Community detection in inflectional networks

Maria Copot (m.copot@surrey.ac.uk)
Surrey Morphology Group, University of Surrey
Stag Hill, Guildford, Surrey, GU2 7XH, England

Andrea D. Sims (sims.120@osu.edu)
Department of Linguistics, The Ohio State University
1712 Neil Avenue, Columbus OH, 43210 USA

Abstract

Inflection classes partition the lexicon, classifying lexemes into groups based on shared inflectional exponence; as such, they are foundational aspects of lexical organisation. The task of identifying good partitions is non-trivial, however. We show that a modularity-based community detection method, borrowed from network theory, allows for identification of intuitive partitions at different granularities of representation. Applied to French verbs and Bosnian/Croatian/Montenegrin/Serbian nouns, the method detects inflection classes at multiple granularities and reveals (imperfect) hierarchical organisation of these. Ultimately, community detection methods facilitate more nuanced understanding of the inflectional organisation of the lexicon.

Keywords: lexicon; morphology; inflection classes; hierarchical organisation; network science; community detection

Introduction

Psycholinguistic studies have provided robust evidence for a memory-rich lexicon structured by associative relations (Kapatsinski, 2018). Within linguistic theory, this has helped fuel development of cognitively-oriented, usage-based models which hypothesise grammar to be emergent from this lexical structure (Rácz et al., 2015). Given this view of language as a complex system, modeling the structure of lexical relations might be expected to be at the top of linguists' agenda. Capturing both the micro-structure of a system and its macro-structure, and the relationship between the two, is relevant to explaining much of the behaviour of complex systems (Amaral & Ottino, 2004). Yet development of tools for this modeling has been surprisingly sparse in linguistics, at least so far as the macro-structure of the inflectional lexicon is concerned. Models of inflectional productivity have shown that at a granular level, the structure of lexical neighbourhoods can exert a powerful influence on generalisability (e.g. Albright & Hayes, 2003; Bybee & Moder, 1983). However, surprisingly little work has asked how the inflectional relations over which analogical inference happens are related to system-level inflectional organisation.

In this paper we home in on one particular issue that is of central importance for modeling of inflectional analogy: the organisation of inflection class (IC) systems. Analogical operations are generally seen as the result of gradient, similarity-based inference across lexemes. Nevertheless, morphologists frequently model the inflectional relations over which this analogical inference happens as hierarchically structured. In other words, ICs exhibit both a micro-structure and a macro-structure, suggesting that modeling of inflectional generalisation may require attention to both. A foundational question is thus: how are inflectional microclasses related to inflectional macroclasses?

We borrow tools from network theory and show their usefulness for modeling IC structure. In particular, we model inflection in terms of a bipartite network connecting inflectional exponents to lexemes, and apply a modularity-based community detection method to the classification of lexemes into ICs. The method can produce more coarse-grained or more fine-grained classifications, allowing different analyses of the same underlying data to be compared. Moreover, while previous methods for IC identification have presupposed hierarchical grouping, in our method the communities (ICs) found at one granularity of representation are independent of those found at another. By comparing the classes found at different granularities, we are thus able to empirically test for hierarchical structure. Applied to French verbs and Bosnian/Croatian/Montenegrin/Serbian (BCMS) nouns, the method reveals (mostly) hierarchical organisation. Our network-based approach thus produces precise representations of both the micro-structure and macro-structure of IC systems, and facilitates more nuanced questions about the dynamics of inflectional generalisation.

Inflection class micro- and macro-structure

ICs are partitions of the lexicon that classify lexemes (or in some approaches, stems) into non-overlapping sets. Partitioning is based on identity (or sometimes less strictly, similarity) of inflectional exponence among some set of lexemes and on how strongly those lexemes' exponence distinguishes them from other lexemes. In canonical ICs, lexemes have identical exponence class-internally and do not share any exponents with other classes (Corbett, 2009). The canonical case is rarely if ever encountered, however. More typically, classes exhibit *partial* inflectional identity.

Morphologists generally model ICs in terms of a tree structure with default inheritance (Brown & Hippisley, 2012; Corbett & Fraser, 1993; Dressler et al., 2006). Inflectional exponence shared widely across classes is specified high in the tree and inherited by default by lower nodes, all the way to individual lexemes at the leaves. Inflectional exponence associated with a specific class is specified at lower nodes. Re-

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

cently, Beniamine (2018, 2021) has argued that inflectional information is better modeled as a semi-lattice, and thus in terms of multiple default inheritance.

Inheritance hierarchies are designed for modeling the structure of inflectional exponence, but the relationship between inheritance hierarchies and IC partitions of the lexicon is not direct, especially if multiple inheritance is allowed. For one thing, nodes in an inheritance hierarchy (whether a tree or lattice) capture inflectional properties shared by a set of lexemes, but any properties that differentiate the lexemes are irrelevant. Nodes thus do not define classes in ways that maximize differentiation to other classes. Additionally, in a multiple inheritance approach in particular, nodes in a lattice do not partition the lexicon into non-overlapping sets. Because a given node can inherit from multiple parent nodes, lexemes belong, in some sense, to multiple, not-hierarchically-related groups at the same time.

Theoretical morphologists' interests have thus been covertly shifting away from lexical partitioning towards modeling the flow of information in the exponent system. However, understanding dynamics that underlie lexical partitioning and how they are to be modeled remains crucial for research interested in understanding operations over lexical partitions. For example, the study of how inflectional systems evolve over time centers on how partitions merge, divide or reorganise, whether it be in the language acquisition process of an individual (e.g. Clahsen et al., 2002) or in the history of a language family (e.g. Hinzelin, 2021).

Thus, despite much interest within morphological theory in hierarchical inflectional structure, there is a disconnect between work on inflectional inheritance and questions about the dynamics of lexical relations in inflectional generalisation, acquisition, and change. Whether ICs (under a traditional definition as lexical partitions) are hierarchically structured, and the extent of this within and across languages, is something that remains to be empirically investigated.

Automated inflection class detection

The task of partitioning lexemes into ICs is non-trivial. As already noted, classes in an IC system tend to exhibit partial inflectional identity. In a large lexicon, optimising partitions may therefore be too complex a task for manual analysis. Additionally, while grammars and theoretical analyses often present IC analyses as straightforward, this is undercut by the fact that it is common to encounter different analyses of the same data. For example, analyses of Russian nouns mostly posit three classes (not always the same three), but sometimes two or four (Corbett, 1982). Parker & Sims (2020) show that in a more fine-grained analysis, up to 82 microclasses can be distinguished. Beniamine (2018) finds 159 classes. Since different analyses are often possible, and each may (or may not) capture important aspects of inflectional organisation, we want to be able to generate and evaluate multiple analyses, and especially, analyses at different granularities.

Partitioning methods that are both automated and capture different granularities of class organisation have mostly been a bottom-up set-theoretic endeavour: approaches have consisted of specifying the smallest inflectional subgroupings, and then progressively combining these into supersets (Beniamine et al., 2017; Lee, 2014; Sagot & Walther, 2013). For example, Beniamine et al. (2017) present a method for bottom-up clustering of microclasses into macroclasses, using Minimum Description Length at each stage of grouping to determine whether a grouping makes the description of the system shorter. While these methods are a vast improvement over manual methods of analysis, they inherently presuppose the hierarchical organisation of classes.

We share a goal with this previous work: a bottom-up method of IC identification that is based entirely on the extent of inflectional identity among lexemes, so that, as Beniamine et al. (2017, 483) put it, "... any preconceived idea about other properties that macroclasses have can be tested empirically". However, hierarchical organisation is itself exactly the kind of "preconceived idea" that we might in fact want to test empirically. An ideal method would thus infer macroclasses at different granularities of representation without presupposing any particular relationship between representations.

An underexplored approach is to attempt the induction of inflectional classes over the lexical network itself. In the following section, we present such a method, showing that it has these desirable properties. We take as inspiration Sims (2020), a network-based comparison of IC systems based on pre-defined microclasses. However, we improve on Sims's approach by modeling the relations between lexemes and exponents directly. By doing so, we are able to do automated detection of the classes themselves.

Conceptual overview

ICs are partitions of the lexicon that seek to maximize identity of inflectional exponence within a class, and also maximize differentiation from other classes. In network terms, this can be formulated as a problem of community detection. Detecting the community structure of networks is a problem that arises in virtually all large networks, with applications in both the hard and social sciences and in computing (for an overview, see Fortunato & Hric, 2016).

Morphological systems as bipartite networks

Network theory, a mathematical framework for dealing with objects and the relationships that exist between them, lends itself well as a formalism for the lexicon. A network is a mathematical structure consisting of objects (nodes) and connections between them (edges). Because we are interested in the exponence patterns that characterise individual lexemes within a morphological system, we employ a bipartite network: nodes are grouped into two partitions (lexemes in one and exponents in the other, with the latter tagged for the paradigm cell they appear in), and edges can only connect nodes belonging to different partitions. Each edge thus represents the fact that a given lexeme takes a given expo-

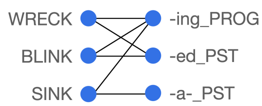

Figure 1: A toy bipartite network of three English verbs, with lexemes in the left partition and exponents in the right.

nent.[1] A toy illustration of this bipartite network structure is given in Figure 1, with three English lexemes (e.g. WRECK) and three exponent-cell combinations (e.g. -ing_PROG). This representation provides a lossless description of the system, where any higher-level structure is secondary to the properties of the connectivity patterns therein. Our representation of morphological systems shares a family resemblance with exemplar-based and analogical approaches, in that it models local relations using surface-based representations.

Lexical partitioning as community detection

We operationalise the task of partitioning the lexicon as one of community detection (Fortunato, 2010), a specialised form of clustering based on network-structured data. Specifically, we employ community detection based on modularity (Girvan & Newman, 2002; Newman & Girvan, 2004). The modularity of a network is defined as the fraction of edges that connect nodes within defined modules, minus the fraction of edges that would be expected to connect these same nodes if edges were distributed at random, i.e. in a network with a random level of clustering. Networks with high modularity have dense connections (edges) internally to modules (i.e. among nodes belonging to the same module), and sparse edges between modules. Detection of communities involves modularity optimisation, for which there are different techniques. The output of the optimisation process (ideally) yields groups of nodes that are more tightly connected to each other than to nodes elsewhere in the network.

Methods

Datasets

We use BCMS nouns and French verbs to illustrate the method.[2] These were chosen because large, high-quality datasets are available and because we have deep knowledge of the languages, facilitating interpretation of results.

BCMS BCMS noun paradigms have 12 cells, reflecting 6 cases and 2 numbers. Nouns are traditionally divided into four macroclasses; these are illustrated in Table 1. However, substantial variations on these patterns exist, ranging from regular subpatterns (e.g. animacy-conditioned syncretism in

	PROZOR 'window(M)'	SELO 'village(N)'	ULICA 'street(F)'	NOĆ 'night(F)'
NOM.SG	prozor	selo	ulica	noć
GEN.SG	prozora	sela	ulice	noći
DAT/LOC.SG	prozoru	selu	ulici	noći
ACC.SG	prozor	selo	ulicu	noć
VOC.SG	prozore	selo	ulico	noći
INS.SG	prozorom	selom	ulicom	noću

Table 1: Four-class description of BCMS noun inflection (singular forms only)

the PROZOR class), to quasi-regularities, to true exceptions. These variations multiply the number of microclasses.

Many exponents are shared across classes. For example, a quasi-regular epenthetic /a/ in genitive plural (*zemlja* 'ground', GEN.PL *zemalja*) may break stem-final consonant clusters in all but the NOĆ class. Overall, BCMS nominal (micro)classes are distinguished more by the particular way they mix and match exponents (suffixes, stem allomorphy) than by the uniqueness of their exponents.

Our dataset comes from Sims & Copot (in prep), a manually corrected version of nouns in the Serbo-Croatian UniMorph dataset (Batsuren et al., 2022), which itself consists of full paradigms for lexemes drawn from Wiktionary. Our dataset contains 10,927 noun lexemes. BCMS orthography reflects a consistent, virtually one-to-one relationship between letters and sounds. This allows us to work with the orthographic form directly.[3]

French French verb paradigms have 51 cells. The finite part has four moods (indicative, subjunctive, conditional, imperative), with the indicative having four synthetic tenses, the subjunctive having two and the other moods having one. Each synthetic tense-mood combination has six forms inflected for person and number, except for the imperative, which has three inflected forms. In addition, there are three nonfinite synthetic forms (present participle, past participle, gerundive).

Verbs are traditionally divided into three macroclasses identified by their suffix in the INF and PRS.PTCP: class 1 (INF: *-e*, e.g. MANGER), class 2 (INF: *-iʁ*, PRS.PTCP: *-isɑ̃*, e.g. FINIR), and all verbs that do not match these pattern belong to class 3, also called the class of "irregular verbs".[4] Table 2 shows these two cells in each of the 3 macroclasses.

Class 3 features the most internal variation, as it contains all the microclasses that do not closely follow the first or second conjugation. These instances can take the form of verbs featuring unique exponents not found elsewhere in the system (e.g. much of the conjugation of ALLER 'go') or combining exponents of the first two classes (e.g. OUVRIR 'open'). However, even the first and second conjugation fea-

[1] See Hofmann et al. (2020) for a conceptually similar approach to modeling derivation.

[2] Bosnian, Croatian, Montenegrin, and Serbian are distinct modern standard languages, but with historical periods of joint development (as Serbo-Croatian/Croato-Serbian). Our use of the term "BCMS" is purely practical: the nominal inflection is common to all four standards. It is not a claim about the statuses of the languages.

[3] However, the letters *lj*, *nj*, *dž*, and *dj* are two graphemes each; we replaced these with unique single graphemes for analysis.

[4] Irregular verbs like ALLER 'go' which, based on the diagnostics above, might be thought to belong to class 1 or 2, are assigned to this class.

	MANGER 'eat' (I)	FINIR 'finish' (II)	METTRE 'put' (III)	AVOIR 'have' (III)
INF	mɑ̃ʒe	finiʁ	mɛtʁ	avwaʁ
PRS.PTCP	mɑ̃ʒɑ̃	finisɑ̃	mɛtɑ̃	ɛjɑ̃

Table 2: Three-class description of French verb inflection, showing the cells traditionally used to distinguish macroclasses

ture small amounts of internal variation, such as lexicalised changes (e.g. ENVOYER 'send' has the irregular FUT/COND stem *enverr-*), leading to small amounts of variability. There is also overlap in the exponence patterns of the classes: for example, all verbs take -ɔ̃ as their suffix in the IND.PRS.1PL (for some, accompanied by an augment). All contrasts are segmental, and involve either suffixation or stem allomorphy. Defective cells have been identified by a special symbol.

We use a phonologically transcribed lexicon (Bonami et al., 2014) containing 5,274 verb lexemes.

Representing inflected words as a network

We translate the inflected words of a language into a network structure based on their exponent pattern. This process consists of finding the phonological material that characterises a paradigm cell for each lexeme, and then transposing lexemes and exponents into a bipartite network.

Word segmentation Starting from each language's inflectional dataset, we align the inflected forms using `morphalign` (Beniamine & Guzmán Naranjo, 2021). We then segment them into minimally discriminative units using `setmorph` (Beniamine & Carroll, 2023), and convert these into units resembling linguistic exponents by merging units that always appear in adjacent positions, and those which if combined would match an exponent found in a parallel form for another lexeme. To illustrate the process, take the French verb OBÉIR 'to obey', which has the PRS.IND.1PL *ɔbeisɔ̃*. After the stem *ɔbe-* has been identified, the formative for PRS.IND.1PL is deemed to be *-isɔ̃*. The formative is segmented into two exponents, *is, ɔ̃*. Following the observation that *-isɔ̃* is an exponent for other verbs in the same cell, and that the two exponents occur adjacently in this cell, they are fused into a single exponent, reflecting a common analysis of 2nd conjugation verbs.

To capture partial similarities between the exponence patterns of different lexemes, the exponent strings are then converted into a list of triphones. For example, MANGER 'eat' has PRS.IND.1PL *mɑ̃ʒɔ̃*, with exponent *-ɔ̃*. By turning *-isɔ̃* and *-ɔ̃* into sets of triphones, it is possible to capture the generalisation that both lexemes share the substring *-ɔ̃* in their exponent for the PRS.IND.1PL, while still capturing that OBÉIR has an augment that sets it apart from MANGER.

Setting up the network We set up a weighed bipartite network with `networkx` (Hagberg et al., 2008), in which lexemes and exponent triphones marked for the cell they belong to constitute the nodes for each partition. Lexemes share an edge with a triphone if the triphone features in their exponence pattern. To account for the fact that some exponents are longer than others, we ensure that the weights of all edges pertaining to the same exponent for a certain lexeme sum to 1. We obtain large, dense networks that capture the inflectional behaviour of French and BCMS lexemes and similarities in lexemes' exponence patterns.

Summary of the data structure

Table 3 summarizes the properties of the BCMS and French data. Edges is the number of edges in the bipartite network, connecting exponent triphones (tagged for the cell they appear in) in one partition to lexemes in the other partition.

	BCMS nouns	French verbs
Paradigm cells	12	51
Lexemes	10,927	5,274
Triphones	1,290	3,917
Edges	202,999	605,896

Table 3: Summary of network data structure

Community detection

We employ the Louvain method (Blondel et al., 2008) over the bipartite structure described above. Louvain is a greedy algorithm for finding communities which initialises every node in its own community and then iteratively groups nodes of lexemes and exponents together into larger communities, optimising for the modularity of the communities, iterating until improvements from each iteration fall below a threshold. We employ hard clustering, meaning that lexemes and exponents are clustered into non-overlapping communities.[5]

Optimization of modularity is known to have a resolution limit (Fortunato & Barthélemy, 2007): small modules are poorly detected in large, dense networks. One solution is to manipulate resolution by rescaling the network with self-loops for every node (Arenas et al., 2008). As Arenas et al. point out, rescaling a network to modify resolution is not just a technical work-around; rather, it allows network structure to be explored at different levels of granularity. This makes it well-suited to questions of hierarchical organisation.

The implementation of the function in `networkx` has a resolution parameter, ranging from 0 to infinity; its role is to bias the Louvain algorithm in favour of smaller or larger communities, with higher values encouraging smaller communities. We apply the method to the network at resolutions ranging from 0.1 to 2, increasing in 0.1 intervals. If there truly is a single, objectively perfect partitioning of the data, the method should return those partitions no matter the resolution. If ICs instead exhibit different organisation at different scales of description, we should observe that the size of identified communities is sensitive to the resolution parameter.

[5] The default version of the Louvain algorithm assumes a unipartite graph. In ongoing work (Copot & Sims, in prep), we reproduce qualitatively the same results using a version of the algorithm optimised for clustering over bipartite graphs.

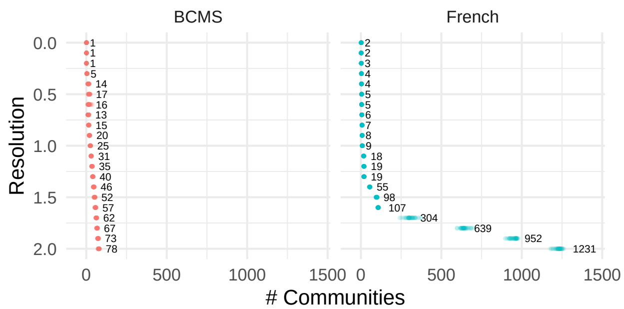

Figure 2: Number of communities for different resolutions. The partition set size for each iteration corresponds to a point (50 iterations, median partition set size indicated).

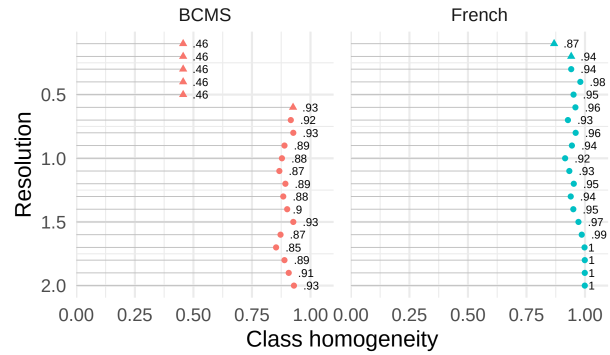

Figure 3: Class homogeneity score for different resolutions, by language. Triangle dots correspond to resolutions where only one class has been identified.

As the Louvain algorithm is nondeterministic, we perform the community detection process 50 times per language per resolution, to check for stability. This produces a distribution of number of communities found per resolution and allows us to verify whether all resolutions produce equally stable results. Once stability has been ascertained, we perform the rest of the analysis on a single representative set of communities.

Results

Community structure

This section describes the observed community structure of the two systems. Figure 2 shows that in BCMS, the method has trouble finding any partitions for low resolutions, reflecting the highly interconnected nature of the exponence system. However, as resolution increases, the number of partitions also increases. Overall, the number of partitions found while keeping other parameters constant is always highly stable for BCMS. In French, even for the lowest resolutions the lexicon features at least two partitions in most iterations, owing to the better delineated nature of classes in the system. The number of partitions increases slowly with resolution, until resolutions greater than 1.6, for which the number of classes shows a dramatic increase compared to the next coarser resolution. This increase is coupled with high variability in the precise number. We expect variability in the number of classes found if the structure at a particular level of granularity does not lend itself well to clear partitioning, and conversely expect consistency in the presence of clearly delineated structure. The results thus suggest that for high resolutions, the method is overpartitioning the lexicon in French, attempting to find distinguishing structure where it is not clear there is any.

Comparison to traditional descriptions

We evaluate these results by comparing the communities found by our method to the macroclasses traditionally identified by linguists. This constitutes a verification that our partitions respect the macro-level description of the linguistic systems. We calculate a class homogeneity score for each partition which quantifies how homogeneous the partition is in terms of its traditional class make up. For each partition, we identify the most frequent traditional class within it and calculate the percentage of lexemes within the partition belonging to said class. This is the partition's homogeneity score. We average this across partitions in the same resolution to obtain a single number characterising class homogeneity.

The results are shown in Figure 3. Overall, high scores show that the partitions found are generally homogeneous in their makeup, suggesting that they either match traditional descriptions of the inflectional systems or (more commonly) make finer distinctions internal to traditional classes.

There are two instances of scores deviating from the 0.85-0.99 range, both of which we interpret as artefacts of how the score is calculated. The comparatively low scores for the lowest resolutions in the BCMS data are reflective of the method finding only one community; the score therefore corresponds to the percentage accounted for by the traditional class with the highest type frequency. Conversely, the perfect homogeneity scores for the highest resolutions in the French data are reflective of the high number of partitions, resulting in many partitions containing single lexemes, which bias the homogeneity scores towards the ceiling due to overpartitioning, as was observed in the previous section.

Hierarchical relationship between micro- and macro-structure

We test for the hierarchy assumption by measuring whether partitions in adjacent resolutions are in sub-/superset relationships, or whether lexemes are partitioned in ways that reflect non-hierarchical organisation. For each pair of adjacent resolutions, we compute a hierarchy score: for each partition within the finer resolution we identify all lexeme pairings and check whether the pair is found in the same partition in the coarser resolution; if this is verified, the partition hierarchy score increases by 1. The score is then divided by the number of items in the partition, and averaged across partitions. We perform the analysis on a representative iteration of the community detection method.

The hierarchy scores are shown in Figure 4. Whenever

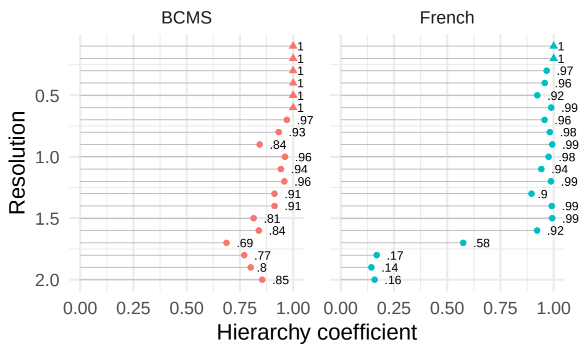

Figure 4: Hierarchy score for adjacent partitions by language. Triangle points represent instances where the coarser-grained resolution has only one partition.

multiple partitions are involved, hierarchy scores are below 1, sometimes considerably, reflecting deviations from hierarchical structure in the inflectional systems. Overall, however, we see evidence for a fairly high amount of hierarchical organisation, especially in the BCMS data. Resolutions above 1.6 in French once again stand out, with exceptionally low hierarchy scores likely being linked to the somewhat aleatory nature of the overpartitioned space.

Discussion

In this paper, we provided a methodology to model a morphological system as a bipartite graph between lexemes and exponents, which provides a lossless relational representation of system-internal patterns. We then applied the Louvain community detection algorithm to exponent networks of French verbs and BCMS nouns, where we found stable community structure at multiple levels of granularity. These empirically-derived communities at multiple levels of granularity are largely compatible with the macrostructure posited by linguists for each system. On the matter of hierarchical representation, we found that while a degree of hierarchical organisation exists in both of the systems we analysed, it is a tendency rather than a strict requirement, and the strength of the tendency does not have a monotonic relationship to the granularity of the distinctions made. In ongoing work (Copot & Sims, in prep), we further examine the hierarchy and stability of identified communities by comparing results to those observed in a randomly generated graph with similar structural properties, increasing the interpretability of our results.

The Louvain community detection algorithm is one of many tools for this purpose, as are the analytical procedures we employ in this paper. We leave the question of how much of a difference is made by specific implementational choices to future research. The Louvain community detection algorithm is known to not be equally good at detecting partitions at all resolutions, a fact borne out by our data: we hypothesise that French is overpartitioned by the algorithm at high granularities, based on a manual inspection of results combined with the low community number stability and low hierarchy scores for French. Moreover, we suspect that the presence of a single community for low granularities in BCMS is an instance of underpartitioning, a result of attempting to partition a system where distinctions between inflectional classes are more gradient. We leave more thorough exploration of the invalidity of certain partitionings to future research.

In larger view, part of our goal with this project is to develop methods for analysing lexical relations that facilitate predictions about the dynamics of inflectional generalisation. In a study of the phonological network structure of the lexicon, Siew & Castro (2023) find that speakers' word similarity judgments depend on the community structure of the network and path length.[6] They interpret this effect in terms of spreading activation mechanism defined by network connections. In the domain of morphology, Baayen (2010) finds that in a directed network of English compound elements (e.g. *black* and *board* in *blackboard*), distant lexical neighbours as defined by network structure have an inhibitory effect on compound word naming. These and other studies thus suggest that speakers are sensitive to the large-scale structure of lexical networks, beyond local neighbourhoods. Given the importance of neighbourhoods and similarity-based inference also in inflectional generalisation, learning, and change, we might wonder whether similar effects exist in this domain. Though this paper focuses on the class-level network structure, one advantage of inducing the classes over the lexical network is that it allows for simultaneous inspection also of node-level properties, including a lexeme's relationship to its class (e.g. whether it is peripheral or central to the class) and also its relationships to other lexemes. This allows for predictions about the generalizability of morphological patterns at the level of the lexeme, to the extent that these derive from the position of lexemes within their communities. While specific predictions are beyond the scope of the present paper, a first step in the direction of asking such questions is to be able to model the network structure. We view community detection methods as promising in this respect.

We hope that this paper offers pointers to the broader potential that network science holds for morphological research. Network theory is a promising toolkit for investigating the internal organisation of the lexicon, focusing on quantifying properties of the system and the role played by items within it. Thanks to these desirable properties, it is beginning to make inroads in modeling of lexical relations at different levels of grammar (Chern et al., 2024; Siew, 2013; Parker et al., 2022; Sims, 2020; Turnbull, 2021, 2023) and modeling of language processing (Baayen, 2010; Pham & Baayen, 2015; Siew & Castro, 2023). In broader terms, a modeling approach rooted in network science has the benefit of connecting linguistic theory to a large literature on the emergent properties of networks and the dynamics of complex systems.

[6]Path length is the distance between two points in a network, specifically, the number of edges that must be traversed in order to get from one target node to another.

Acknowledgements

The research presented in this paper was completed thanks to an Ohio State President's Postdoctoral Scholarship and a British Academy International Fellowship (IF24\100591 - *Seeing both the forest and the trees: lexical organisation beyond lexemes and cells*), both awarded to Maria Copot, and a sabbatical granted to Andrea Sims by The Ohio State University. We gratefully acknowledge this support. We also thank the two anonymous reviewers for their feedback.

References

Albright, A. C., & Hayes, B. P. (2003). Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition*, *90*, 119–161.

Amaral, L. A., & Ottino, J. M. (2004). Complex networks: Augmenting the framework for the study of complex systems. *The European Physical Journal B*, *38*, 147–162.

Arenas, A., Fernández, A., & Gómez, S. (2008). Analysis of the structure of complex networks at different resolution levels. *New Journal of Physics*, *10*, 053039.

Baayen, R. H. (2010). The directed compound graph of English: An exploration of lexical connectivity and its processing consequences. In S. Olsen (Ed.), *New impulses in word-formation* (pp. 383–402). Buske.

Batsuren, K., Goldman, O., Khalifa, S., Habash, N., Kieraś, W., Bella, G., … Vylomova, E. (2022). UniMorph 4.0: Universal Morphology. In N. Calzolari et al. (Eds.), *Proceedings of the Thirteenth Language Resources and Evaluation Conference* (pp. 840–855). Marseille, France: European Language Resources Association.

Beniamine, S. (2018). *Classifications flexionnelles. Étude quantitative des structures de paradigmes*. Thèse de doctorat, Université Sorbonne Paris Cité - Université Paris Diderot (Paris 7), Paris, France.

Beniamine, S. (2021). One lexeme, many classes: Inflection class systems as lattices. In B. Crysmann & M. Sailer (Eds.), *One-to-many relations in morphology, syntax and semantics* (pp. 23–51). Berlin: Language Science Press.

Beniamine, S., Bonami, O., & Sagôt, B. (2017). Inferring inflection classes with description length. *Journal of Language Modelling*, *5*(3), 465–525.

Beniamine, S., & Carroll, M. (2023). The other perspective on exponence. Presented as the International Symposium of Morphology, Nantes, France.

Beniamine, S., & Guzmán Naranjo, M. (2021). Multiple alignments of inflectional paradigms. *Proceedings of the Society for Computation in Linguistics*, *4*, 216–227. doi: 10.7275/ymc0-p491

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, *2008*(10), P10008. doi: 10.1088/1742-5468/2008/10/p10008

Bonami, O., Caron, G., & Plancq, C. (2014). Construction d'un lexique flexionnel phonétisé libre du français. In F. Neveu, P. Blumenthal, L. Hriba, A. Gerstenberg, J. Meinschaefer, & S. Prévost (Eds.), *Actes du quatrième congrès mondial de linguistique française* (pp. 2583–2596).

Brown, D., & Hippisley, A. (2012). *Network morphology: A defaults-based theory of word structure*. Cambridge: Cambridge University Press.

Bybee, J., & Moder, C. (1983). Morphological classes as natural categories. *Language*, *59*(2), 251–270.

Chern, J., Castro, N., & Siew, C. S. Q. (2024). Evidence of community structure in phonological networks of multiple languages. *Canadian Journal of Experimental Psychology / Revue canadienne de psychologie expérimentale*, *ahead of print*. doi: 10.1037/cep0000357

Clahsen, H., Aveledo, F., & Roca, I. (2002). The development of regular and irregular verb inflection in Spanish child language. *Journal of Child Language*, *29*(3), 591–622. doi: 10.1017/s0305000902005172

Copot, M., & Sims, A. D. (in prep). *Crosslinguistic comparison of community structure in lexical networks.*

Corbett, G. G. (1982). Gender in Russian: An account of gender specification and its relationship to declension. *Russian Linguistics*, *6*, 197–232.

Corbett, G. G. (2009). Canonical inflection classes. In F. Montermini, G. Boyé, & J. Tseng (Eds.), *Selected proceedings of the 6th Décembrettes* (pp. 1–11). Somerville, MA: Cascadilla.

Corbett, G. G., & Fraser, N. (1993). Network Morphology: A DATR account of Russian nominal inflection. *Journal of Linguistics*, *29*, 113–142.

Dressler, W. U., Kilani-Schoch, M., Gagarina, N., Pestal, L., & Pöchtrager, M. (2006). On the typology of inflection class systems. *Folia Linguistica*, *40*(1-2), 51–74. doi: doi:10.1515/flin.40.1-2.51

Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, *486*(3–5), 75–174. doi: 10.1016/j.physrep.2009.11.002

Fortunato, S., & Barthélemy, M. (2007). Resolution limit in community detection. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(1), 36–41.

Fortunato, S., & Hric, D. (2016). Community detection in networks: A user guide. *Physics Reports*, *659*, 1–44.

Girvan, M., & Newman, M. E. J. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, *99*(12), 7821–7826. doi: 10.1073/pnas.122653799

Hagberg, A. A., Schult, D. A., & Swart, P. J. (2008). Exploring network structure, dynamics, and function using networkx. In G. Varoquaux, T. Vaught, & J. Millman (Eds.), *Proceedings of the 7th Python in Science Conference (SciPy2008)* (pp. 11–15). Pasadena, CA.

Hinzelin, M.-O. (2021). Inflection classes in verbs in the Romance languages. In M. Aronoff (Ed.), *The Oxford research encyclopedia of linguistics.* Oxford: Oxford University Press. doi: 10.1093/acrefore/9780199384655.013.693

Hofmann, V., Schütze, H., & Pierrehumbert, J. B. (2020). A graph auto-encoder model of derivational morphology. In D. Jurafsky, J. Chai, N. Schluter, & J. Tetreault (Eds.), *Proceedings of the 58th Meeting of the Association for Computational Linguistics* (pp. 1127–1138). Online: Association for Computational Linguistics.

Kapatsinski, V. (2018). Words versus rules (storage versus online production/processing) in morphology. In M. Aronoff (Ed.), *The Oxford research encyclopedia of linguistics.* Oxford: Oxford University Press. doi: 10.1093/acrefore/9780199384655.013.598

Lee, J. L. (2014). *Automatic morphological alignment and clustering* (Technical Report No. TR-2014-07). Department of Computer Science, University of Chicago.

Newman, M. E. J., & Girvan, M. (2004). Finding and evaluating community structure in networks. *Physical Review E*, *69*, 026113.

Parker, J., Reynolds, R., & Sims, A. D. (2022). Network structure and inflection class predictability: Modeling the emergence of Marginal Detraction. In A. D. Sims, A. Ussishkin, J. Parker, & S. Wray (Eds.), *Morphological diversity and linguistic cognition* (pp. 247–281). Cambridge: Cambridge University Press.

Parker, J., & Sims, A. D. (2020). Irregularity, paradigmatic layers, and the complexity of inflection class systems: A study of Russian nouns. In F. Gardani & P. Arkadiev (Eds.), *The complexities of morphology* (pp. 23–51). Oxford: Oxford University Press.

Pham, H., & Baayen, H. (2015). Vietnamese compounds show an anti-frequency effect in visual lexical decision. *Language, Cognition and Neuroscience*, *30*(9), 1077–1095. doi: 10.1080/23273798.2015.1054844

Rácz, P., Pierrehumbert, J., Hay, J., & Papp, V. (2015). Morphological emergence. In B. MacWhinney & W. O'Grady (Eds.), *The handbook of language emergence* (pp. 123–146). Malden, MA: Wiley Blackwell.

Sagot, B., & Walther, G. (2013). Implementing a formal model of inflectional morphology. In *Proceedings of the Third International Workshop on Systems and Frameworks for Computational Morphology (SFCM 2013)* (pp. 115–134). Berlin: Humboldt-Universität and Springer.

Siew, C. S. Q. (2013). Community structure in the phonological network. *Frontiers in Psychology*, *4*, Article 553.

Siew, C. S. Q., & Castro, N. (2023). Phonological similarity judgments of word pairs reflect sensitivity to large-scale structure of the phonological lexicon. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *49*(12), 1989–2002.

Sims, A. D. (2020). Inflectional networks: Graph-theoretic tools for inflectional typology. *Proceedings of the Society for Computation in Linguistics*, *3*, 10. doi: 10.7275/c1f4-pg94

Sims, A. D., & Copot, M. (in prep). *A lexicon of Bosnian, Croatian, Montenegrin, and Serbian noun paradigms.*

Turnbull, R. (2021). Graph-theoretic properties of the class of phonological neighbourhood networks. In E. Chersoni, N. Hollenstein, C. Jacobs, Y. Oseki, L. Prévot, & E. Santus (Eds.), *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2021)* (p. 233-240). Online: Association for Computational Linguistics.

Turnbull, R. (2023). Phonological network properties of non-words influence their learnability. In R. Skarnitzl & J. Volín (Eds.), *Proceedings of the 20th International Congress of Phonetic Sciences (ICPhS 2023)* (pp. 4056–4060). Prague: International Phonetic Association.