Contrastive morphological typology and logical hierarchies

John Sylak-Glassman[1] and Ryan Cotterell[1,2]
[1]Johns Hopkins University, [2]Ludwig Maximilian University of Munich

1 Introduction

One of the central goals of research in morphological typology is to discover a list of categories (e.g. Tense) and features (e.g. present) that capture the concepts encoded by inflectional morphology across the world's languages.[1] Such a list would represent a substantive claim about the content of the inflectional morphological component of grammar. Defining its contents would clarify the interfaces of morphology with syntax and semantics by more clearly highlighting a division of labor among them. To build such a list, a principled method is necessary to justify the existence of these features and to both discover and define their content. We argue that a method based on overt surface contrast within a language can be used to justify the existence of features, and that comparison of similar features across languages can define the core meaning of that feature.

However, discovering such a list of features is not sufficient for understanding the structure of morphological distinctions. Prior work by Harley & Ritter (2002) highlighted the existence of inherent organization among inflectional features, demonstrating that a simple unordered list of morphological features is insufficient for accurately predicting typological and acquisitional patterns. The hierarchies discovered by Harley & Ritter (2002) are termed *dependency hierarchies* based on the principle underlying their organization, i.e. patterns of dependency observed consistently across languages. We argue that additional factors provide inherent organization to inflectional features. Based on evidence from facultative feature use, superclassing, grammaticalization patterns, and translational equivalence between forms, we argue specifically that inflectional features are also organized according to the specificity of their conceptual-semantic content. Features are organized into *logical hierarchies* based on the ability of certain features to subsume the meaning of other features.

First, we describe the method that we propose can be used to justify the existence of inflectional features and discover their content (§2). In addition to describing the method itself, we discuss the theoretical status of the features that the method uncovers. §3.1 reviews Harley & Ritter's (2002) proposal about morphosyntactic features and their organization. §3.2 motivates the existence of logical hierarchies and describes their organization based on language-internal evidence, diachronic evidence from grammaticalization pathways, and evidence from translational equivalence among inflected forms across languages with differing morphological contrasts within a category.

[1]*Acknowledgements*: Many thanks to the organizers of CLS 52 and to the anonymous reviewers and the audience for helpful suggestions.

2 Contrastive morphological typology

The goal of finding a list of features that captures the cross-linguistic concepts encoded by inflectional morphology requires a language-independent method with clear principles for delineating features and determining their content. The method proposed here relies on evidence from surface contrasts within single languages, similar to the contrastive method of discovering phonemes and phonological features outlined in Jakobson *et al.* (1952). The method we propose is as follows:

> If a morphological distinction is encoded by two overtly contrasting morphemes in a language (of which one may be phonologically null) and the meaning encoded by at least one of the morphemes is not decomposed further in any other language, then that non-decomposable (i.e. basic) meaning is represented by a feature.

A guiding idea behind the method proposed here is that features in the final set should include only inflectional morphological distinctions that are as basic as possible in that they cannot be further subdivided into other attested distinctions. Adhering to this principle allows more complex features to be built up additively or disjunctively.

Nominal case marking can be used to illustrate how the contrastive morphological feature discovery method proceeds and how features may be discovered to be non-basic. Suppose that the search for a list of nominal case features begins with a system like that of English, in which only a nominative and oblique case are distinguished overtly, as in *I*/*me*, *he*/*him*, *she*/*her*, *we*/*us*, and *they*/*them*. The two cases can be defined as encoding subjecthood (for nominative case) and any non-subject relation (oblique). From this evidence, nominative and oblique can be taken to be basic case features. However, upon further investigation, languages like Russian overtly mark cases, such as genitive, dative, accusative, instrumental, and prepositional, which capture relations that are encoded by the oblique in an English-like system. From this evidence, it can be concluded that oblique is not a basic nominal case feature. [2] Instead, nominative is a basic case along with the specific cases identified in Russian, for example. 'Oblique' may cease to be a case feature at all, or may be found to be an indeterminate feature (in the sense proposed by Corbett (2000:39), originally for number). That is, the meaning of an indeterminate feature like 'oblique' would depend on the existence of other features with more specific meanings, either on the level of a single language or cross-linguistically.

Using this method, a cross-linguistic survey of the inflectional morphological distinctions made by the world's languages was conducted. The survey started by finding categories to which agreement features belong, and then finding additional categories according to the part of speech on which they are most often expressed (e.g. tense on verbs). To find basic features, languages with the largest known number of distinctions within a category were examined. For example, Sursurunga makes very fine-grained distinctions in number, and distinguishes the singular, dual,

[2] Upon further investigation, some of the cases of Russian are also not basic for all of the functions in which they are used. For example, Uralic and Nakh-Daghestanian languages use extensive local case marking systems which designate the type of motion as well as the location (Radkevich 2010), and these further differentiate the functions of Russian's dative, accusative, and prepositional cases as they are used with location- or motion-designating prepositions.

paucal, greater paucal, and plural (Corbett 2000:26-29). These distinctions finely subdivide a scale defined by the number of countable entities. In addition to distinctions such as these, distinctions which are basic, but do not fall along a scale, also receive features. For example, in languages with a 'greater' or 'global' plural, the distinction between plurals cannot be described as a difference in where a division falls on a numerical scale, but as a distinction in an additional conceptual dimension, denoting not just more than one entity, but an abundance of entities or all possible entities (Corbett 2000:30). In keeping with the contrastive morphological feature discovery method, the scope of the survey was limited by considering overt affixal morphology and contrasts that were expressed paradigmatically.[3]

The survey resulted in a set of 277 features distributed across 25 morphological categories. These categories include: Aktionsart, animacy, argument marking, aspect, case, comparison/grade, definiteness, deixis, evidentiality, finiteness, gender, information structure, interrogativity, mood, number, POS, person, polarity, politeness, possession, switch-reference, tense, valency, voice, and a language-specific category. The large number of features was due not only to true cross-linguistic variation in distinctions, but to two other factors. First, because gender and noun-class systems are not necessarily organized according to semantic principles and those that are make different divisions, many gender and noun-class features are required, even when they are abstracted to the highest genealogical level possible (e.g. features are defined according to a linguistic stock, rather than an individual language). Second, some features, such as those for possession marking, are compositional in that they reflect systematic distinctions that may also be independent features. For example, to capture the marking of possession by a second person singular possessor, the compositional feature PSS2S is used. The feature's label reflects this compositionality, in that PSS- signifies 'quality of being possessed,' -2- signifies the second person, and -S signifies singular. Such compositional features are necessary to properly attribute characteristics to the possessor versus the possessed thing, e.g. one value of number will not suffice to express possession of a plural number of things by a singular possessor.

The theoretical status of these features is intermediate between what Haspelmath (2010) terms *universal cross-linguistic categories* and *comparative concepts*. Haspelmath (2010) describes universal cross-linguistic categories as being a finite, relatively small set of substantive categories that are universally available to all languages, can be equated across them, and are used for both description/analysis and typological comparison. By contrast, comparative concepts are defined explicitly by typologists to be used for comparison only, are not used for description or analysis of a particular language, and are not "psychologically real" (Haspelmath 2010:664-665). The features discovered by the contrastive method are assumed to be finite, substantive, and universally applicable like universal cross-linguistic categories, but like comparative concepts, they are assumed not to be psychologically

[3]Because of the intended use of the resulting feature set, some contrasts which are not marked by dedicated morphology and which may be lexical, such as animacy contrasts, are given features. For a similar reason, cross-linguistic features which may not be basic, such as indicative, subjunctive, realis, and irrealis are also included. These possibly non-basic features may be reducible to identifiable basic features, or may be truly indeterminate, capturing some (modal) senses that are not conveyed by dedicated morphology in any language.

real nor necessarily needed to describe particular languages. The most important similarity of the features proposed here with comparative concepts is that they have the limited range of conceptual-semantic definition used for comparative concepts. The features discovered via the contrastive method proposed here are assumed to share the cross-linguistically valid *core* (i.e. the intersection) of the concepts they encode in each language, not the entire set of meanings (i.e. the union) encoded cross-linguistically. For example, the feature for dative case means "encoding an indirect object, recipient, or patient," and this meaning is assumed to be present on any element associated with the DAT (dative) feature. Therefore, while many different languages may all use a category termed "dative" (Haspelmath 2010:665-666), a morpheme associated with the feature DAT is assumed to have at least the meaning of that feature, but not exclusively that meaning.[4]

Related to their intermediate role between universal cross-linguistic categories and comparative concepts, this set of features can be applied as a standardized form of glossing for inflectional morphemes across languages (like a more extensive set of the abbreviations set forth in the Leipzig Glossing Rules; Comrie *et al.* 2008). In this role, the set of features, called the UniMorph Schema, was used to label inflected word forms gathered from the English edition of Wiktionary[5] in a consistent way across languages. Full details on this specific use of the feature set gathered using the contrastive method described here can be found in Kirov *et al.* (2016) and Sylak-Glassman *et al.* (2015b). Full details on the UniMorph Schema itself and the features that comprise it are available as a user guide at `unimorph.org` and in Sylak-Glassman *et al.* (2015a).

3 Morphological feature organization

Although a list of morphological features is an important element in defining inflectional morphology and its interfaces with other components of grammar, the inherent organization of those features must be discovered. To the authors' knowledge, no theory of morphology has ever assumed completely unstructured bundles of features since all have made assumptions about basic category membership. These assumptions are so uncontroversial that it is easy to overlook category membership as being a form of feature organization. For example, features such as 1, 2, and 3 are assumed to be members of a coherent 'person' category, and features such as singular, dual, and plural are assumed to belong to a 'number' category. Feature organization can be much more complex than this. In the following sections, a previous approach to the hierarchical organization of features by Harley & Ritter (2002) is reviewed (§3.1), and a new approach that adds to those findings is presented in §3.2.

[4]It is for this reason that the Turkish 'dative' is glossed as DAT/ALL in Table 4. The Turkish dative is used not only to encode recipients, but also destinations of motion, which is covered by the allative feature (ALL).

[5]`en.wiktionary.org`

3.1 Dependency hierarchies

In an influential study of pronominal paradigms across the world's languages, Harley & Ritter (2002) propose a set of universal morphosyntactic features organized into a feature geometry that is used to account for typological and acquisitional generalizations. Based on data from a sample of "110 languages representing over 90 distinct families or subfamilies, four isolates, and three creoles," Harley & Ritter (2002:496) propose high-level, monovalent (i.e. privative) morphosyntactic features that encode distinctions in person, number, and gender under PARTICIPANT, INDIVIDUATION, and CLASS nodes, respectively. These nodes and their subtrees correspond to the "grammaticization of natural cognitive categories," namely reference, plurality, and taxonomy (ibid.:482,485).

The organization of features adopted by Harley & Ritter (2002) follows the principles first established by feature geometries in phonology. The presence of a feature lower down in a hierarchy entails the presence of a higher feature directly on the path to the geometry's root. In other words, the lower feature depends on the presence of any higher features. In addition, Harley & Ritter (2002) assume that non-terminal nodes without any terminal dependents receive a default interpretation.

Harley & Ritter (2002:497-498) use dependency to encode typological generalizations like those noted by Greenberg, such as the fact that "a language will not have a dual number if it does not have a plural number." Interestingly, the relationship of a lower feature to a higher feature is not used to directly encode the dual number in the expected way, i.e. an independent dual feature is not nested below a plural feature. Instead, the default interpretation of the individuation node, the feature Minimal, is used in conjunction with the Group feature, representing plural, to convey the dual. This makes the presence of dual directly contingent on the presence of plural, and explains a pattern termed 'constructed number' in Corbett (2000:169) that is used to form the dual in Hopi: The combination of a plural marked noun with a singular marked verb induces a dual interpretation. In Harley & Ritter's (2002) framework, dependency relationships are not only encoded directly through representational structure, but through representational complexity.

3.2 Logical hierarchies

Although dependency is a typical organizing principle not just in feature geometries, but in linguistics as a whole, it is not the only possible hierarchical organizing principle. While the hierarchy presented by Harley & Ritter (2002) accurately models typological generalizations, entailments on the existence of features in a system, and acquisitional generalizations, it does not capture intrinsic relationships between the conceptual-semantic content of features, e.g. that dual or paucal is a logical subset of plural. Harley & Ritter's (2002:485) argument that these logical relationships need not be represented in a feature geometry is valid and well-addressed. Nonetheless, language-internal evidence indicates that such logical relationships among features influence grammatical patterns. In addition, the relationships between morphosemantic features, which do not participate in agreement and government (Kibort 2010), must be inferred from evidence about the nature of their conceptual-semantic content.

In this section, we present evidence that morphological features are organized not just by dependency, but by specificity relationships among their conceptual-semantic content. The hierarchies built according to this principle are termed *logical hierarchies*. In them, features lower down in the hierarchy specify the parameter encoded by the dominating feature in a more specific way. For example, dual number, conveying 'exactly two,' is more specific than the plural number, which typically conveys only 'multiple' or 'more than 1.' The 'substance' that constrains these logical hierarchies is the specificity-generality relationships which arise from the content of morphological features. This section will proceed by first presenting two types of language-internal evidence for logical hierarchies, and will continue on by presenting diachronic and cross-lingual evidence.

3.2.1 Language-internal evidence

One source of language-internal evidence for logical hierarchies is facultative feature use, which occurs in Larike-Wakasihu (Austronesian), Slovene (Slavic), and Ngan'gityemerri (Southern Daly) (Corbett 2012:21-30). For example, in Larike-Wakasihu, four number distinctions are used: singular, dual, trial, and plural. Dual and trial are used for exactly two and three referents, respectively, and are not paucal. However, given three referents, speakers have the option of using either the trial or plural, and similarly, given two referents, speakers may use either the dual or the plural. In these cases, the dual and trial can be viewed as more specific values whose meaning is subsumed under the more general plural feature. In Ngan'gityemerri, which uses the same number distinctions as Larike-Wakasihu, the use of the dual is obligatory, but the use of the trial is optional, with plural able to replace it. In Slovene, which has only singular, dual, and plural, the dual can optionally be used in the same manner as in Larike-Wakasihu, i.e. it may be replaced by the plural (especially with less animate or salient referents). This data points to a logical hierarchy within the number category in which singular and plural are opposed and appear as top-level features with dual and trial opposed to each other and both dominated by plural, as shown in (1).

(1)

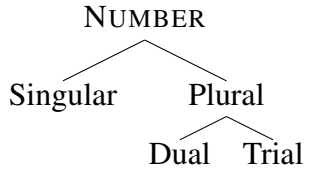

Another source of language-internal evidence for logical hierarchies comes from a phenomenon called superclassing, which is conceptually similar to facultative feature use in that it involves the optional use of available distinctions. The example cited in Corbett (2012:22-24) is from Jingulu (Mirndi; Australia), with a very similar case attested in Bininj Gun Wok (Mayali; Australia). Jingulu has four gender values: Masculine, feminine, neuter, and vegetable (for edible plants). Full gender agreement is possible on adjectives, with all four values distinguished. However, a reduced pattern of agreement is also possible in which masculine can be used in place of feminine and neuter can be used in place of vegetable, creating a system in which only masculine and neuter are opposed. Finally, agreement can be done

away with all together, with only a masculine ending used (23). This type of evidence can be interpreted similarly to that from facultative feature use as a case in which a less specific feature is used in place of a more specific one. Although the determinacy of masculine and feminine genders makes this interpretation problematic on conceptual-semantic grounds, the evidence from this phenomenon points to a logical hierarchy like that in (2) for gender.[6]

(2)

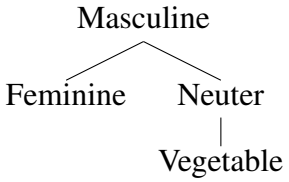

Both facultative feature use and superclassing offer evidence that the distinctions conveyed by certain features can be subsumed into others, thereby implying a specific-general relationship.

3.2.2 Generality relationships in grammaticalization pathways

The historical genesis or loss of distinctions provides another source of evidence for the kind of specific-general relationships that underlie logical hierarchies. We surveyed the *World Lexicon of Grammaticalization* (Heine & Kuteva 2002), a database of over 400 grammaticalization processes, for any processes that resulted in a change in specificity between the features involved. The survey was limited to processes in which at least either the source or target in the process clearly belonged to morphological categories identified in the UniMorph Schema using the contrastive method and survey described in §2. The processes found this way were further filtered by limiting comparison to processes in which both the source and target features likely belonged to the same category. The survey excluded processes in which one side was a free lexeme (e.g. 'say' $\rightarrow$ hearsay evidential) or a purely syntactic category (e.g. VP-and $\rightarrow$ subordinator).

This survey resulted in 16 processes, of which 4 were reported by the data source as involving a specific$\rightarrow$general relationship, none were reported to have a general$\rightarrow$specific relationship, and 12 were judged by the authors to involve either kind of specificity relationship. These processes are shown in Table 1.

Of these 16 processes, 9 involve case, 3 involve mood, 2 involve valency, and 2 involve aspect. If we restrict this evidence to that which can be directly represented by UniMorph features, the fragments of hierarchies in (3) emerge. In these fragments, we have not included the link between benefactive and purposive since the relationship between dative and purposive is unclear. We have also not included manner, anticausative, material, concessive, and patient, since further evidence is needed to assess whether features for these concepts should be included in the UniMorph Schema. Finally, we have taken the liberty of including 'until' in relation to ablative by encoding it as the case feature 'terminative' (TERM). In these fragments,

[6]Due to the problematic nature of applying a specific-general relationship to the masculine-feminine opposition and to the diversity of principles organizing gender across the world's languages, the applicability of this particular hierarchy may not be expected to extend beyond this language, and hierarchies for gender will not be discussed further.

Reported S→G	Judged S→G	Judged G→S
benefactive → dative	reflexive → anticausative	ability → permissive
benefactive → purpose	reflexive → reciprocal	ability → possibility
instrument → manner		ablative → material
iterative → habitual		ablative → partitive
		allative → until
		comitative → temporal
		conditional → concessive
		dative → comparative
		dative → patient
		perfect → perfective

Table 1: Grammaticalization processes found in which the degree of generality differed between features in the same category. Each process is given in the format *source* → *target* to indicate diachronic directionality, following the data source (Heine & Kuteva 2002).

the label for the morphological category has been included to unify disparate hierarchical fragments of features that belong to that category.[7] The theoretical need for a root node and a determination of its content in logical hierarchies has yet to be explored.[8]

(3) a.

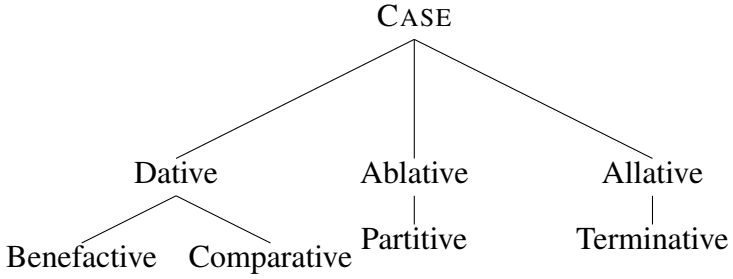

b.

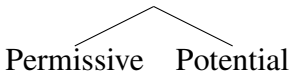

c.

d.

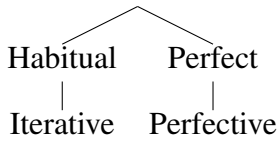

3.2.3 Cross-lingual morphological contrast mismatches

Another source of evidence for cases of the kinds of specific-general relationships that underlie logical hierarchies is the translational mappings between languages whose morphological contrasts are mismatched. For example, dual number in any

[7] The fragments in (3b) and (3c) are from the mood and valency categories, respectively.

[8] The ability to effectively do away with adjective agreement in Jingulu, as described in §3.2.1, can be seen as evidence for the ability to specify a default interpretation for a non-terminal node (e.g. masculine as a default for gender), which is in accordance with findings from phonological feature geometries productively carried over into the morphosyntactic feature geometry of Harley & Ritter (2002).

language is translated into English simply as plural. These kinds of mismatches are similar to language-internal facultative feature use and superclassing in that a more detailed set of distinctions is mapped to a less detailed set. The difference is that these two levels represent different languages.

While cross-lingual evidence is not often used to justify a theoretical construct, it represents a type of linguistic competence that bi- or multi-lingual speakers possess. Kornai (2008:140) notes that evidence from translation "has largely fallen into disfavor" due to the fact that monolingual speakers never need access to information from translation. While this latter point is certainly true, the evidence from language-internal phenomena in which speakers use mappings from a high-contrast feature set to a lower-contrast feature set indicates that the kind of organization proposed here for logical hierarchies on the basis of cross-lingual evidence is also relevant in language-internal contexts. Moreover, cross-lingual evidence sheds light on the organization of what Kibort (2010:81) terms morphosemantic categories. Features in these categories (which comprise the majority of morphological categories, based on the figure in Kibort 2010, 83) do not transparently interact with the syntax of a language, which means that evidence from syntactic patterns like agreement and government is typically unavailable. Therefore, evidence from cross-lingual contrast mismatches (as well as from historical patterns like grammaticalization processes) is valuable for determining the internal organization of morphological (and particularly morphosemantic) features.

We examined translational mappings within a morphological category between a higher-contrast feature set in one language to a lower-contrast feature set in the other by examining inflected word pairs derived from automatic word alignments. The guiding idea was that among aligned words in parallel text (i.e. two translations of the same content, such as the Bible) from two languages, if one or more features in the higher contrast language mapped to a feature in the lower contrast language, then the feature in the lower contrast language likely dominates the feature(s) in the higher contrast language in a logical hierarchy. Setting up translation pairings in this way creates the type of high-contrast vs. low-contrast levels that exist language-internally in cases of facultative feature use and superclassing.

The parallel text data and the word alignments within its sentences were provided freely on the OPUS website in the form of bilingual dictionaries (Tiedemann 2012).[9] Specifically, the parallel text data came from the OpenSubtitles[10] corpora from 2012 and 2013 (Tiedemann 2012), the European Medicines Agency corpus[11] (Tiedemann 2009), and the EU Bookshop corpus[12] (Skadiņš *et al.* 2014). See Table 2 for full details for each language pair.

Word alignments were pre-compiled for each source by the OPUS website maintainers using GIZA++, an automatic statistical word alignment tool described in detail in Och & Ney (2003).[13] Word alignment can be visualized as in Figure 1, which shows a hypothetical example of how the alignment of word pairs across two languages, Finnish and Turkish, would look with an attested example pair supplied

[9] `http://opus.lingfil.uu.se/`
[10] `http://www.opensubtitles.org/`
[11] `http://www.emea.europa.eu/`
[12] `http://bookshop.europa.eu`
[13] `http://www.statmt.org/moses/giza/GIZA++.html`

with a subset of the morphological features available for each word. These morphological features were obtained from the UniMorph Wiktionary database described in Sylak-Glassman *et al.* (2015b) and Kirov *et al.* (2016).[14]

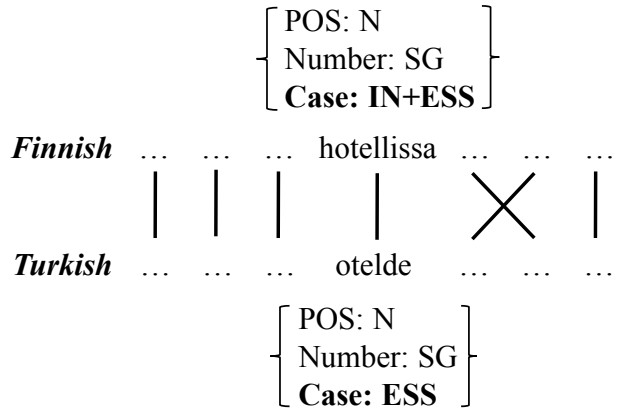

Figure 1: Illustration of possible alignment patterns among of a hypothetical portion of parallel text with an attested aligned pair embedded within. The pair includes the Finnish word *hotellissa* and the Turkish word *otelde*, both meaning "in a hotel." The morphological features in braces were supplied from the UniMorph Wiktionary database.

The measure of word alignment quality adopted by Och & Ney (2003:33) is alignment error rate (AER), which is calculated as in (4) for a test data set that also has manual annotations available.

(4)
$$AER(S, P; A) = 1 - \frac{|A \cap S| + |A \cap P|}{|A| + |S|}$$

In (4), $|A \cap S|$ is the number of proposed alignments (A) which have been judged as sure alignments (S) from manual annotation, and, similarly, $|A \cap P|$ is the number of proposed alignments which have been manually judged as probable alignments (P). These quantities are summed and divided by the sum of all proposed alignments and all manually-annotated sure alignments. This quotient is subtracted from one, yielding a percentage alignment error rate.

Cross-lingual feature contrast mismatches within the categories of case and number were examined within six pairs of languages, shown in Table 2. A lack of sufficient parallel text and morphological feature specifications precluded a broader examination across additional languages that could provide insight into the organization of features within categories likely to exhibit logical hierarchies, such as evidentiality, possession, and voice.

The overall results of our examination of the alignment of number features across these language pairs corresponds to the evidence encountered with facultative feature use and superclassing that dual number is dominated by plural in a

[14]This full database, with approximately ∼977,000 lexemes and ∼14 mllion inflected forms, is available at `unimorph.org`.

Language Pair (Source-Target)	*Data Source*	*Sent-ences*	*Source Tokens*	*Target Tokens*	*Alignment Quality Range (in % AER)*
Italian - Maltese	EMEA	1	12.7	16.3	8.7-12.5%
Italian - Slovene	OS '13	4	30.7	25.7	<8.7%
Russian - Slovene	OS '12	2.3	16.4	15.1	<8.7%
Estonian - Turkish	OS '12	5.4	34.8	31.5	<8.7%
Finnish - Turkish	OS '12	8.5	47.3	51.7	<8.7%
Hungarian - Turkish	OS '12	12.2	77.0	72.9	<8.7%

Table 2: Language pairs with feature contrast mismatches within the categories of number (top three pairs) and case (bottom three pairs). Quantities of sentences and (word) tokens are given in the millions. Alignment quality calculations were not provided for the data sources. The approximate AER range for GIZA++ provided by Och & Ney (2003:37) for the quantity of sentences in the source is given instead in the rightmost column. Abbreviations: EMEA = European Medicines Agency, OS = OpenSubtitles.

logical hierarchy. With respect to case, our finding is that specific local cases, especially in combination with motion designations (i.e. ESS = essive, stationary location; ALL = allative, motion toward; ABL = ablative, motion from), map to cases in which only the type of motion is designated. This result is in accordance with the typological generalization, observed by Radkevich (2010:36) in an extensive cross-linguistic survey of local case systems, that a motion-designating directional component "is always present in the structure of PPs, even in the case of PPs with locational meanings."[15]

Table 3 shows the results of comparing the number features of aligned words in Italian, Maltese, Slovene, and Russian. Italian (Indo-European) contrasts only singular and plural while Maltese (Semitic) and Slovene (Indo-European) contrast singular, dual, and plural. While there are very few instances of the dual number in Maltese due to its marginal status (see Borg & Azzopardi-Alexander 1997 for details), the majority of its appearances do align with plural nouns in Italian (shown in 3a).[16] The Slovene dual, which is used more robustly than that of Maltese (particularly with higher animacy referents; Corbett 2012, 28-29), also aligns to the plural in the majority of cases in which it occurs (3b). With Russian, we note a problematic result: Dual appears to align more commonly with singular than with plural. This is due to the fact that the Russian number *dva*, 'two,' requires the noun that it quantifies to take a morphological form that is identical to the genitive singular. However, this form can be more precisely analyzed as a paucal since it is also required with the numbers *tri*, 'three,' and *četyre*, 'four' (Paperno 2012). In this light, dual likely aligns to either paucal or plural.

Table 4 shows the results of comparing the case features of aligned words in Finnish, Hungarian, Estonian, and Turkish. For each of these Uralic languages, in the majority of occurrences, the case features, which designate both location and motion, map to the more general motion designation marked by Turkish case fea-

[15] We assume that single words with local case marking may constitute PPs.

[16] Labeled morphological data for Maltese was obtained from the Ġabra database (Camilleri 2013) with the original source's features mapped to UniMorph features.

a.

↓Ma/It→	SG	PL	X
SG	1183	187	462
DU	1	6	0
PL	171	757	81
{SG/PL}	1	5	37

b.

↓Sl/It→	SG	PL	X
SG	3010	777	517
DU	655	1081	417
PL	415	2140	376
X	682	853	334

c.

↓Sl/Ru→	SG	PL	X
SG	10550	2767	739
DU	3263 (!)	3027	440
PL	3128	5960	317
X	2805	2712	261

Table 3: Results of comparing number features specified for words in aligned pairs in a) Italian and Maltese, b) Italian and Slovene, and c) Russian and Slovene. 'X' indicates that no feature value was available for a given token due to the lack of a feature specification in the annotation source.

tures. These results allow for the formulation of the fragments of logical hierarchies shown in (5).

(5) a. PL — DU

b. ESS → IN+ ESS, ON+ ESS, AT+ ESS

c. ALL → IN+ ALL, ON+ ALL, AT+ ALL

d. ABL → IN+ ABL, ON+ ABL, AT+ ABL

While the relationships that underlie logical hierarchies lead to the expectation of perfect correspondence between specific and more general morphological features, two main sources of error affect this correspondence. First, automatic alignment necessarily introduces error, but likely only up to the thresholds identified in Table 2. Apart from this, linguistic idiosyncrasies, such as the syncretism of the genitive singular and paucal in Russian and the archaic status of the dual in Maltese, can also cause unexpected deviations in the correspondence between morphological features.

However, gathering evidence from cross-lingual morphological contrast mismatches, explored via word alignments, is an important technique for uncovering relationships among morphosemantic features, which do not participate in agreement or government, offering little language-internal evidence for their organization. Using a systematic, contrastive method to establish cross-linguistic distinctions within a morphological category is an important step in uncovering morphosemantic features themselves, but without external evidence about their organization, relationships among them can only be proposed on the basis of subjective interpretation of their contents.

4 Conclusion

This study has presented a proposal for discovering morphological features using a method that prioritizes attested surface contrasts as evidence for morphological features that reflect the distinctions captured by those contrasts. These features are

a.

↓Fi/Tu→	NOM	GEN	DAT/ALL	ACC	ESS	ABL	X
NOM	1998	36	41	675	16	13	644
ACC	3219	261	84	1425	26	30	1249
GEN	344	238	29	450	6	20	13
PRT	862	31	100	309	4	32	295
IN+ESS	26	8	11	7	303	6	0
ON+ESS	31	21	10	10	93	0	5
IN+ALL	31	16	287	22	14	28	9
ON+ALL	17	4	175	9	0	8	1
IN+ABL	119	28	23	94	24	235	16
ON+ABL	9	4	6	3	5	55	0
X	0	0	0	0	0	0	1483

b.

↓Hu/Tu→	NOM	GEN	DAT/ALL	ACC	ESS	ABL	X
NOM	1444	79	42	173	11	13	585
ACC	515	5	38	466	2	19	69
DAT	15	26	22	10	6	4	4
IN+ESS	10	6	27	0	147	3	5
ON+ESS	9	4	10	10	64	15	1
AT+ESS	2	1	2	2	21	8	0
IN+ALL	3	1	79	0	5	8	1
ON+ALL	22	6	157	20	7	1	1
ALL	6	0	21	4	1	1	1
IN+ABL	1	0	3	11	1	59	1
ON+ABL	14	3	3	10	1	43	1
ABL	5	6	5	7	0	42	0
X	1433	250	43	367	19	25	1470

c.

↓Es/Tu→	NOM	GEN	DAT/ALL	ACC	ESS	ABL	X
NOM	401	8	7	106	1	1	133
ACC	222	32	6	66	0	1	103
GEN	358	120	36	195	5	3	8
PRT	362	24	66	268	0	13	23
IN+ESS	16	4	5	4	46	5	1
ON+ESS	11	2	2	5	17	0	0
ESS	1	0	0	0	0	0	0
IN+ALL	0	0	13	1	0	0	0
ALL	5	1	27	5	1	0	1
IN+ABL	5	4	2	1	2	46	0
ABL	2	0	0	0	1	11	0
X	6	0	8	2	0	0	377

Table 4: Results of comparing number features specified for words in aligned pairs in a) Finnish and Turkish, b) Hungarian and Turkish, and c) Estonian and Turkish.

intermediate in status between universal cross-linguistic categories and comparative concepts, in the terms of Haspelmath (2010), and are lower-level than the type of

cognitively-inspired features proposed by Harley & Ritter (2002) in that they more closely model distinctions captured by language-specific descriptive categories. We argued that in addition to dependency hierarchies that capture typological and acquisitional generalizations, logical hierarchies, which capture intrinsic specificity relationships among the conceptual-semantic content of features, are necessary to explain language-internal phenomena like facultative feature use and superclassing. Logical hierarchies are also necessary to explain how speakers are able to map from a higher-contrast system of morphological distinctions in one language to a lower-contrast system in another. Based partly on analyzing attested patterns in these mappings, logical hierarchies can be deduced, and these shed light on the organization of morphosemantic categories which are not amenable to investigation using traditional morphosyntactic patterns, such as agreement and government.

References

Borg, A., & M. Azzopardi-Alexander. 1997. *Maltese*. London: Routledge.

Camilleri, J. J. 2013. A computational grammar and lexicon for Maltese. Master's thesis, Chalmers University of Technology. Gothenburg, Sweden.

Comrie, B., M. Haspelmath, & B. Bickel, 2008. The Leipzig Glossing Rules: Conventions for interlinear morpheme-by-morpheme glosses. http://www.eva.mpg.de/lingua/resources/glossing-rules.php.

Corbett, G. G. 2000. *Number*. Cambridge: CUP.

Corbett, G. G. 2012. *Features*. Cambridge: CUP.

Harley, H., & E. Ritter. 2002. Person and number in pronouns: A feature-geometric analysis. *Language* 78.482–526.

Haspelmath, M. 2010. Comparative concepts and descriptive categories in crosslinguistic studies. *Language* 86.663–687.

Heine, B., & T. Kuteva. 2002. *World Lexicon of Grammaticalization*. Cambridge: Cambridge University Press.

Jakobson, R., G. Fant, & M. Halle. 1952. *Preliminaries to Speech Analysis*. Cambridge, MA: MIT Press.

Kibort, A. 2010. Towards a typology of grammatical features. In *Features: Perspectives on a Key Notion in Linguistics*, ed. by A. Kibort & G. G. Corbett, 64–106. Oxford: OUP.

Kirov, C., J. Sylak-Glassman, R. Que, & D. Yarowsky. 2016. Very-large scale parsing and normalization of Wiktionary morphological paradigms. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*, 3121–3126, Paris, France. European Language Resources Association (ELRA).

Kornai, A. 2008. *Mathematical Linguistics*. London: Springer.

Och, F. J., & H. Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics* 29.19–51.

Paperno, D. 2012. Quantification in standard Russian. In *Handbook of Quantifiers in Natural Language*, ed. by E. L. Keenan & D. Paperno, 729–780, Dordrecht, Netherlands. Springer.

Radkevich, N. V. *On Location: The Structure of Case and Adpositions*. Storrs, CT: University of Connecticut dissertation.

Skadiņš, R., J. Tiedemann, R. Rozis, & D. Deksne. 2014. Billions of parallel words for free. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC 2014)*, Reykjavik, Iceland. European Language Resources Association (ELRA).

Sylak-Glassman, J., C. Kirov, M. Post, R. Que, & D. Yarowsky. 2015a. A universal feature schema for rich morphological annotation and fine-grained cross-lingual part-of-speech tagging. In *Proceedings of the 4th Workshop on Systems and Frameworks for Computational Morphology*

(SFCM), ed. by C. Mahlow & M. Piotrowski, Communications in Computer and Information Science, 72–93. Berlin: Springer.

Sylak-Glassman, J., C. Kirov, D. Yarowsky, & R. Que. 2015b. A language-independent feature schema for inflectional morphology. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, 674–680, Beijing. Association for Computational Linguistics.

Tiedemann, J. 2009. News from OPUS - a collection of multilingual parallel corpora with tools and interfaces. In *Recent Advances in Natural Language Processing 5*, ed. by N. Nicolov, K. Bontcheva, G. Angelova, & R. Mitkov, 237–248. Amsterdam: John Benjamins.

Tiedemann, J. 2012. Parallel data, tools and interfaces in OPUS. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)*, Istanbul. European Language Resources Association (ELRA).