/) is more common in tender poems (Fonagy 1961) but contradicts another previous finding that plosives correlate with unpleasant words (Whissell 1999). Clearly the relevant discussion in literature and psychology is far from mature, but there is growing evidence that phonological iconicity is a real phenomenon worth investigating in the context of names.

- Soundex phonetic distance and Levenshtein lexicographic distance to positive or negative terms in SentiWordNet (Esuli & Sebastiani 2006), a linguistic resource for *sentiment analysis* that annotates WordNet terms with an estimated degree of positive, negative or neutral hue. This makes character polarity prediction aware of the negative sentiment in *Darth Vader* through its similarity to the term *dark*.

- Socio-linguistic pragmatics, such as familiarity and gender. How familiar a name sounds is estimated via the frequency of its appearance in the *Social Security Death Index (SSDI)*, the publicly available database of all deceased US citizens since 1936. First names were matched against gender by scraping male and female first names from the multitude of Web sites that list baby names for prospective parents. This gives character polarity prediction access to the information that *Jane Smith* is a common female name and *Hannibal Lecter* is a rare male name.

Papantoniou & Konstantopoulos (2016) used these features to learn from their manually annotated corpus a decision tree that predicts character polarity (Table 2). Movie metadata such as genre and crediting order were expected to be very good discriminants, and were also included to the feature set for comparison only.

By comparing the performance of all features (F = 82%), only metadata (F = 71%), and all name-intrinsic features (excluding metadata, F = 80%), we can immediately understand that name-intrinsic features are better discriminants than metadata. This validates the core hypothesis that there is a correlation between what fictional character names look and sound like and the role they play in the plot of the fictional work they appear in. And among all intrinsic features, the phonetic ones appear to be the best discriminants. In fact, removing any other feature category *increases* performance, leading us to believe that all other features are actually adding noise (rather than discriminatory power) to the feature space.

Table 2: Performance of polarity prediction for different feature settings.

	Recall	Precision	$F_{\beta=1}$ score
Without IMDB metadata	80%	80%	80%
Only metadata	73%	70%	71%
Only phonetic features	79%	79%	79%
Without poetic features	84%	83%	83%
Without consonance feature	82%	82%	82%
Without SentiWordNet features	81%	81%	81%
Without phonetic features	80%	79%	79%
Without social features	81%	80%	80%
All features	82%	82%	82%

5 Concluding remarks

An interesting result was that the 'unfamiliar sounding' feature is not discriminative, refuting the hypothesis that the concept of the 'other' is stereotyped negatively. A more thorough investigation (and, in fact, one that is more inline with prior theories) will refine the 'unfamiliar' class into different ethnic backgrounds. Although not directly targeting any linguistic conclusions, from the wider humanities perspective such an investigation could give a tool for exploring whether 'bad guy' names in major US productions follow political developments to shift from German-sounding to Slavic-sounding to Arabic-sounding.

The result that was most relevant to language was the discriminative power of phonetic features. Although the current level of theoretical understanding of iconicity and its underlying mechanisms is far from complete, it helped formulate features and verify that they are discriminative. On the computational linguistics front, the findings presented here are also too focused on a particular language and domain to be a sound basis for grand generalizations, but they do point to various interesting directions. It would, for example, be interesting to extend the experiments to written literature to observe if there are 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 pushing earlier than the relatively young age of motion pictures.

Acknowledgements

I would like to gratefully acknowledge having used in this paper the results and insights obtained from the excellent work on language identification by Eirini Florou and on polarity prediction by Katerina Papantoniou for their MSc theses.

References

Algeo, John. 2010. Is a theory of names possible? *Names* 58(2). 90–96. http://dx.doi.org/10.1179/002777310X12682237915106.

Ashley, Leonard R. N. 2003. *Names in literature*. B: Authorhouse.

Auracher, Jan, Sabine Albers, Yuhui Zhai, Gulnara Gareeva & 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. http://dx.doi.org/10.1080/01638531003674894.

Butler, James Odelle. 2013. *Name, place, and emotional space: themed semantics in literary onomastic research*. University of Glasgow PhD thesis.

Chen, Lindsey N. 2008. Ethnic marked names as a reflection of United States isolationist attitudes in Uncle $crooge comic books. *Names* 56(1). 19–22. http://dx.doi.org/10.1179/175622708X282901.

Esuli, Andrea & 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)*, 417–22.

Florou, Eirini & Stasinos Konstantopoulos. 2011. A quantitative and qualitative analysis of Nordic surnames. In *Proceedings of the 18th Nordic Conference of Computational Linguistics (NODALIDA 2011), Riga, Latvia, 11–13 May 2011*, vol. 11 (NEALT Proceedings Series), 74–81. http://hdl.handle.net/10062/17292.

Fonagy, Ivan. 1961. *Communication in poetry*. William Clowes.

Font Llitjós, Ariadna & Alan W. Black. 2001. Knowledge of language origin improves pronunciation accuracy of proper names. In *Proceedings of Eurospeech 2001, Aalborg, Denmark*.

Gutiérrez, Elkin Darío, Roger Levy & Benjamin K. Bergen. 2016. Finding non-arbitrary form-meaning systematicity using string-metric learning for kernel regression. In *Proceedings of the 54th Annual Meeting of the ACL (ACL 2016), Berlin, Germany, 7–12 August 2016*, 2379–88.

Huang, Fei. 2005. Cluster-specific named entity transliteration. In *Proceedings of Human Language Technology Conference and Empirical Methods in Natural Language Processing Conference (HLT/EMNLP 2005), Vancouver, BC, Canada*, 435–442.

Kaplan, David M. & 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)*, 553–8. http://dx.doi.org/10.1109/ICDM.2007.76.

Konstantopoulos, Stasinos. 2007. What's in a name? In *Proceedings of Computational Phonology Workshop, held at the International Conference on Recent Advances in NLP (RANLP 07), Borovets, Bulgaria, September 2007*.

Konstantopoulos, Stasinos. 2010. Learning language identification models: a comparative analysis of the distinctive features of names and common words. In *Proceedings of the 7th International Conference on Language Resources and Evaluation (LREC-2010), Valletta, Malta, 19–21 May 2010*.

Markey, T. L. 1982. Crisis and cognition in onomastics. *Names* 30(3). 129–142. http://dx.doi.org/10.1179/nam.1982.30.3.129.

Miall, David. 2001. Sounds of contrast: an empirical approach to phonemic iconicity. *Poetics* 29(1). 55–70.

Monaghan, Padraic, Richard Shillcock & Morten H. Christiansen. 2014. How arbitrary is language? *Phil. Trans. of the Royal Society of London B* 369(1651).

Nicolaisen, William F. H. 2008. On names in literature. *Nomina* 31. 89–98.

Papantoniou, Katerina & Stasinos Konstantopoulos. 2016. Unravelling names of fictional characters. In *Proceedings of the 54th Annual Meeting of the ACL (ACL 2016), Berlin, Germany, 7–12 August 2016*.

Ramachandran, Vilayanur S. & Edward. 2001. Synaesthesia: a window into perception, thought and language. *Journal of Consciousness Studies* 8(12). 3–34.

Rudnyckyj, Jaroslav B. 1959. Function of proper names in literary works. *Internationalen Vereinigung für moderne Sprachen und Literaturen* 61. 378–383.

Shillcock, Richard, Simon Kirby, Scott McDonald & Cris Brew. 2001. Filled pauses and their status in the mental lexicon. In *Proceedings of the ISCA Tutorial and Research Workshop on Disfluency in Spontaneous Speech*.

Spiegel, Murray F. 1985. Pronouncing surnames automatically. In *Proceedings of the 1985 Conference of the American Voice Input/Output Society*, 109–132.

Tamariz, Monica. 2008. *Exploring the adaptive structure of the mental lexicon*. Edinburgh University PhD thesis.

Virga, Paola & Sanjeev Khudanpur. 2003. Transliteration of proper names in cross-lingual information retrieval. In *Proceedings of the Workshop on Multilingual and Mixed-language Named Entity Recognition, held at ACL 2003*, 57–64.

Whissell, Cynthia. 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. http://dx.doi.org/10.2466/pms.1999.89.1.19.