The algebraic structure of morphosyntactic features

Sebastian Bank
University of Leipzig

Abstract

The most common way of separating homophony from syncretism – which is a basic challenge for any inflectional analysis: to distinguish between accidental and systematic form-identity – is attributing only the latter to a coherent feature combination instantiating a natural class. Features predetermine which form-identities can or cannot be analyzed as *natural-class syncretism*. Hence, they are crucial for the restrictiveness and predictions of morphological grammar. However, most current theoretical frameworks (e.g. Anderson 1992, Corbett and Fraser 1993, Halle and Marantz 1993, Stump 2001) do not make explicit their assumptions regarding the formal status of features. They miss out on state-of-the-art formalisms to introduce feature notations like *Formal Concept Analysis* (FCA, going back to Wille 1982, Ganter and Wille 1999) which provides a formal model of conceptualization in general. In this paper, I will show how FCA provides an all-embracing terminology to reproduce, visualize, and compare feature systems from different morphological frameworks, enables more precise and consistent morphological analyses, and crucially serves to rule out excessively powerful notations where the feature combinatorics are decoupled from the distributional facts they represent.

1. Introduction

One of the core problems of morphological analysis is the question whether different occurrences of the same inflectional marker – e.g., a suffixed form – are just accidentally identical homophones or actually instances of the same syncretic morpheme. Whenever different occurrences of a form stand out by a *consistent set of common morphosyntactic features*, they are usually analyzed as being syncretic:[1] There is only one underlying lexical entry for them – for the native speaker they are indistinguishable. Consider the verbal tense and agreement suffixes in (1).[2] The six occurrences of the *-te* suffix are uniquely identified by the property of being past tense and therefore a canonical example of syncretism originating from an underspecified lexical element, i.e. filling multiple paradigm cells that are both necessarily and sufficiently described by a common feature specification:[3]

(1) **Present and past tense forms of German *spielen* 'to play'**

	SG	PL
1	spiel-**e**	spiel-**(e)n**
2	spiel-**st**	spiel-**t**
3	spiel-**t**	spiel-**(e)n**

PRESENT

	SG	PL
1	spiel-**te**	spiel-**te-n**
2	spiel-**te-st**	spiel-**te-t**
3	spiel-**te**	spiel-**te-n**

PAST

In other words, the six environments with *-te* constitute a *natural class*. Formally, this is represented by the well-formedness of a corresponding feature specification, say [PAST]. For the distinction between natural-class-syncretism and homophony, it is thus essential what features exactly are available – and also which of their combinations: If, for example, the feature '−HEARER' is provided, then the four occurrences of *-n* can be analyzed as being a natural-class syncretism. They can stem from a single lexical item with a meaning using this feature: [−HEARER +PL]. It has, however, been argued empirically that there is no

[1] Note that the converse is not necessarily true: Morphological frameworks often characterize additional types of systematic form-identity by the means of extra formal machinery like impoverishment rules (Halle 1997) or rules rules of referral (Zwicky 1985).

[2] See Section 4 for a discussion of the concrete analysis of these data in Müller (2006).

[3] Although an anonymous reviewer demands a more restricted usage of the term *underspecification* by which *-te* might also be called *fully-specified* as long as it is analyzed to fill an 'only tense slot'.

© Sebastian Bank. *Nordlyd* 41.2: 239-259, special issue on Features, edited by Martin Krämer, Sandra Ronai and Peter Svenonius. University of Tromsø – The Arctic University of Norway. http://septentrio.uit.no/index.php/nordlyd

such 'non-hearer' feature because the kind of person syncretism this would allow – spanning the first and the third person excluding the second person – is actually comparatively rare in the languages of the world (e.g. Zwicky 1985, Harley and Ritter 2002, Harbour 2006). Having a richer set of morphological features incorporating −HEARER fails to explain this rareness while the more restricted inventory allows one to predict it. For the empirical progress of morphological models, it is thus required to keep the feature set inventory and the natural classes that unfold from their combinatorics in line with typological observations on the frequency and rarity of syncretism types in natural languages (Cysouw 2010). As this is in fact a rather biased distribution (apart from some common types, many of the logically possible combinations are rare, cf. e.g. Pertsova 2007, Sauerland and Bobaljik 2013), this modeling demands feature parsimony.[4]

In general, morphological analyses follow these considerations by giving at least a brief description of their inventory of features and their foundation – for example, a semantic or syntactic background. Usually, it is taken for granted that the reader can infer their combination rules by declaring them to be similar to those of ordinary *sets* – referring to them as 'feature sets' or feature specifications.[5] However, there are still two different ways to understand a set-like algebra of features: Either collections of features are just sets of *symbols* which by themselves are meaningless or they are abbreviations for the sets of linguistic *objects* that they refer to. While this may appear to be a subtle difference at first glance, it may have radical consequences on the expressive power of feature notations and what types of form-identity are therefore representable as syncretism. For lack of terminology, I will call the former feature *autonomy* and the latter *extensionalism*. Consider, for example, the feature inventory in (2), as used for Sierra Popoluca by Müller (2006). Although it is almost minimal (3 binary features providing $2 \times 2 \times 2 = 8$ combinations for 7 different objects to describe), there is a logical dependence in the feature inventory as defined: As the paradigm lacks an inclusive-singular cell, $+1$ and $+2$ together entail $+$pl (inclusiveness entails plurality).[6] Therefore, there are two different feature notations for the 1INCL pronominal context, viz. [$+1$ $+2$] and [$+1$ $+2$ $+$pl]. The question of feature autonomy vs. extensionalsm is simply whether these two notations are effectively the same or not. In other words, can grammar differentiate them? Within extensionalism, they are the same if and only if they pick out the same set of paradigm cells (which is true in (2) and gets false once e.g. a [$+1$ $+2$ $+$du] inclusive dual cell is added). With feature autonomy, they are different exactly because they contain different feature symbols, i.e. regardless of what the paradigm looks like and regardless of which cells they eventually pick out.

(2) **Binary feature decomposition for pronominal elements with clusivity**

	1EXCL	1INCL	2	3
SG	$+1$ -2 $-$pl	•	-1 $+2$ $-$pl	-1 -2 $-$pl
PL	$+1$ -2 $+$pl	$+1$ $+2$ $+$pl	-1 $+2$ $+$pl	-1 -2 $+$pl

The difference between the two views increases, as feature systems depart more from their combinatorial minimum.[7] Consider, for example (3), which adds a binary ± 3 feature, such that $+3$ subsumes the third person cells (adding just a shorthand for -1 -2) and -3 allows to refer to all cells involving a speech-act-participant.[8]

[4]Such empirical requirements for positing features preclude stipulating for arbitrary paradigm cells to form a natural class, e.g. subsuming the 3SG and 2PL *-t* occurrences under an ad-hoc feature.

[5] A notable exception are analyzes within frameworks with more emphasis on implementability like GPSG (Gazdar et al. 1985) and HPSG (Pollard and Sag 1994). See Petersen and Kilbury (2005) on the close connection between their feature structure notations and Formal Concept Analysis.

[6]Because the binary values of each feature in (2) are complementary the following are equivalent: $+1$ and $+2$ together entail $-$pl. $+1$ and $-$pl together entail -2. $+2$ and $-$pl together entail -1.

[7]Observe that with binary features, a feature system describing a set of objects whose number does not happen to be a power of two will always be non-minimal, i.e. have notational variants.

[8]Note that irrespective of the labeling ($+3$ = −PARTICIPANT; -3 = +PARTICIPANT), third person inherently refers to *complete absence* of speech act participants. This is why 1INCL contexts *can*, 1EXCL.PL and 2PL *must* involve a non-participant 'although' they are subsumed by a -3 label.

(3) **Extended binary feature decomposition**

	1EXCL	1INCL	2	3
SG	$+1\ -2\ -3\ -$pl	•	$-1\ +2\ -3\ -$pl	$-1\ -2\ +3\ -$pl
PL	$+1\ -2\ -3\ +$pl	$+1\ +2\ -3\ +$pl	$-1\ +2\ -3\ +$pl	$-1\ -2\ +3\ +$pl

Adding more general features creates more natural classes to be captured by the feature system. At the same time it adds logical interdependencies between the features (excluding certain combinatorial possibilities): In (3), additionally, $+1$ entails -3, $-1\ -2$ is equivalent with $+3$, $+1$ is incompatible with $+3$, etc. Due to those relations, there are again different ways to notate one and the same set of pronominal contexts: Different sets of symbols (e.g. $[+1]$ and $[+1\ -3]$) refer to the same set of linguistic objects (the paradigm cells 1EXCL.SG, 1EXCL.PL, 1INCL). While these are interchangeable notational variants in an extensionalistic feature algebra, the notations are actually differentiated in an autonomous model. Morphological grammar can exploit this to treat them differently although they are referentially identical (an example from the literature will be discussed later).

So whenever the feature inventory provides different ways to refer to the same set of objects, the two kinds of feature algebra will be different. As demonstrated in (4), they give contradicting relations between non-maximally specified feature sets (equivalence, set containment) and yield different results for the operations that combine them (intersection, union):[9]

(4) **Disagreement between two possible kinds of feature algebra**

	feature autonomy	*feature extensionalism*
a.	$[+3] \neq [-1-2]$	$[+3] = [-1-2]$
b.	$[-1] \not\sqsubseteq [+3]$	$[-1] \sqsubseteq [+3]$
c.	$[+1] \sqcap [+3] = \emptyset$	$[+1] \sqcap [+3] = [-2]$
d.	$[+1] \sqcup [-1] \neq [+2] \sqcup [-2]$	$[+1] \sqcup [-1] = [+2] \sqcup [-2]$

As soon as there are dependencies inside the feature inventory – one feature value being logically dependent on the presence or absence of others –, its autonomous feature algebra will contain extra notations that boost its expressive power in comparison to the more restricted extensionalistic one. As I will argue in this paper, this extra power undermines the restrictiveness of morphological grammar and also challenges its learnability: Feature autonomy introduces analytical ambiguity for cases usually analyzed as involving extended exponence or impoverishment, resulting in less specific empirical predictions. As long as the predictions that follow from the more restrictive feature model have not been falsified, there is no reason to use the less restrictive autonomous model. Furthermore, the learnability of choices between feature specifications that are referentially identical but only discernible from their grammatical side-effects are at least questionable.

This paper is structured as follows: The next section will briefly look over the most widely used feature notations used in morphological grammar and discuss their significance for representing patterns of natural-class syncretism. The third section will give an informal introduction into Formal Concept Analysis (FCA, Wille 1982, Ganter and Wille 1999), which is a fully worked out model of extensionalistic feature algebra that was developed as a practical application of mathematical order and lattice theory (Birkhoff 1940). FCA can be viewed as a formalized description of conceptualization in general – whenever conceptualization means needing to define clear-cut categories of objects in terms of their shared properties. In other words, most of the feature systems used in morphology and phonology can straightforwardly be translated into FCA.[10] Crucially, the model provides the terminology to identify the use of feature autonomy in morphological (or other) analyses. Finally, I will reveal some of the unwanted effects that autonomous feature algebras may have for morphological grammar and argue in favor of their renouncement.

[9] In this paper, I will use square symbols ($\sqsubseteq$, $\sqcap$, $\sqcup$) for feature set operators to distinguish them from operators for the corresponding object sets ($\subseteq$, $\cap$, $\cup$). Note the correspondence between their mirror-images, e.g. $[+1] \sqsubseteq [+1+2]$ is equivalent to $\{$1EXCL.SG, 1EXCL.PL, 1INCL$\} \supseteq \{$1INCL$\}$.

[10] See Petersen (2008) for an application of FCA to phonology.

2. Feature notations in morphological grammar

In morphology, there are two basic flavors of feature notations: *feature-value pairs* as used in Paradigm Function Morphology (PFM, Stump 2001) and Network Morphology (NM, Corbett and Fraser 1993) and *binary* and/or *privative features* as used in A-Morphous Morphology (AM, Anderson 1992) as well as Distributed Morphology (DM, Halle and Marantz 1993). [11] I will refer to them as morphosyntactic feature specifications. In inflectional analysis such feature specifications fulfill two functions: Firstly, they provide a formal representation for the meaning of each individual paradigm cell. Secondly, they define which sets of paradigm cells correspond to more general meanings, i.e., are natural classes.

In the following, I will give a short usage-example of each of the feature notations, following matching examples in each framework's literature as close as possible. Each example will build a feature system for the seven-cell pronominal paradigm from the introduction pointing out potential drawbacks – notably regarding natural classes and the issue of notational variants. At the end of Section 3, all of them will be reimplemented in FCA making their differences explicit in a precise way.

2.1. Feature-value pairs

PFM-style feature-value notations are created by partitioning all available property descriptions into mutually exclusive groups and assigning these groups ('features' consisting of possible 'values') a category name (cf. Stump 2001:39, 88):

(5) **Morphosyntactic features with permissible values**

a.	PER	1, 2, 3	c.	NUM	sg, pl
b.	INCL	yes, no	d.	GEN	masc, fem, neut

Individual feature specifications are sets of 'CATEGORY:value' pairs drawn from such an inventory, for example:

(6) **Morphosyntactic feature specifications as feature-value pairs**

a.	{PER:1, NUM:sg}	c.	{PER:3, NUM:sg, GEN:neut}
b.	{NUM:pl}	d.	{}

Each category needs to be paired with no more than a single value (specifications are partial functions from categories to values) and a value cannot be paired with a category it does not belong to. Hence the following notations are undefined:

(7) **Ill-formed feature specifications**

a.	*{PER:1, PER:2}	(not a partial function)
b.	*{NUM:masc}	(category-value mismatch)

In other words, by assumption all values from a category are always incompatible with each other (they are nominal scales) and all values from different categories are compatible with each other by default (they are logically independent/orthogonal).

In an ideally symmetrical world, these two assumptions would always hold and all paradigms could be represented as *n*-ary rectangles. In reality, however, features of different kinds occasionally interact: For example, gender or obviation may only be combinable with third person – or first person inclusive incompatible with singular number. In these cases the default assumption of logical independence is too lax. It can be overridden by adding specific well-formedness conditions:

[11] For an FCA view on GPSG and HPSG feature notations, see Petersen and Kilbury (2005).

(8) **Feature-value cooccurrence restrictions for non-orthogonal categories**

a. $\sigma \sqsupseteq \{\text{PER:1}\} \vee \sigma \sqsupseteq \{\text{PER:2}\} \longrightarrow \sigma \not\sqsupseteq \{\text{GEN:}\alpha\}$

where $\alpha \in \{\text{masc}, \text{gen}, \text{neut}\}$

b. $\sigma \sqsupseteq \{\text{PER:1, INCL:yes}\} \longrightarrow \sigma \sqsupseteq \{\text{NUM:pl}\}$

The statement in (8a) bans specifications containing first or second person features also having a gender value. (8b) can be thought of as a rewrite rule adding the entailed plural value to specifications containing first person inclusive. While such rule systems can become difficult to create and verify (cf. the complex rule set for Potawatomi in Stump 2001:88f), they eliminate notational variants (provided all interdependencies are covered) and serve to override default orthogonality.

In cases, however, where reality does not seem to comply with the strong assumption of mutual exclusivity, definition gets cumbersome: Feature-value-pairs ban a person notation where first and second person can simply be combined to represent first person inclusive (7a).[12] Instead, clusivity needs to be represented by another feature category. Stump (2001) chooses the yes/no feature INCL (e.g. for Potawatomi), which of course needs to be interacting heavily with the three 'primary' person features PER and the number features. So this requires adding more cooccurrence restrictions (undermining again the orthogonality ideal):

(9) **Additional cooccurrence restrictions to implement clusivity**

a. $\sigma \sqsupseteq \{\text{PER:2}\} \longrightarrow \sigma \sqsupseteq \{\text{INCL:yes}\}$

b. $\sigma \sqsupseteq \{\text{PER:1, NUM:sg}\} \vee \sigma \sqsupseteq \{\text{PER:3}\} \longrightarrow \sigma \sqsupseteq \{\text{INCL:no}\}$

Counting disjunction as rule duplication, the PFM-notation needs as much as four rules to state that the feature representation for the basic paradigm cells are as given in (10).[13]

(10) **Feature specifications for pronominal elements with clusivity**

	1EXCL	1INCL	2	3
SG	PER:1 INCL:no NUM:sg	•	PER:2 INCL:yes NUM:sg	PER:3 INCL:no NUM:sg
PL	PER:1 INCL:no NUM:pl	PER:1 INCL:yes NUM:pl	PER:2 INCL:yes NUM:pl	PER:3 INCL:no NUM:pl

This makes it somewhat less straightforward to add more general features to represent paradigm cells that comprise natural classes: For example, a feature that picks out cells involving speech act participants cannot be defined as a person feature because of mutual exclusivity. It needs to be defined a feature of its own category, which due to the default orthogonality demands more cooccurrence rules to represent the dependencies. While this is perfectly possible technically, the rule set will get more complex with every feature and the category division less meaningful.[14]

2.2. *Ordered attribute paths*

In Network Morphology the meanings of the individual paradigm cells and markers are represented by DATR path expressions (Evans and Gazdar 1996), which consist of feature-value pairs just as defined above (mutually exclusive values of orthogonal features). Instead of sets, the path expressions are *sequences*: By assumption, they impose the additional constraint on feature values to comply with a predefined total ordering of morphosyntactic categories.[15] Feature systems are created by dividing all available attribute

[12] Following Gazdar et al. (1985), Stump (2001) also allows for *set-valued* as opposed to the *atom-valued* features used for person and number: With person values as a set of features, first person inclusive can in principle be modeled within the PFM-notation as {PER:{1:yes, 2:yes}}. However, Stump (2001) does not seem to consider set-valued features for categories like person and number.

[13] One may also object that the second person value really refers to a second person exclusive distribution while the inclusive feature bears the second person distribution.

[14] The natural endpoint of such decomposition is a system where each value is a feature of its own (with just one value) or features have two complementary values, viz. privative/binary features.

[15] With KATR (Finkel et al. 2002), there is also a variant of DATR in which the path expressions are unordered so that the ordering-assumption is not needed technically. Brown and Hippisley (2012:57–64) however clarify that the ordering of features is to be regarded

outcomes into an *ordered partition* of mutually exclusive values (cf. Brown and Hippisley 2012:182):

(11) **Template for PER ≺ NUM ≺ GEN feature-value path expressions**

$$\left\langle \begin{array}{lll} \texttt{1st} & & \\ \texttt{1st\_excl} & \texttt{sg} & \texttt{masc} \\ \texttt{1st\_incl} & \texttt{pl} & \texttt{fem} \\ \texttt{2nd} & & \texttt{neut} \\ \texttt{3rd} & & \end{array} \right\rangle$$

Path expressions are constructed by following the template strictly from left to right with the option of omitting values only at the end of the expression, for example:

(12) **Morphosyntactic feature specifications as path expressions**

a. <1st sg> c. <3rd sg neut>
b. <1st> d. <>

In other words, paths have to be initial substrings of the template. Hence the following paths are undefined given the order in (11):

(13) **Ill-formed path expressions**

a. *<pl> b. *<masc> c. *<3rd masc> d. *<1st 2nd>

Each paradigm cell can be represented by a fully specified path expression:

(14) **Feature specifications for pronominal elements with clusivity**

	1EXCL	1INCL	2	3
SG	1st sg	•	2nd sg	3rd sg
PL	1st_excl pl	1st_incl pl	2nd pl	3rd pl

For the lack of a device similar to cooccurrence restrictions, the resulting feature systems contain notational variants in case of feature interdependencies, e.g. <1st_incl> and <1st_incl pl>. These may be exploited to treat them differently.

When it comes to the use of more general features to represent natural classes, this notation is very restricted. As exemplified by (13a), it is not possible to refer to all plural cells as a natural class as long as the person features are ordered before number. To refer to all three first person cells, one needs to add a slot for that particular feature ordered to the left of the more specific person features. However, if one also wants to refer to the cells involving the hearer these two new features – not being mutually exclusive – would need to be ordered with regard to each other resulting in only one of them being usable in path expressions to refer to the corresponding cells. Due to these limitations, natural classes are sometimes expressed with the help of DATR variables over feature values – similar to α-notation:

(15) **Path expression to denote all plural cells using variables**

```
#vars $person:  1st 1st_excl 1st_incl 2nd 3rd
<$person pl>
```

Note that the inclusion of such a variable notation in principle makes the notation completely unrestrictive – it makes it possible to express each and every combination of paradigm cells. In other words, if variable notation is not restricted somehow, path expressions no longer serve to distinguishes natural from unnatural classes.

an inherent restriction of the framework rather than a limitation of the notation.

2.3. *Binary and privative feature sets*

In A-Morphous Morphology and Distributed Morphology feature systems are usually defined by decomposing the attribute outcomes for each grammatical category into a combination of more general binary or privative features, e.g.:

(16) **Decompositions of person, number, and gender values**

1EXCL = +1 −2	SG = −pl	MASC = *masc*	**or**	MASC = +m −f
1INCL = +1 +2	PL = +pl	FEM = *fem*		FEM = −m +f
2 = −1 +2		NEUT = *neut*		NEUT = −m −f
3 = −1 −2				

Like feature-value pairs, binary features employ a distinction between 'features', for example '1' (speaker) and '2' (hearer) and their two mutually exclusive 'values', viz. '+' and '−'. Privative features on the other hand fulfill the role of the feature and its singleton value at the same time and therefore suspend the distinction. While each binary features introduces a three-way contrast (+, −, and unvalued), privative features only distinguish between their presence and absence.

Which combinations can be regarded as natural classes depends on the exact decomposition. For example if three genders are decomposed into the privative features *masc, fem, neut* from (16), the masculine, feminine and neuter cells can be referred to directly but no gender combination is regarded a potential natural class. If, on the other hand, they are decomposed into the binary features ±m, ±f from (16), {MASC, NEUT} and {FEM, NEUT} are possible natural class gender groupings (−f and −m).

As long as the decomposed attribute outcomes are not freely combinable, it is usually necessary to give the exact feature combination for each of the possible combinations, i.e. each paradigm cell.[16] Consider for example (17) as used e.g. in Müller (2006), repeated from (2):

(17) **Feature specifications for pronominal elements with clusivity**

	1EXCL	1INCL	2	3
SG	+1 −2 −pl	•	−1 +2 −pl	−1 −2 −pl
PL	+1 −2 +pl	+1 +2 +pl	−1 +2 +pl	−1 −2 +pl

Enumerating the fully specified feature notation for each paradigm cell allows one to infer some implicit well-formedness conditions.[17] All feature sets that are no subset of any of the full notations refer to no cell at all and are therefore contradictions:

(18) **Contradicting feature specifications subsuming no cell**

a. [+1 − 1 + 2 − 2] b. [+1 + 2 − pl] c. [+pl − pl]

However, without explicit statement it is not clear if these expressions are generally excluded and if they are equivalent or not: While feature sets with conflicting binary features like (18a) and (18c) can most-likely be assumed invalid, the case is not so obvious with combinations like (18b).[18] Due to the lack of completely spelled-out cooccurrence restrictions it is usually ambiguous, whether notational variants like [+1 + 2] and

[16]Consider a paradigm where obviative is not combinable with first and second person: These cells could be thought of as [+prox], [−obv], or unspecified. In case of dependencies, binary feature pairs may not be complementary, e.g. when [−masc] does not include first and second person, i.e. is restricted to non-masculine third person contexts.

[17]For a more explicit approach, see e.g. the feature-geometries used in Harley and Ritter (2002). Note that under extensionalist assumptions (i.e. notational variants are not differentiated), defining a feature system in terms of an exhaustive set of restrictions (logical dependencies) or by giving the full specification of each paradigm cell is completely equivalent (although the latter might be less error-prone). Crucially, FCA provides algorithms to translate between them.

[18]Note that the validity of a feature specification [+1 +2 −pl] ultimately depends on whether number in the language under consideration is analyzed to be of the minimal/augmented or singular/plural type. Given this, an analysis might link special behavior to the expression in languages where the combination is impossible (in the sense as it is interpreted as not-surfacing).

$[+1+2+\text{pl}]$ are to be treated equivalently (extensionalism) or can have the potential to behave differently in morphological grammar (autonomy).

3. Formal Concept Analysis

In what I called the extensionalistic view, feature notations are abbreviations for sets of objects (pronominal contexts, phonemes, etc.). Accordingly, each feature system provides two things: The definition of a set of well-formed expressions and a procedure to translate these expressions into the set of objects they represent. This has two advantages over directly listing sets of objects: Firstly, feature notation promotes brevity by being able to refer to large (potentially infinite) sets of objects just by writing down their common features. Secondly, feature notation may restrict the sets of objects that can directly be referred to, such that for some object combination there is no corresponding expression.

Consider a domain of three objects. There are $2^3 = 8$ different ways of combining them into (unordered) sets. Set inclusion arranges them into the hierarchy in Figure 1 where an edge indicates that the upper node is a superset of the lower one.

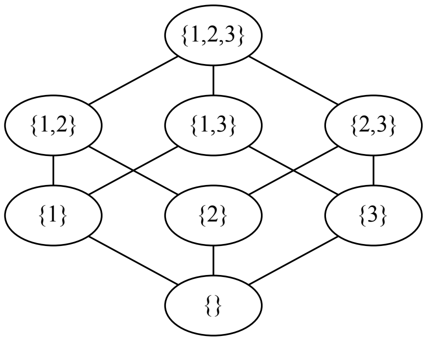

Figure 1: Hasse diagram of the powerset of $\{1,2,3\}$

As this *boolean lattice* already contains all possible combinations, a feature system – in the extensionalistic view – cannot be more fine-grained than this. Feature systems that spare expressions for some of the combinations are restrictive in the sense that not every combination is a natural class. For example, a feature system that lacks a direct notation for the set $\{1,3\}$, while $\{1,2\}$ and $\{2,3\}$ are regarded as natural classes provides the impoverished lattice on the left side of Figure 2. Finally, a feature system that only allows to refer to each individual object or all of them is isomorphic to the lattice on the right.

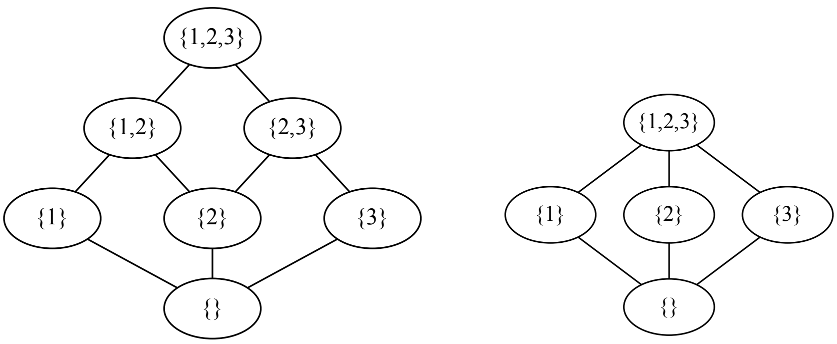

Figure 2: Hasse diagrams of reduced subsets of $\{1,2,3\}$

Formal Concept Analysis (FCA) basically determines which of these kinds of subset hierarchies follow from exactly which feature inventory definition: What kinds of subsets are expressible given an inventory of features and which inventory of features is needed to express a given set of distinctions?[19]

3.1. *Formal contexts, derivation, and formal concepts*

FCA examines the connection between a set of *objects* ($\mathcal{O}$) to be referred to and a set of *attributes* ($\mathcal{A}$) that describe them. The connection is established by a cross-table called *formal context*. It defines the binary relation between both sets ($\mathcal{R} \subseteq \mathcal{O} \times \mathcal{A}$), i.e. which objects has which attribute and vice versa. The context in (19) follows the extended binary feature decomposition of the pronominal paradigm in (3):

(19) **Formal context relating 7 objects with 8 attributes**

	+1	−1	+2	−2	+3	−3	−pl	+pl
1EXCL.SG	×			×		×	×	
1EXCL.PL	×			×		×		×
1INCL.PL	×		×			×		×
2SG		×	×			×	×	
2PL		×	×			×		×
3SG		×		×	×		×	
3PL		×		×	×			×

Attributes can be regarded as boolean flags that each object either has (×) or doesn't – like privative features. Many-valued attributes like binary features or nominal scales need to be translated into flags, which is called *conceptual scaling*. With the resulting table one can do two things: Determine the common attributes O' for a set of objects ($O \subseteq \mathcal{O}$) and the common objects A' for a set of attributes ($A \subseteq \mathcal{A}$). These *derivation* operations are usually notated with the same symbol, the prime ($'$):

(20) **The derivation operator $'$ yielding common attributes/objects**

a. O' of $O \subseteq \mathcal{O} := \{a \in \mathcal{A} \mid \forall o \in O : \langle o,a\rangle \in \mathcal{R}\}$ (common attributes)

b. A' of $A \subseteq \mathcal{A} := \{o \in \mathcal{O} \mid \forall a \in A : \langle o,a\rangle \in \mathcal{R}\}$ (common objects)

With the relation in (19), for example, $\{1\text{EXCL.SG}, 1\text{INCL.PL}\}' = \{+1,-3\}$ and $\{1\text{EXCL.SG}, 2.\text{PL}\}' = \{-3\}$. Furthermore $\{+1,+2\}' = \{1\text{INCL.PL}\}$ and $\{-1,+3\}' = \{3\text{SG}, 3\text{PL}\}$.[20] Observe that the derivation of an empty set $\emptyset'$ yields all attributes $\mathcal{A}$ applying (20a) and all objects $\mathcal{O}$ applying (20b).

Recall that due to the interdependencies of the attributes in (19) there are different ways to derive the same set of objects, e.g. $\{+1,+2\}' = \{1\text{INCL.PL}\}$ and $\{+1,+2,+\text{pl}\}' = \{1\text{INCL.PL}\}$. However, when the derived *extent* is derived again, it yields the maximal set of features that derive it, its *intent* which crucially is unique: $\{1\text{INCL.PL}\}' = \{+1,+2,-3,+\text{pl}\}$ and $\{+1,+2,-3,+\text{pl}\}' = \{1\text{INCL.PL}\}$. Such a pair of extent and intent where the the objects have exactly the features in common and the attributes correspond exactly to the objects is called a *formal concept*:

(21) **Formal concept as pair of *extent* and *intent***

$\langle O,A\rangle$ with $O' = A$ and $A' = O$

[19]Together with this paper I also release `concepts`, an open-source FCA implementation written in Python and `features`, an FCA-based feature set algebra for linguistics. They are available from the Python Package Index (PyPI), see `http://pypi.python.org/pypi/concepts` and `http://pypi.python.org/pypi/features`.

[20]Note that in the selected cases, applying the derivation once more to the result on the right side does not get back the initially derived set on the left-hand side of the equation: $\{+1,-3\}' = \{1\text{EXCL.SG}, 1\text{EXCL.PL}, 1\text{INCL.PL}\}$ and $\{-3\}' = \{1\text{EXCL.SG}, 1\text{EXCL.PL}, 1\text{INCL.PL}, 2\text{SG}, 2\text{PL}\}$. Furthermore $\{1\text{INCL.PL}\}' = \{+1,+2,-3,+\text{pl}\}$ and $\{3\text{SG}, 3\text{PL}\}' = \{-1,-2,+3\}$. However, it is always a superset. Those pairs where double-derivation returns the original value are formal concepts.

The simplest way to get them is to generate all pairs $\{\langle O'', O' \rangle \mid O \subseteq \mathcal{O}\}$ or equivalently $\{\langle A', A'' \rangle \mid A \subseteq \mathcal{A}\}$. In the context table, they correspond to maximal rectangles of crosses that can be formed by reordering the rows and the columns. Hence there is a formal concept in (19) covering the $\{1\textsc{excl.sg}, 2\textsc{sg}\}$ cells with the features $\{-3, -\text{pl}\}$, which form a square if the 2SG row is moved to the top. On the other hand, there is no such concept for $\{1\textsc{excl.sg}, 2\textsc{pl}\}$ as their shared feature set $\{-3\}$ derives $\{1\textsc{excl.sg}, 1\textsc{excl.pl}, 1\textsc{incl.pl}, 2\textsc{sg}, 2\textsc{pl}\}$. There is a concept having exactly the features $\{+1, -2, -3\}$ corresponding to the cells $\{1\textsc{excl.sg}, 1\textsc{excl.pl}\}$. Yet, there is no such concept for $\{+1, -1\}$ because its empty extent $\emptyset$ is also the extent of any other incompatible features, e.g. $\{+1, +3\}$.

3.2. Concept lattices

When all concepts are identified, they form a hierarchy of superconcepts and subconcepts: the *concept lattice* of the formal context. Figure 3 gives the lattice for the concepts in the formal context in (19).

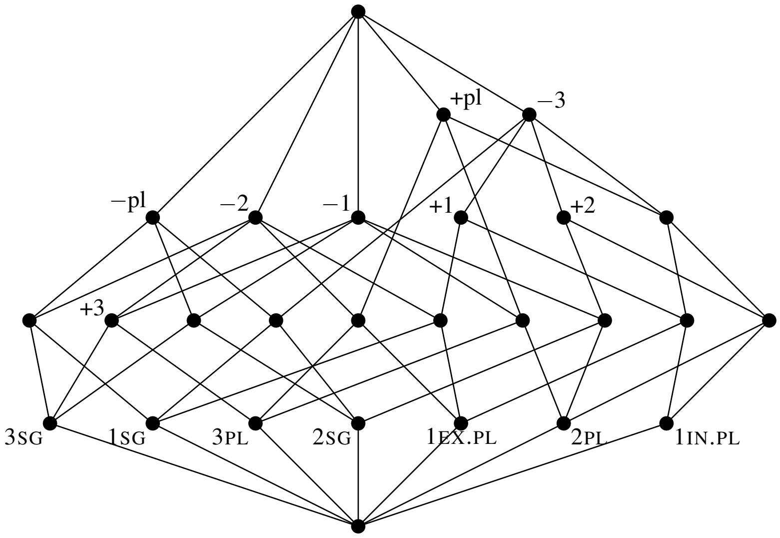

Figure 3: Lattice of 27 concepts for pronominal elements with clusivity

Each node represents a concept subsuming all lower concepts that it directly or indirectly dominates by an edge. The extent of a node is retrieved by following all downward edges collecting the labels below all visited nodes. The intent is retrieved by collecting the labels above nodes following upward edges. For example, the bottom concept, which corresponds to all contradicting attributes, is the $\langle \textit{extent}, \textit{intent} \rangle$ pair in (22a). Going upward from there at the right edge of the lattice visits the increasingly general concepts in (22b-e) up to the most general top concept (22f).

(22) **Implicational hierarchy of formal concepts (a < b < c < d < e < f)**

a. $\langle \emptyset, \{+1, -1, +2, -2, +3, -3, -\text{pl}, +\text{pl}\} \rangle$ *infimum*, $\bot$
b. $\langle \{1\textsc{in.pl}\}, \{+1, +2, -3, +\text{pl}\} \rangle$ *atom*
c. $\langle \{1\textsc{in.pl}, 2\textsc{pl}\}, \{+2, -3, +\text{pl}\} \rangle$
d. $\langle \{1\textsc{ex.pl}, 1\textsc{in.pl}, 2\textsc{pl}\}, \{-3, +\text{pl}\} \rangle$
e. $\langle \{1\textsc{sg}, 2\textsc{sg}, 1\textsc{ex.pl}, 1\textsc{in.pl}, 2\textsc{pl}\}, \{-3\} \rangle$ *coatom*
f. $\langle \{1\textsc{sg}, 2\textsc{sg}, 1\textsc{ex.pl}, 1\textsc{in.pl}, 2\textsc{pl}, 3\textsc{sg}, 3\textsc{pl}\}, \emptyset \rangle$ *supremum*, $\top$

Crucially, the ordering of the lattice is identical to the set inclusion between both the object sets as well as the attribute sets of its nodes:

(23) **Partial ordering $\leq$ between formal concepts**

$\langle O_1, A_1 \rangle \leq \langle O_2, A_2 \rangle$ if and only if $O_1 \subseteq O_2$ or equivalently $A_1 \supseteq A_2$

Therefore the lattice has exactly the same hierarchical structure as the directed graphs from Figures 1 and 2 – meeting the extensionalistic requirements: It is an inclusion hierarchy of possible subsets from a fixed domain of objects. Regardless of the number of attributes it can never have more distinctions than there are subsets in the domain. Each set of objects is paired with an attribute set such that object and attribute set operations can be substituted equivalently for each other. Finally, the derivation operator substitutes notational variants with the same set of objects or attributes which never behave differently. A provisional way to avoid feature autonomy is therefore to always use the most specific notational variant for all feature specifications (applying double derivation $''$ to them)

3.3. Feature algebra with concept lattices

Once the concept lattice is generated, the combinatorics of its members follow from their lattice position. The lattice provides two basic operations: *join* ($\vee$) and *meet* ($\wedge$). The join (*generalization*) yields the closest concept subsuming all joined ones, which is the concept derived from their extent union and intent intersection:

(24) **Join** $[+1 -\text{pl}] \vee [+1 +2 +\text{pl}] = [+1]$
extent: $(\{1\text{SG}\} \cup \{1\text{IN.PL}\})'' = \{1\text{SG}, 1\text{EX.PL}, 1\text{IN.PL}\}$
intent: $(\{+1, -2, -3, -\text{pl}\} \cap \{+1, +2, -3, +\text{pl}\})'' = \{+1, -3\}$

The meet (*unification*) yield the closest concept implying all met ones, which is the concept derived from their extent intersection and intent union:

(25) **Meet** $[+1] \wedge [-\text{pl}] = [+1 -\text{pl}]$
extent: $(\{1\text{SG}, 1\text{EX.PL}, 1\text{IN.PL}\} \cap \{1\text{SG}, 2\text{SG}, 3\text{SG}\})'' = \{1\text{SG}\}$
intent: $(\{+1, -2, -3\} \cup \{-\text{pl}\})'' = \{+1, -2, -3, -\text{pl}\}$

Two concepts are incompatible if their meet is the bottom concept (contradiction): $[+1] \wedge [+3] = \bot$. Two concepts are complementary if moreover their join is the top concept (tautology): $[-\text{pl}] \vee [+\text{pl}] = \top$. If two concepts are compatible and their join is the top concept, they are subcontrary: $[-1] \vee [-3] = \top$. A concept subsumes another one if there is a downward path from the former to the latter, i.e. the former is their join and the latter their meet: $[-3] \vee [+1] = [-3]$. If none of these relations hold between two concepts, they are orthogonal, e.g. $[+1]$ and $[+2]$ in (19) and the corresponding lattice.[21]

3.4. Examples

To illustrate that Formal Concept Analysis is not another feature notation but a general model of describing and analyzing extensionalistic feature notations of all kinds, I will implement the feature notations from Section 2 in FCA.

The following table contains the PFM-style feature-value-pair inventory summarized by the paradigm in (10) translated into a formal context. Crucially, all information is preserved:

[21] Note that these statements presuppose that the bottom concept has an empty extent, i.e. there is no object in the formal context that has every attribute.

(26) **Formal context relating 7 objects with 7 attributes**

	PER:1	PER:2	PER:3	INCL:yes	INCL:no	NUM:sg	NUM:pl
1EXCL.SG	×				×	×	
1EXCL.PL	×				×		×
1INCL.PL	×			×			×
2SG		×		×		×	
2PL		×		×			×
3SG			×		×	×	
3PL			×		×		×

Mutual exclusivity holds between feature-values that have no common object. Feature-value groupings and cooccurrence restrictions do not need to be stated separately. They are simply table patterns, e.g. that PER:2 entails INCL:yes. The same can be read from the corresponding concept lattice, which is given in Figure 4.[22]

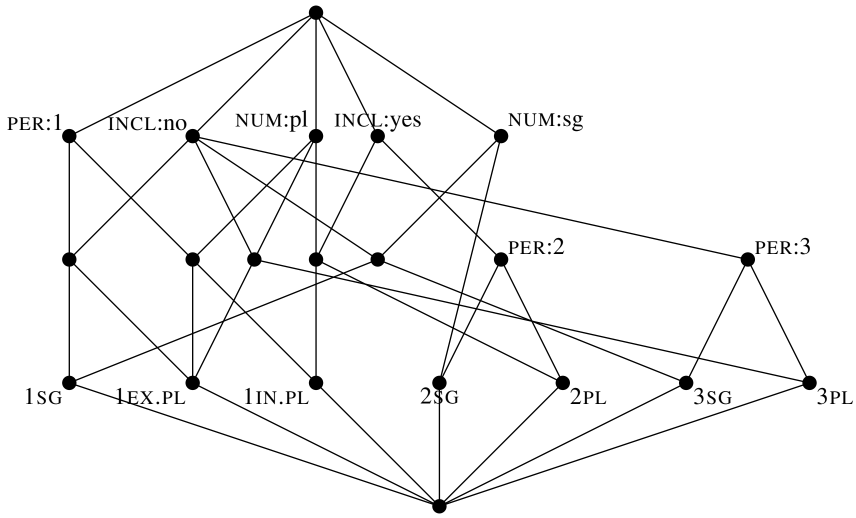

Figure 4: Lattice of 21 concepts for pronominal elements with clusivity

The next table translates the DATR ordered attribute path expressions summarized by (14) into a formal context. As the relation illustrates, it is essentially a nominal scale referring to each paradigm cell plus two more general person features. Observe that the high specificity of the attributes in the table directly follows from the implemented ordered path notation being able to refer to exactly 1) to 10): 1) to 7) each individual cell, 8) second person exclusive, 9) third person, 10) all cells. Together with the emty-set denoting bottom element, this makes exactly 11 concepts.

[22] To facilitate comparison, the node-placement of the graphs in this section was fixed so that the bottom nodes come in a uniform order. They look much smoother when drawn with an algorithm that minimizes crossing lines (but also much more different from each other).

(27) **Formal context relating 7 objects with 9 attributes**

	1st	1st_excl	1st_incl	2nd sg	2nd pl	3rd sg	3rd pl	2nd	3rd
1EXCL.SG	×								
1EXCL.PL		×							
1INCL.PL			×						
2SG				×				×	
2PL					×			×	
3SG						×			×
3PL							×		×

As the three first person cells are already fully identified by their person attribute, the context table can leave out an additional attribute for their number value. Figure 5 contains the corresponding concept lattice.

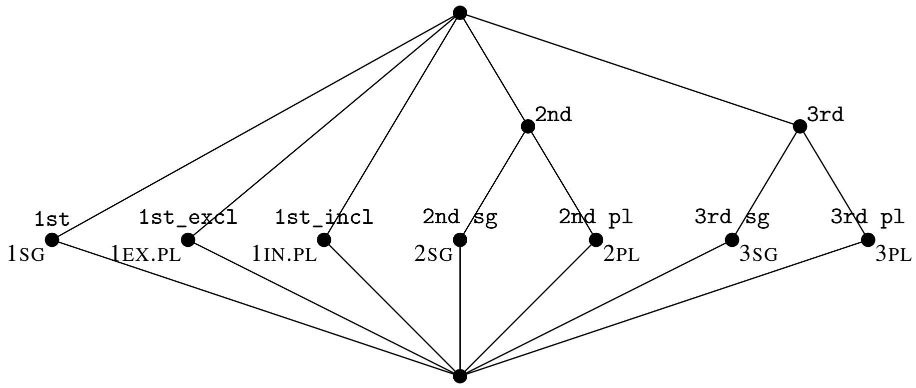

Figure 5: Lattice of 11 concepts for pronominal elements with clusivity

The last formal context implements the feature decomposition from (17), which repeated (2), the initial Sierra Popoluca feature inventory from Müller (2006):

(28) **Formal context relating 7 objects with 6 attributes**

	+1	−1	+2	−2	−pl	+pl
1EXCL.SG	×			×	×	
1EXCL.PL	×			×		×
1INCL.PL	×		×			×
2SG		×	×		×	
2PL		×	×			×
3SG		×		×	×	
3PL		×		×		×

The context already illustrates that the feature system is very similar to the one in (27). Obviously, PER:1 is the same as +1, NUM:sg is −pl, and NUM:pl is +pl. Also, INCL:yes is the same as +2, and INCL:no is −2. The lattices in Figure 4 and Figure 6 allow to confirm this visually and also show that PER:3 is the same as [−1 −2]. In fact, the binary-feature lattice contains all distinctions from the PFM-lattice adding exactly three nodes, i.e. three more possible natural classes: [−1] ({2SG, 2PL, 3SG, 3PL}), [−1 −pl] ({2SG, 3SG}), and [−1 +pl] ({2PL, 3PL}).

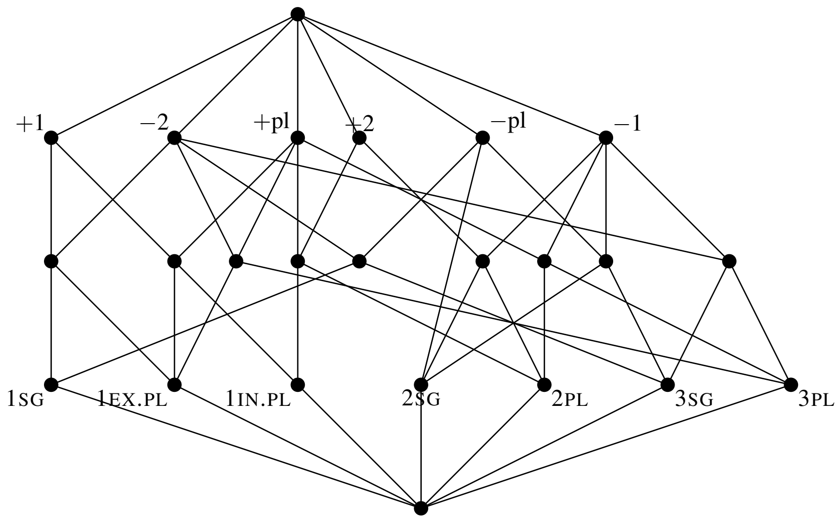

Figure 6: Lattice of 24 concepts for pronominal elements with clusivity

4. Feature algebra in morphological analysis

Morphological analyses distinguish between homophony, where different insertion runs accidentally produce identical forms and syncretism, where different occurrences of a form can be traced back to their common origin (e.g. a single insertion rule, lexical item, etc.). The most straightforward means to have an exponent appear in multiple paradigm cells is to *underspecify* the features that govern its insertion. Consider for example the distribution of the different auxiliary forms in (29):

(29) **Present and past tense forms of English 'to be'**

	SG	PL
1	am	are
2	are	are
3	is	are

PRESENT

	SG	PL
1	was	were
2	were	were
3	was	were

PAST

The feature specifications determining the (non-)insertion of *am* and *is* need to be fully specified: they occur in exactly one paradigm cell. The forms *was*, *are*, and *were* on the other hand occur in multiple cells. They are candidates for two different kinds of syncretism that result from underspecification.

(30) **Insertion feature specifications for syncretic distributions**

a. **fully specified**	b. **natural-class**	c. **elsewhere**
am ↔ [+1 −pl −past]	*was* ↔ [−2 −pl +past]	*are* ↔ [−past]
is ↔ [+3 −pl −past]		*were* ↔ [+past]

For distributions like (30a) and (30b) insertion needs to compare the meaning of the marker with the feature specification of every paradigm cell and insert it if the former subsumes the latter. Elsewhere distributions like (30c) additionally require resolving the competition between multiple compatible markers. To correctly recreate such distributions, where a more specific marker like *was* occurs in a subset of the cells that a more general marker like *were* also matches (1SG/3SG PAST vs. all PAST cells), the more specific marker needs to *block* the more general one in these cells.[23] It is usually assumed that with this Pāṇinian relation between two markers blocking happens automatically – enforced by a blocking principle, subset principle, elsewhere principle, etc., which is part of the grammatical core: As the more specific exponent *was* already indicates past tense, the additional insertion of the blocked exponent *were* into these two past tense cells would fail to add any information anyway.

4.1. Extended exponence

Regarding Pāṇinian blocking as a morphological principle does not mean that it is impossible for languages to violate it. There are in fact cases where exponents occur in a superset of contexts of more specific ones – making them redundant with regards to information-transfer, viz. *extended exponence*. Consider for example the distribution of the *-wa* and *-gi* suffixes in the following paradigm:

(31) **Animate intransitive agreement affixes of Fox** (Bloomfield 1927)

	1EXCL		1INCL		2		3
SG	ne-			•	ke-		-wa
PL	ne-	-pena	ke-	-pena	ke-	-pwa	-wa-gi

According to the blocking principle, the insertion of *-gi* should block the insertion of *-wa* in the 3PL cell: The cells of the former are a proper subset of the extent of the latter (32a). Equivalently, its features are a superset of the other (32b).

(32) **Blocker/blockee relation between the distributions of *-gi* and *-wa***

a. $\{3\text{PL}\} \subset \{3\text{SG}, 3\text{PL}\}$ b. $[+3 +\text{pl}] \sqsupset [+3]$

In other words, the occurrence of a third person marker *-wa* is unexpected as the presence of *-gi* already indicates this information – even within a more specific information. Overriding the blocking principle for such cases of extended exponence requires the use of additional grammatical machinery like feature copying, rule blocks ('slots'), contextual features, marker sensitivity, or enrichment.[24] I will sketch the contextual feature solution: Blocking is assumed to result from discharging a markers' features after it has been inserted – making them unavailable for further insertion (Noyer 1992). In the case of extended exponence, however, a subset of these features are marked as being contextual: Like inherent features, contextual features govern the (non-)insertion of the marker. Unlike them, they are not discharged afterwards and are thus still available for further exponence.

(33) **Contextual feature implementation of extended exponence**

a. *-gi* $\leftrightarrow$ [+pl] / [+3] b. *-wa* $\leftrightarrow$ [+3]

Having the *-gi* suffix in (33a) specify the third person feature as contextual instead of inherent feature overrides its blocking discharge and effects the insertion of the extended third person exponent *-wa*.[25]

Independent of the concrete technical solution, such a division of labor between principled blocking and an overriding formalism for special cases represents the *markedness* of extended exponence: Being a deviation from the expected distribution and requiring additional information in the grammar makes morphologies with this redundant marking more complex and thus harder to acquire. The distinction between blocking and extended exponcence thus contributes to the restrictiveness and empirical predictions of mor-

[23]Under the additional premise that only one insertion can apply ('a single slot'), this is simply a logical requirement: If the more general marker blocked the more specific one, the latter would not be inserted at all.

[24]Note that the concrete implementation and terminology is not relevant here (only the prediction from overriding default blocking having additional analytical costs).

[25]One might challenge the claim that it is the marker failing to block (the primary instead of the extended exponent) that is regarded the special case in this solution.

phological theory.[26] The use of autonomous feature algebra, however, jeopardizes the empirical significance of this distinction because it decouples the notion from observable distributional facts: In the following example, within autonomy, the analysis can be switched between having and lacking extended exponence by the choice of a notational variant for a marker.

Consider the distributions of the suffixes *-s* and *-t* in the analysis cited in (34). Clearly, they are in a Pāṇinian relation: The cells with *-s* are a proper subset of the cells with *-t*: $\{2\text{SG}\} \subset \{2\text{SG}, 2\text{PL}, 3\text{SG}\}$

(34) **Present tense agreement affixes of German** (Müller 2006)[27]

	SG			PL		
						/s/ $\leftrightarrow$ [$+2$ $-$pl]
1	$+1$ -2 $-$pl	*-e*	1	$\cancel{+1}$ -2 $+$pl	*-n*	/n/ $\leftrightarrow$ [-2 $+$pl]
2	-1 $+2$ $-$pl	*-s-t*	2	-1 $+2$ $+$pl	*-t*	/t/ $\leftrightarrow$ [-1]
3	-1 -2 $-$pl	*-t*	3	$\cancel{-1}$ -2 $+$pl	*-n*	/e/ $\leftrightarrow$ [ ]

According to the blocking principle, the more specific marker *-s* should block the insertion of *-t*. This is violated, so the occurrence of *-t* in the 2SG cell is to be regarded as extended exponence. Employing the inherent/contextual features distinction, this gives the following vocabulary items for the two markers:

(35) **Contextual feature implementation of extended exponence**

a. /s/ $\leftrightarrow$ [$-$pl] / [-1 $+2$] b. /t/ $\leftrightarrow$ [-1]

However, if you compare these with the attached list of vocabulary items, the analysis follows a different reasoning that silently requires autonomous feature algebra: To generate the paradigm in (34) with the given vocabulary items, the feature set comparisons by the subset principle needs to treat them as opaque sets of symbols. Only then will the notation [$+2$ $-$pl] fail to be a superset of [-1] and therefore *-s* not block the insertion of *-t* with that meaning. As this contradicts the relation between the corresponding extents $\{2\text{SG}\} \subset \{2\text{SG}, 2\text{PL}, 3\text{SG}, 3\text{PL}\}$, this is only possible with an autonomous feature algebra, where [$+2$ $-$pl] and [-1 $+2$ $-$pl] are actually differentiated although they refer to the same set of paradigm cells.

The issue is not so much that this analysis manages to account for a case that is usually regarded as extended exponence without needing to use extended exponence machinery. The point is that the autonomous feature algebra (being a more powerful superset of the extensionalistic model) is just as compatible with the extended exponence analysis in (35) as with the other one – resulting in analytical ambiguity.[28] It turns the question whether there is extended exponence in the paradigm into whether one chooses to notate [$+2$ $-$pl] or [-1 $+2$ $-$pl] for *-s*, which due to their extensional identity is not decidable on an empirical basis.[29]

I will briefly illustrate that it is not possible to exploit notation to conceal extended exponence in the same way without feature autonomy: In the extensionalistic model, [$+2$ $-$pl] and [-1 $+2$ $-$pl] are equivalent because in (34) $+2$ entails -1. In other words, there is no paradigm cell that has $+2$ but lacks -1. The only way to change that is to introduce such a cell. For example, one could reanalyze German to have an underlying clusivity distinction introducing a $+1$ $+2$ cell:

[26] Apart from the pure markedness of extended exponence, morphological theories occasionally also make predictions about the relative order between primary and extended exponents.

[27] In Müller's (2006) analysis, the crossed out -1 features are the result of the prior application of an impoverishment rule [± 1] $\rightarrow$ $\varnothing$ / [-2 $+$pl] not concerning the current point.

[28] To resolve this ambiguity, one would probably need to replace all grammatical machinery that can be emulated with autonomous features (introducing notational variants as necessary). This does not seem to go without also undermining the natural-class restrictiveness of feature systems.

[29] Admittedly, this does not entail that feature autonomy is wrong but that it makes analysis unnecessarily ambiguous and thus harder to falsify. An extended exponence analysis enforced by extensionalism for example suggests functional pressures to drop the information-wise redundant *-t* from the 2SG cell in (34). Without such expected outcomes, it seems impossible to decide on the empirical adequacy of general notions like the 'markedness of extended exponence' hypothesis.

(36) **Hypothetical clusivity reanalysis of German agreement affixes**

	SG			PL	
1	$+1\,-2\,-\text{pl}$	*-e*	1	$+1\,-2\,+\text{pl}$	*-n*
			12	$+1\,+2\,+\text{pl}$	**-e*
2	$-1\,+2\,-\text{pl}$	*-s-t*	2	$-1\,+2\,+\text{pl}$	*-t*
3	$-1\,-2\,-\text{pl}$	*-t*	3	$-1\,-2\,+\text{pl}$	*-n*

While this makes the spurious prediction of *-e* instead of *-n* in the new cell – requiring further adaptations –, it eliminates the entailment between $+2$ and -1. However, it does not change that *-s* is in a Pāṇinian relation with *-t* and therefore cannot be used to mask their extended exponence. The only way to achieve this is to introduce a paradigm cell with a completely new word form – namely where *-s* occurs without *-t* and the meaning is different from second person singular.[30] Feature extensionalism thus makes it impossible to alter the extended exponence fact without also describing an empirically different language.

4.2. Impoverishment

A second example illustrating how the use of autonomous feature algebra undermines the restrictiveness of morphological formalisms is *impoverishment* (Halle and Marantz 1993, Frampton 2002). In Distributed Morphology, impoverishment rules are deletion operations that alter the morphological input before the insertion of vocabulary items. As they are able to delete features from the input, they have the effect that a marker that would normally be inserted for a specific paradigm cell is bled because some of its insertion-essential features have been impoverished. In that case, it is possible that a different marker subsuming the impoverished features is inserted instead. As impoverishment can only delete features, the alternative marker can never be more specific than the prevented one. Accordingly, it is only possible to reproduce a limited set of syncretic distributions by impoverishment rules, namely those where distinctions are neutralized and more general markers replace more specific markers ('retreat to the general case'). However, this restriction does not hold with autonomous feature algebra.

Consider the abstract example in (37) with two markers *A* and *B* whose distributions are in a Pāṇinian relation such that *A* is regarded less specific than *B*.

(37) **Abstract paradigm with extended and primary exponent**

	SG	PL
3	A	A

PRESENT

	SG	PL
3	AB	AB

PAST

A $\leftrightarrow$ [+3]
B $\leftrightarrow$ [+past] / [+3]

With an extensionalistic feature algebra exactly two changes can be made impoverishing a set of cells where both markers occur. *Firstly*, it is possible to delete features that distinguish the more specific marker from the less specific marker:

(38) **Impoverish [+past] $\rightarrow \varnothing$ / [+3 +pl]**

	SG	PL
3	A	A

PRESENT

	SG	PL
3	AB	A

PAST

Accordingly, the more specific marker is blocked while the less specific marker is still present. *Secondly*, it is possible to delete features that both markers share:

[30] Reanalyzing the number distinction as minimal/augmented, a first person minimal inclusive **-s*.

(39) **Impoverish [+3] → ∅ / [+pl +past]**

	SG	PL
3	A	A
	PRESENT	

	SG	PL
3	AB	
	PAST	

Accordingly, both markers are blocked.

The third logical possibility – preventing the insertion of less specific marker without affecting the insertion of the more specific marker – is regarded impossible.[31] However, impoverishment has the power to generate this distribution if feature specifications are treated as opaque sets of meaningless symbols (autonomy):

(40) **Impoverish [−1 −2] → ∅ / [+3 +pl +past]**

	SG	PL
3	A	A
	PRESENT	

	SG	PL
3	AB	B
	PAST	

A ↔ [−1 −2]
B ↔ [+3 +past]

As long as the feature system provides the notational variation for the markers' features, it is possible to prevent a less specific marker without affecting a more specific one. Feature autonomy thus makes it possible to override 'retreat to the general case' in the sense that 'general case' refers to the observable distributions.[32]

5. Remarks on feature deletion

Morphological models like Distributed Morphology rely on feature deletion for many of their elementary operations (insertion as feature discharge, impoverishment, fission). To make such operations follow extensionalistic feature algebra, it needs to be determined how feature specifications are to be *subtracted* from one another extensionalistically: The use of plain set difference over the maximal sets of feature symbols would bring back the excessive power of feature autonomy, e.g. the deletion of [−1] from [+3 +pl] resulting in a notational variant having +3 but lacking −1. Below, I will sketch a possible solution based on the idea that feature deletion is essentially pre-emptive zero-insertion (Trommer 1999; 2003)

Consider the impoverishment rule [±1] → ∅ / [+pl] applied to the 3PL and 1PL cell of a pronominal paradigm of the same shape as English (three persons, two numbers, no clusivity). Two of these four subtractions are trivial. Trying to remove features from a specification that are not entailed by it (not even in it) does not change anything:

(41) **Non-superconcept subtraction as identity operation**

a. [+3 +pl] − [+1] = [+3 +pl]
b. [+1 +pl] − [−1] = [+1 +pl]

In other words, the set of markers that can be inserted into the cells is not affected by these deletions. Figure 7 gives the lattices of marker meanings that are compatible with each of them (so that markers can be inserted into the cell).

[31] Crucially, this means that patterns like in (40) need to be analyzed differently, e.g. by having a different segmentation with three markers (A, AB, and B) instead of two. In other words, the restriction translates into concrete predictions contributing to the empirical falsifiability of the theory.

[32] Both the extended exponence and the impoverishment issue arise from a more general question regarding the empirical foundation of the notion of *specificity*. In Distributed Morphology the specificity ordering between two markers follows either from the subset relation between their features (Halle and Marantz 1993), or the cardinality of their features (Halle 1997), or a combination of both – possibly enhanced or broken down by a feature hierarchy (Noyer 1992, Müller 2005). What all variations of these specificity principles have in common is that they are *extensions* of the subset relation between the particular paradigm cell sets: If marker *A* matches a proper subset of paradigm cells of marker *B*, then *A* is also more specific than marker *B*, while the converse is not necessarily true. However, this only applies to an extensionalistic feature algebra. With feature autonomy, markers in Pāṇinian relation may be analyzed to be in any specificity relation by using different notational variants. Accordingly, the notion of specificity loses its distributional significance.

If one regards the subtraction of [−1] from the [+3 +pl] lattice on the left as insertion of a [−1] zero-marker, two things follow: *Firstly*, any marker whose meaning is a superconcept of [−1] cannot be inserted any more because that markers' distribution would be in Pāṇinian relation with the zero-marker and therefore blocked. *Secondly*, any marker whose meaning is a subconcept of [−1] (i.e. [+3 +pl], [−1 +pl], and [+3]) cannot be inserted any more because this would contradict the insertion of the zero-marker which would have been blocked by the more specific markers.

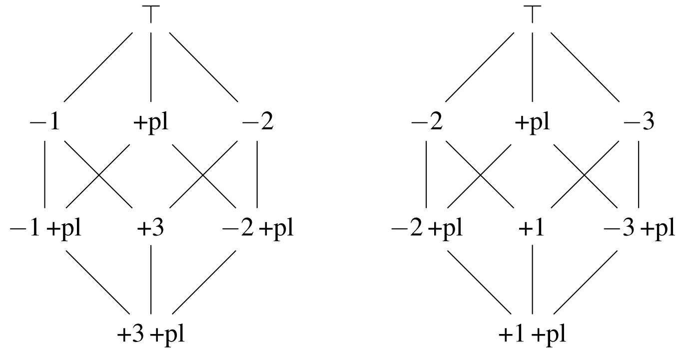

Figure 7: 3PL and 1PL cell specifications with their superconcepts

Therefore, only [−2 +pl], [+pl], and [−2] are possible insertion meanings after the deletion. This set of meanings may be abbreviated as [−2 +pl] as this specification entails exactly the said meanings and will therefore allow their insertions:

(42) **Subtract a superconcept having a complementary concept**
$[+3\,+\mathrm{pl}] - [-1] = [-2\,+\mathrm{pl}]$
$\{\cancel{+3\,+\mathrm{pl}}, \cancel{-1\,+\mathrm{pl}}, \cancel{+3}, -2\,+\mathrm{pl}, \cancel{-1}, +\mathrm{pl}, -2\} = \{-2\,+\mathrm{pl}, +\mathrm{pl}, -2\}$

Applying the same logic to the subtraction of [+1] from the [+1 +pl] lattice on the right, [−2 +pl], [−3 +pl], and [+pl] are possible insertions after the deletion:

(43) **Subtract a superconcept lacking a complementary concept**
$[+1\,+\mathrm{pl}] - [+1] = \{-2\,+\mathrm{pl}, -3\,+\mathrm{pl}, +\mathrm{pl}\}$
$\{\cancel{+1\,+\mathrm{pl}}, -2\,+\mathrm{pl}, \cancel{+1}, -3\,+\mathrm{pl}, \cancel{-2}, +\mathrm{pl}, \cancel{-3}\}$

As the lattice contains no concept entailing exactly these three meanings, they cannot be abbreviated by a single meaning (only a boolean lattice would ensure this).

The following paradigm tables illustrate that these rules for feature set subtraction reflect the distributional facts in the desired way: Removing [−1] from [+3 +pl] results in zero-insertion in the cells {2SG, 2PL, 3SG, 3PL}:

(44) **Distributions corresponding to [+3 +pl] − [−1]**

	SG	PL
1		
2		
3		

+3 +pl

	SG	PL
1		
2	∅	∅
3	∅	∅

−1

	SG	PL
1		
2	∅	∅
3	∅	∅

−2 +pl

	SG	PL
1		
2	∅	∅
3	∅	∅

−2

Afterwards, [−2 +pl] and [−2] markers can still be inserted as they neither are blocked by the zero-marker nor blocking it themselves. Being compatible with [+3 +pl], they are in *overlapping distribution* with the impoverished cells. This gives exactly the kind of markers that match an impoverished cell meaning: subsuming the cell and disjoint or overlapping with the deleted features' distribution.

Removing [+1] from [+1 +pl] results in zero-insertion in the {1SG, 1PL} cells. Accordingly, a marker with the meaning [−2 +pl], which is in overlapping distribution with the zero-marker and therefore still informative, can be inserted:

(45) **Distributions corresponding to [+1 +pl] − [+1]**

	SG	PL
1		
2		
3		

+1 +pl

	SG	PL
1	∅	∅
2		
3		

+1

	SG	PL
1	∅	∅
2		
3		

−2 +pl

	SG	PL
1	*∅	*∅
2		
3		

−2

A [−2] marker, however, is in Pāṇinian relation with the zero marker and therefore blocked in the {1SG, 1PL} cells. Consequently, it is not possible to insert a marker [−2] marker into the 1PL cell after this cell has been impoverished the [+1] feature.

6. Conclusion

As morphological analysis strives to reconstruct the generalizations that native speakers make about the internal structure of words, it is faced with recurring agglomerations: Some combination of word forms make better generalizable clusters than others – they are natural classes that have some inherent connection to be discovered. In inflectional grammar, these clusters are represented by feature specifications: While other combinations of paradigm cells cannot be affected by a single rule or constraint, natural classes may be subsumed under a single feature specification and therefore straightforwardly filled with a syncretic marker. Usually, these morphosyntactic specifications are comprised of semantic or syntactic features, so that they are motivated independently of the inflectional clustering at hand. The features selected by the need to represent the right natural classes will therefore rarely provide a minimalistic notation for the different contrast present in the paradigm. As soon as these features interact logically in one or another way, it is inevitable that inflectional grammar is facing different notational possibilities to refer to one and the same set of paradigm cells – as for example with the clusivity/number interaction. As long as these variants are treated indiscriminately by the grammatical formalism, they only concern stylistic variation and make no empirical difference. If they are, however, differentiated by the grammar – treating them as sets of opaque symbols – formal machinery generally considered restricted may acquire substantial power. As I argued, it is undesirable to make use of this extra power as long as the predictions that follow from a more restricted morphological grammar doing without it are not falsified. If only the most specific notational variant is used to refer to linguistic objects in grammatical formalisms, this effect can be avoided.

References

Anderson, Stephen R. 1992. *A-Morphous Morphology*. Cambridge: Cambridge University Press.

Birkhoff, Garret. 1940. *Lattice Theory*. No. 25 in American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, R.I., first edn.

Bloomfield, Leonard. 1927. Notes on the fox language. *International Journal of American Linguistics* 4 2: 181–219.

Brown, Dunstan and Andrew Hippisley. 2012. *Network morphology : A defaults-based theory of word structure*. Cambridge : Cambridge University Press.

Corbett, Greville G. and Norman M. Fraser. 1993. Network Morphology: a DATR account of Russian nominal inflection. *Journal of Linguistics* 29: 113–142.

Cysouw, Michael. 2010. The expression of person and number: a typologist's perspective. *Morphology* 21 2: 419–443.

Evans, Roger and Gerald Gazdar. 1996. DATR: A language for lexical knowledge representation. *Computational Linguistics* 22 2: 167–216.

Finkel, Raphael, Lei Shen, Gregory Stump, and Suresh Thesayi. 2002. KATR: A set-based extension of DATR. Technical Report 346-02, University of Kentucky.

Frampton, John. 2002. Syncretism, impoverishment, and the structure of person features. In *CLS 38*, Papers from the 2002 Chicago Linguistic Society Meeting, pp. 207–222.

Ganter, Bernhard and Rudolf Wille. 1999. *Formal Concept Analysis: Mathematical Foundations*. Springer, Berlin.

Gazdar, Gerald, Ewan Klein, Geoff Pullum, and Ivan Sag. 1985. *Generalized Phrase Structure Grammar*. Blackwell, Oxford.

Halle, Morris. 1997. Distributed Morphology: Impoverishment and fission. In *Papers at the Interface*, edited by Yoonjung Kang Benjamin Bruening and Martha McGinnis, vol. 30 of *MIT Working Papers in Linguistics*, pp. 425–449. Cambridge MA: MITWPL.

Halle, Morris and Alec Marantz. 1993. Distributed Morphology and the pieces of inflection. In *The View from Building 20*, edited by Kenneth Hale and S. Jay Keyser, pp. 111–176. Cambridge MA: MIT Press.

Harbour, Daniel. 2006. Person hierarchies and geometries without hierarchies or geometries. Handout.

Harley, Heidi and Elizabeth Ritter. 2002. A feature-geometric analysis of person and number. *Language* 78 3: 482–526.

Müller, Gereon. 2005. Syncretism and iconicity in Icelandic noun declensions: A distributed morphology approach. *Yearbook of Morphology* 2004: 229–271.

Müller, Gereon. 2006. Subanalyse verbaler Flexionsmarker. In *Grammatische Untersuchungen, Analysen und Reflexionen. Festschrift für Gisela Zifonun*, edited by Eva Breindl, Lutz Gunkel, and Bruno Strecker, no. 36 in Studien zur Deutschen Sprache, pp. 183–203. Narr, Tübingen.

Noyer, Robert R. 1992. *Features, Positions and Affixes in Autonomous Morphological Structure*. Ph.D. thesis, MIT.

Pertsova, Katya. 2007. *Learning Form-Meaning Mappings in Presence of Homonymy*. Ph.D. thesis, UCLA.

Petersen, Wiebke. 2008. *Zur Minimalität von Paninis Sivasutras - Eine Untersuchung mit Methoden der Formalen Begriffsanalyse*. Ph.D. thesis, Universität Düsseldorf.

Petersen, Wiebke and James Kilbury. 2005. What feature co-occurrence restrictions have to do with type signatures. In *Proceedings of the Joint Conference on Formal Grammar and Mathematics of Language (FG/MOL-05)*, pp. 125–139.

Pollard, Carl and Ivan A. Sag. 1994. *Head-driven phrase structure grammar*. University of Chicago Press.

Sauerland, Uli and Jonathan Bobaljik. 2013. Syncretism distribution modeling and person paradigms. In *Proceedings of GLOW in Asia 9*.

Stump, Gregory T. 2001. *Inflectional Morphology*. Cambridge: Cambridge University Press.

Trommer, Jochen. 1999. Morphology consuming syntax' resources: Generation and parsing in a minimalist version of Distributed Morphology. In *Proceedings of the ESSLI Workshop on Resource Logics and Minimalist Grammars*, edited by Christian Retoré and Edward Stabler.

Trommer, Jochen. 2003. Feature (non-)insertion in a minimalist approach to spellout. In *Proceedings of CLS 39*, pp. 469–480.

Wille, Rudolf. 1982. Restructuring lattice theory: an approach based on hierarchies of concepts. In *Ordered sets*, edited by Ivan Rival, pp. 445–470. Reidel, Dordrecht–Boston.

Zwicky, Arnold M. 1985. How to describe inflection. In *BLS*, vol. 11, pp. 372–386.