



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/93706/>

Version: Submitted Version

Book Section:

Brown, Dunstan Patrick and Chumakina, Marina (2012) What there is and what there might be. In: Brown, Dunstan, Chumakina, Marina and Corbett, Greville, (eds.) Canonical Morphology and Syntax. Oxford University Press, Oxford.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

1. What there might be and what there is: an introduction to Canonical Typology¹

Dunstan Brown and Marina Chumakina

KEY WORDS: canonical typology, morphology, syntax, correspondence problem

Abstract

This chapter introduces the volume and discusses the issue of data comparability in linguistic typology. Key concepts of Canonical Typology, including the base, criteria and the canonical ideal (canon), are introduced and exemplified. Desirable properties of canonical typologies are outlined, and the framework is discussed in relation to other approaches. Summaries are given of each of the chapters in the volume.

1.1 Introduction

Data comparability is a fundamental problem within typology and it has far reaching implications for our ability to define natural languages. Haspelmath (2007) characterises the problem in the following way:

Which are the right categories for a given language? ... For descriptive linguists, there would have to be a list that contains the pre-established categories that general linguists have figured out in some way. These would not necessarily have to be innate, but they would have to be universal in the sense that a descriptive linguist can be sure that the

¹ With thanks to Greville Corbett for sharing his ideas on the basic underpinnings of Canonical Typology, including the analogy of cardinal vowels and the Venus Precept, to Roger Evans for discussion of the nature of the canonical space, to Patricia Cabredo Hofherr and Enrique Palancar on prototypes and the canonical ideal, and to Patricia Cabredo Hofherr and Matthew Baerman for comments on an earlier draft of this chapter. We are also grateful to participants at the Surrey workshop for useful discussion.

categories needed for describing his/her language are on the list. (Haspelmath 2007: 120-121)

The issue of correspondence between similarly named features in different languages is a substantial part of the data comparability issue. Corbett (2008a) calls this the Correspondence Problem:

The Correspondence problem: as typologists we need to be able to justify treating features and their values as comparable across languages. This is not straightforward, and yet a good deal of typology, including enterprises such as the *World Atlas of Language Structures* (WALS), depends upon it. (Corbett 2008a)

In relation to morphosyntactic features, the correspondence problem can be illustrated by the notion of past tense. Languages can divide the notional space of anteriority in different ways, taking into account the remoteness of the event (the most striking example from Dahl and Velupillai's (2008) chapter in WALS is the Yagua language that distinguishes five degrees of remoteness) and/or the aspectual characteristics of the event, and it is impossible to know by the label what sort of oppositions the tense in a certain language is part of. So when we see in a typological database that Russian and Tigrinya have simple past tense, we should not infer that the feature covers the same space in both languages: in Russian, simple past occurs with perfective and imperfective aspect and is opposed to non-past. In Tigrinya simple past denotes only resultative events, and is opposed to stative past (Dahl, 1985: 115-120). As Dahl puts it, "the distribution of PAST depends largely upon what is left when the other categories have taken their share of the pie" (Dahl, 1985: 117).

The Correspondence problem is not just an issue for our treatment of features cross-linguistically, it also affects comparison within languages, because a given phenomenon may vary in its properties in the same language. We illustrate the issue in relation to the notion ‘agreement’ as applied to a specific language, Russian. In the examples in (1), the controller of agreement, a pronoun, is overtly masculine in (1a), and feminine in (1b).

- (1) a. *on pisal*
3SG.M write.PST[SG.M]
‘He was writing.’
- b. *ona pisal-a*
3SG.F write.PST-SG.F
‘She was writing.’

It might be taken as a defining characteristic of agreement that the controller and target of agreement mark the same feature value. So (1a), and (1b) are instances of agreement, because the value on the verb matches that on the controller pronoun. In contrast, this might lead one to discount (2a) and (2b) as agreement, because of the absence of overt gender on the controller.

- (2) a. *ja pisal*
1SG write.PST[SG.M]
‘I was writing.’ (man talking)

- b. *ja* *pisal-a*
1SG write.PST-SG.F
‘I was writing.’ (woman talking)

The point is that constructions such as (2a) and (2b) share some, but not all, of the properties we associate with the notion ‘agreement’. This is one of the reasons why Corbett (2006) used the canonical approach to define agreement.

Canonical Typology seeks to avoid the tendency to use linguistic terms with vague and shifting definitions by placing emphasis on the criteria used to associate particular linguistic phenomena with cross-linguistic categories. It therefore demands greater detail and rigour in terms of description, because it requires the typologist to be clear about the basis on which a phenomenon might be considered an instance of a particular concept. In other words, Canonical Typology addresses the issue of how cross-linguistic concepts can be accurately related to specific categories in a given language (the relationship between *comparative concepts* and *descriptive concepts* as they are called in Haspelmath (2010a: 363) and references there. Canonical Typology shares the focus on the need for detail, as underlined by Haspelmath (2007: 125), for example. However, it places greater emphasis on the use of consistent criteria and multiple dimensions as a way of addressing the Correspondence Problem. We now turn to the description of the Canonical Typology method and its basic conceptual apparatus.

1.2 Doing Canonical Typology

Key concepts of the canonical method are: i) the *base*; ii) *criteria*; and iii) the *canonical ideal* (or *canon*). The base defines the broad space of particular linguistic phenomenon to be described by the typologist. It is defined in such a way that it will

include a wide variety of instances, some of which may be considered to be quite far from the ideal example of the particular category under investigation. We discuss the base in the next section, §0. The canonical method allows the typologist to account for the set of possible instances in languages, by employing sets of criteria to describe how well they approximate to the ideal instance of the particular category. We explain criteria further in §0. As we explain in §0, taken together, the dimensions define a space for which there is one point which is the canonical ideal (or canon). We now turn to the starting point of a canonical typology, the base.

1.2.1 The base

Canonical typologies plot a multidimensional space in which particular linguistic objects differ in terms of their proximity to a point of convergence which is the canonical ideal. In practical terms the question naturally arises how that space is defined more broadly, even before we consider how it is to be divided up. There are two possibilities here. Either the theoretical space, the base, is defined by a long-standing debate about what counts as an instance of the phenomenon in question, or a notional starting point (a *base* or first approximation) needs to be defined. Bond (this volume) illustrates how one can create a notional starting point for the investigation of negation. He identifies 5 key properties which a base should have. Paraphrased, these are:

- it should be broad in coverage so as to be inclusive
- it should be minimal description of the phenomenological domain
- it should contain sufficient information to determine whether the phenomenon exists in the languages under investigation

- it should be supported by empirical evidence from other linguistics sub-disciplines

An example base or first approximation for, say, inflectional classes might look something like the following:

INFLECTIONAL CLASS

For a given language an inflectional class is one of a set of classes which cross-cut syntax and which define the forms for words belonging to the same syntactic category.

We can see that this is a relatively minimal definition, which could cover a broad range of phenomena, some corresponding to our intuitions of inflectional class less than others. It expresses something about the relationship between word forms and syntactic categories.

Once the domain under investigation is broadly defined by the base, we use criteria to define the dimensions which cover this domain, and so we turn to these.

1.2.2 Criteria

The Canonical Typology approach to the correspondence problem is to define sets of dimensions for the phenomenon in question in order to calibrate the position of a particular construction in relation to the ideal definition. *Criteria* are used to define these dimensions, and they state which points represent the more canonical, and less canonical, instances along the dimension. This allows for systematic treatment and can be used to discriminate between instances. To illustrate, we return to the Russian

examples in (1) and (2), considering them in terms of two of the criteria defined by Corbett (2006) for agreement. (Here ‘>’ means *more canonical than*.)

C-1: controller present > controller absent (Corbett 2006: 10)

The examples in (1) and (2) are canonical from the perspective of this criterion, as there is a controller present in all of them. This contrasts with C-2.

C-2: controller has overt expression of agreement features > controller has covert expression of agreement features (Corbett 2006: 11)

There is no gender distinction in the controller *ja* ‘I’ in (2a) and (2b), as opposed to examples (1a) and (1b), where the pronouns are overtly masculine or feminine. Consequently the examples in (1a) and (1b) are more canonical than (2a) and (2b) from the perspective of criterion C-2. Canonical Typology addresses the correspondence problem by defining different dimensions and so deals with it by treating correspondence as a matter of *convergence* on an ideal, with some instances further from the ideal than others. For agreement, for example, (2) is further from the ideal for canonical controllers than (1). Criteria can be used to define a set of possibilities which account for the broad typological space covered by the base, and they allow us to determine where a particular construction is located within this space.

Convergence on a canonical ideal indicates that it is possible to define a phenomenon consistently. A particular challenge in defining criteria so that there is convergence on an ideal is determining the number and nature of different dimensions. For instance, in their discussion of the canonical clitic, Spencer and Luís

(this volume) argue that this needs to be understood in terms of the orthogonal dimensions which define affixes and function words respectively. They argue that there is no unique set of properties which defines clitics and excludes other types of element. Instead, a clitic which best approximates to the canonical ideal is one which has the canonical form properties of an affix and the canonical distribution properties of a function word. The strategy adopted in addressing the problem is similar to that of Brown, Chumakina, Corbett, Spencer and Popova (forthcoming) where periphrasis is defined in terms of the dimensions which define syntax and morphology, so that there is no need to make appeal to unique specific properties. Instead, the phenomenon is understood in terms of the criteria which define the syntax and morphology components more broadly.

In their simplest form criteria define dimensions for which there are two possible values: the canonical value and the non-canonical value. For example, Spencer and Luís use the notion of phonological dependence to defining the distributional independence (DwdInd) of a word.

DwdInd: A function word is not phonologically dependent on another word (Spencer and Luís, this volume)

Likewise Corbett's work on morphosyntactic features defines compositionality in terms of predictability.

Compositionality: Given the lexical semantics of a lexical item and a specification of its feature values, the meaning of the whole is fully predictable. (Corbett, this volume)

We can think of the dimensions defined by these two criteria as having two points associated with them. That is, they may or may not hold. In this simple case we have a dimension $\langle 0,1 \rangle$. People are naturally drawn to thinking of dimensions as lines, but we must bear in mind that we are dealing with linguistic objects. In the case of certain dimensions, as we have noted, there may well be no mid-point between the two ends of the dimension, because either the criterion holds, or it does not. In principle, however, the space defined by the criteria may contain dimensions which differ in their nature (i.e. in terms of the number of points in the dimension, and whether the dimension is discrete or gradient).

While we can recognize an instance which is totally canonical (i.e. a linguistic object which is at the canonical end of every dimension), when we speak of one thing being less canonical than something else, this often has to be understood in relation to a particular dimension. As we explain in our discussion of the canonical ideal in the next section, constructions may be equally non-canonical for different reasons, and non-canonicity involves a partial ranking.


1.2.3 The canonical ideal

For a given linguistic phenomenon once we have a base and a set of criteria which describe the space covered by the base, we are in a position to determine the canonical ideal. We can conceive of the canonical ideal as the point of convergence, and the criteria as dimensions of the space covered by the base. In some approaches, the criteria may operationalize particular overarching *principles*. For example, in Corbett (2006: 10-12) the overarching Principle 1, which states that agreement is redundant as opposed to informative, is operationalized by four criteria. In other

words we can visualize the notion that agreement is canonically redundant in terms of four dimensions in the space.

In Figure 1 we use a Boolean lattice to illustrate the notion of the canonical ideal for four dimensions, each defined by a criterion. As it is a Boolean lattice, a criterion either holds or does not hold. We illustrate with the four canonical criteria used by Corbett (2006: 10-12) to define the space of agreement controllers. We have already seen criteria 1 and criteria 2 earlier in our discussion of Russian examples (1) and (2). Criterion 3 requires that the agreement controller be consistent in the values it takes in order to be canonical for this dimension (2006: 11), and criterion 4 states that the part of speech of the controller should be irrelevant in order to be canonical for that dimension (2006: 12). Both (1) and (2) are canonical with respect to criteria 1, 3 and 4, but differ in relation to criterion 2.

Figure 1: Proximity to the canonical ideal



At the top of the lattice all four criteria hold, and this represents the canonical ideal for the four dimensions which deal with the space of agreement controllers in Corbett's typology. Both examples in (1) are canonical with respect to all four criteria, while example (2), as we have seen, has no overt expression of agreement features on its controller and is therefore non-canonical with respect to criterion 2. So it is located at the point on the lattice where criteria 1, 3 and 4 hold, but 2 does not (namely, C1/C3/C4). It is located one level below example (1), and this gives a measurement of its non-canonical nature. It is important to note that there are three other possibilities which would be equally as non-canonical as example (2), but for different reasons, the three points labelled C1/C2/C4, C1/C2/C3 and C2/C3/C4 on the

lattice. While there is only one way to be totally canonical, it is possible for constructions to be equally non-canonical in different ways. But it is also possible for some constructions to be more non-canonical than others. At the level down from where example (2) is located on the lattice we could place (3) at the C3/C4.

(3) *da, pisal*

yes write.PST[SG.M]

‘Yes, I wrote it.’

(In reply to the question, ‘Did you write the letter?’)

Here the subject pronoun is absent, and so this construction is also non-canonical with respect to criterion 1, being canonical only with respect to criteria 3 and 4. This means that it is less canonical than (2), as well as (1).

The canonical space is therefore populated by variants of constructions. As we have seen, the instances in (1), (2) and (3), although belonging to the same language (Russian), are variants of the agreement construction, but appear at different points in the space for the canonical typology of agreement. Informally, we may talk in holistic terms of one language being more or less canonical than another with respect to a particular construction, but what we really mean is that certain constructions of a particular language are more canonical than those of another. Canonical Typology is a method for comparing languages. While allowing us to move beyond the bottleneck of terminology, it also permits us to move towards predictive measures that are not dependent on straightforward hierarchies. For instance, it would be possible to ask whether canonicity in one dimension for a particular language might predict canonicity in another dimension. It also raises the question of the extent to which

particular parts of the lattice will be favoured in comparison with others. In Canonical Typology it is possible for there to be many different dimensions defining a particular phenomenon. This is because it concentrates on outlining the logical possibilities before examining all the instances. When we come to look at the instances, it is more than likely that a much smaller number of dimensions of variation are doing the work. Indeed, the natural thing to do would be to determine whether particular values for certain dimensions predict other ones, and this may well mean that, when it comes to the actual cross-linguistic reality, certain dimensions, as defined by the criteria, are more important than others. Indeed, it is possible to envisage analyses of data where the dimensions created using the canonical method are reduced by statistical analyses to those which do the most predictive work, as with statistical analysis using Classification and Regression Trees (CART), for example (see Baayen 2008: 148-154).

The criteria for a canonical typology should themselves be grounded in cross-linguistic observations about data, so that for a given construction we can always determine if the canonical criterion holds, if it does not, or if it is irrelevant for the particular construction. This provides the framework with an empirical grounding from which it is possible to move to the more abstract analyses required for typology. It has always been a methodological step in typology to consider individually observable instances and determine which logical combinations of them actually occur. The absence of a particular combination of otherwise attested possibilities is important. The canonical ideal may be rare or non-existent. In §0 we contrast the canonical ideal with prototypes, which we would expect to be frequent. The fact that a canonical ideal may be rare or non-existent is informative and tells us something about the phenomenon under investigation. Either it turns out that languages do not

cluster around this point in the typological space (the set of possibilities as defined by the criteria), or the detail suggested by the typology indicates that there is more work to do. Evans (this volume), for example shows that both canonical direct and canonical biperspectival speech appear to be rare. Either it will turn out that there are specific dimensions along which it is more straightforward to represent the shared world required for indirect and biperspectival speech, or we are still missing significant knowledge of this area. Having used the criteria to map out the space of possibilities, theoretical approaches can concentrate on the sets of criteria where languages and structures cluster to explain why this is, and linguists can search deeper to determine whether the clusterings really are meaningful.

Our typical expectation would be that the criteria will converge. On the other hand, the most common or prototypical instance of a phenomenon may exhibit certain inconsistencies with regard to the canonical ideal. This difference is important, because it highlights something about languages. In §2.4, for instance, we argue that Burmeso shows an instance of the canonical inflectional class, because the forms across the classes differ in each cell. There is a contrast here between this rare ideal and the prototypical inflectional class system, where we often expect to encounter sharing of the morphology between the inflectional classes.

1.2.4 Desirable properties of a canonical typology

We have noted that the criteria of canonical typologies should have a basis in observable data, even though the canonical ideals that they define may not exist. By defining the criteria in relation to otherwise observable phenomena, we can ensure an important property of the canonical approach, namely that we will be able to recognize the canonical examples if we were to come across them. We call this the *Recognizability Precept*.

Recognizability Precept

A canonical ideal must be defined in such a way that it can be generally accepted. This is achievable, because the ideal is itself a combination of properties which individually can be observed in languages, so that we can recognize it, if we come across it.

This property of Canonical Typology also ensures that those instances which approximate to the ideal can be calibrated against it. Finally, we should also be wary of taking the most easily visible and most familiar examples of a particular phenomenon as being the best examples. Greville Corbett refers to this as the *Venus Precept*, because the easy observability of the planet Venus does not mean that it is the best example of a planet.

Venus Precept

The most readily available examples and most commonly discussed instances of particular linguistic phenomena may not actually be the best examples.

As an example of the Venus Precept we could consider the notion of inflectional class. The most readily available examples for linguists come from Indo-European languages, and yet these may not actually be the ones which come closest to the canonical ideal. As Corbett (2008b:8) points out, Burmeso, a language of the Mamberamo river area of Western New Guinea (Donohue, 2001; Ross, 2005) actually has a better example of an inflectional class system, because the inventory of forms differs entirely across the classes.

Table 1.1: verbal inflectional classes in Burmeso (Corbett, 2008b; Donohue, 2001: 100, 102)

	assignment	inflectional class 1		inflectional class 2	
		e.g. -ihi- ‘see’		e.g. -akwa- ‘bite’	
		SG	PL	SG	PL
I	male	<i>j-</i>	<i>s-</i>	<i>b-</i>	<i>t-</i>
II	female, animate	<i>g-</i>	<i>s-</i>	<i>n-</i>	<i>t-</i>
III	miscellaneous	<i>g-</i>	<i>j-</i>	<i>n-</i>	<i>b-</i>
IV	mass nouns	<i>j-</i>	<i>j-</i>	<i>b-</i>	<i>b-</i>
V	banana, sago tree	<i>j-</i>	<i>g-</i>	<i>b-</i>	<i>n-</i>
VI	arrows, coconuts	<i>g-</i>	<i>g-</i>	<i>n-</i>	<i>n-</i>

The Venus Precept naturally forces us to look beyond the most obvious examples, and it should remind us that the languages and families often cited as providing the best instances of a particular phenomenon could be blocking our view.

Another desirable property of a canonical typology is that the dimensions should be logically independent of each other. This can often be achieved by increasing the number of points in the dimension. Consider criterion 10 from Corbett’s (2007a:26) canonical typology of suppletion.

Criterion 10 (Corbett's 2007a canonical typology of suppletion)

non-overlapping > non-directional overlapping > directional overlapping

Corbett uses this dimension to describe the sharing of forms between lexemes. The term ‘overlapping’ is taken from Juge (1999:183), where it is used to describe the pattern of sharing which exists between the verbs *ser* ‘to be’ and *ir* ‘to go’ in Spanish, whereby these use the same forms (*fui* etc.) in the preterite. Corbett points out that these are non-directional overlaps in that it is not clear synchronically which lexeme the set of suppletive forms should primarily be associated with. On the other hand, the ablative and genitive forms of Latin *nemo* ‘nobody’ are based on the ablative and genitive forms of *nullus* ‘none’. Corbett terms this ‘directional overlapping suppletion’. What is important to note here is that we have one dimension with three points on it. It would have been possible to define two dimensions: one in which non-overlapping suppletion was more canonical than overlapping suppletion; and one in which non-directional overlapping suppletion was more canonical than directional overlapping suppletion. But this would have introduced additional complexity into the space in that the second dimension would be dependent on the first. This should be avoided, if possible, when defining a canonical typology. We call this the *Precept of Independence*.

Precept of Independence

Define criteria so that the values of one dimension can be determined independently of the values of another.

In principle, we should be able to create canonical typologies where any combination of values along the defined dimensions is possible. Whether the particular combinations are actually to be found is another matter, of course.

An ideal property of linguistic typologies is that they should fit into a large space and be consistent with analyses of other phenomena, and also point to areas in which there is inconsistency or lack of clarity in our knowledge. This is important because the bigger long-term aim of typology is to define a space for linguistic theorizing wherein a particular account covers as large a set of phenomena as possible. In order to do this, however, we need to fit the elements together. A good example of this is Nikolaeva and Spencer's (this volume) chapter where they show how the canonical typologies for modification and possession fit together, as well as highlighting the key areas in which they contrast. (Canonical modifiers are adjectives, for instance, while canonical possessors are nouns.)

Given the overarching aim to combine typologies, this leads us to another desirable property:

Recyclability

Criteria for one typology should be created with their recyclability for others in mind.

From the foregoing it should be apparent that this represents a cautious approach to the creation of typologies in which the categories in which linguists deal can be checked against the reality. We do not start out with the assumption that what linguists imagine there should be actually exists. Equally, we allow for the range of possibilities to emerge.

1.3 Canonical Typology in context

While it represents a new move in terms of the set of methodological steps it employs, Canonical Typology naturally has affinities with certain current approaches, as well

as employing reasoning which is recognized from well established areas. For instance, the vowel space as defined by IPA could be understood in terms of the notion of the canonical ideal. Ladefoged and Maddieson (1996: 289) point out that Jones's (1956) cardinal vowel scheme has four levels of height, and the IPA (1989) set has an implied seven levels. They argue that while there are clearly more than three levels of height, it is doubtful that any language uses this full range. The vowel set therefore constitutes logical points of reference, but this does not mean that it represents what is cross-linguistically common. This makes it like a canonical ideal; it defines points (ideals) which may be approached only rarely in reality. As we have argued, the canonical ideal also contrasts with the notion of prototype. In the earliest work on prototypes, such as Rosch (1973), prototypes are associated with perceptual salience and are therefore natural categories from a cognitive perspective. We can contrast this with the notion of a canonical ideal, where the ideal constitutes a point of convergence for logically consistent definitions.

The tradition of work on the notion of 'category squish', starting with Ross (1972, 1973/2004), is an important part of the intellectual heritage of the canonical method. In this approach syntactic categories, such as noun and verb, are dissolved into a number of processes and the familiar word classes are 'cardinal points' within the continuum of variation that can be found (Sasse 2001: 495). These points are also taken to be prototypes. Sasse (2001: 507) notes in concluding his discussion of squishes that there is a bias within linguistics towards taking it as given that there is necessarily a prototypical relationship between form and meaning. In contrast, adopting the methodological position that the points of convergence which are canonical ideals are not the same as prototypes, this bias is avoided in Canonical Typology. Because it is a typological method, rather than a theory, it allows both for

the possibility that certain of the abstractions in which linguists deal fail to be instantiated by recognizable groupings of properties and for the possibility that others can be.

The use of multiple dimensions is not unique to the canonical method. For instance, their application can be seen in the work of the Cologne UNITYP project, as discussed in Seiler (2001). Seiler (2001: 324, 338) emphasizes UNITYP's 'teleonomic' view of language, namely that it has a goal-driven function to represent cognitive-conceptual content. The starting point in such typologies is that different structural phenomena can be grouped together because they a 'common functional denominator' (Seiler 2001: 323). While particular typologies created using the canonical method will make reference to criteria about semantics or particular functions of a construction, the canonical method does not take this assumption as its starting point, partly because there is evidence that certain systematic structures within languages may not readily fit with the view that structure is always determined by a need to represent cognitive-conceptual content. Inflectional classes, because they create morphological complexity that cross-cuts syntax and semantics, are an example of this. There may well be sound reasons why they exist which have a cognitive explanation, but the representation of cognitive-conceptual content need not be one of them.

Idea that a typological phenomenon needs to be described using multiple dimensions is employed by Givón whose account of grammaticalization processes of pronominal verb agreement uses a set of predictive implicational hierarchies (grammatical relations, semantic role, animacy, humanity, definiteness and topicalization). The situation in individual languages is defined by different point on the hierarchy. The hierarchies are predictive: "If a language would develop obligatory

pronominal agreement, it would develop it first for the highest (left-most) category in each hierarchic scale (Givón 2001: 414).

The difference from the canonical typology perspective is that these hierarchies are defined according to frequency: “these hierarchies must surely be motivated by the text-frequency of anaphoric pronominalization” (Givón 2001:426). In contrast, canonical typology posits an ideal or best example which may be rare.

Among current approaches to typology, there are a number which share Canonical Typology’s use of dimensions. In Hyman’s (2009) property-driven approach to typology the properties are used to place ‘pitch accent’ within a typological space. Hyman argues that ‘pitch accent’ is a concept which has no clear and consistently identifiable status, other than being associated with a variety of properties otherwise associated with stress or tone. If the properties for stress and tone are defined, then pitch-accent systems are positioned at a point far from the best examples of either. Hyman’s properties play a role akin to the criteria in Canonical Typology. A point of contrast is that Hyman (2009: 215) talks of the prototypical being a clustering of the relevant properties which is associated with the “best” or “clearest” category. However, prototype is a notion which has psychological import, and the canonical method works on the basis that it is best to make a distinction between the prototypical and the canonical ideal. As we argued in §0 the canonical ideal is the best example there could be, whereas the prototype may be the best or most accessible instance we can find in a given language. An additional question which arises is whether we should allow for certain dimensions to be privileged over others. For example, in Hyman’s (2009: 216) set of properties the structural property, whereby assignment is determined by metrical structure, is considered definitional for stress. In other phenomena, however, we might assume that no one criterion is

necessary or sufficient. This will become an increasingly important issue as the Canonical Typology framework develops, and we noted in §0 that it is best if criteria are defined so that the dimensions are independent of each other.

The criteria of the canonical method can be compared with the use of cross-linguistic variables in multivariate-based typology, as proposed by Bickel (2010), for instance, in his analysis of clause linkage. Bickel uses multivariate analysis to uncover correlations in a dataset which differs in terms of a number of variables, such as illocutionary scope or constituent extraction, which cover the diversity in clause linkage. Bickel argues that the variables should be based on inductive development (such as in Bickel and Nichol's 2002 work on AUTOTYP), rather than *a priori* assumptions about what a grammar should be like. There is a subtle point of contrast here with Canonical Typology which identifies criteria that have been observed as important for identifying particular phenomena and treats their combination as defining an ideal which may be rarely obtained. In this sense, there is greater emphasis on what could possibly occur, as opposed to what there is. However, an important contribution of Bickel's analysis of clause linkage is that it demonstrates the value of analyzing the variation within a set of related structures to see what is cross-linguistically prevalent.

1.4 Summary

We noted that Canonical Typology is a response to the Correspondence Problem, wherein we cannot be sure that the categories that linguists talk of are actually the same or similar cross-linguistically. We take the definitions of the phenomenon in question to their logical conclusions to define a space into which we can map

linguistic objects from different languages.² Once this has been done we can see how they cluster cross-linguistically. Canonical Typology therefore represents a cautiously optimistic answer to the question of data comparability both within languages and cross-linguistically. We now go on to discuss each of the chapters from the volume.

1.5 Outline of chapters

In the first chapter, ‘A base for canonical negation’, **Oliver Bond** outlines the methodological objectives of canonical typology (CT) and tries them out in the domain of linguistic negation. The first step in any investigation within the canonical approach must start with working out the base, which comes before the elaboration of the criteria and before examining the linguistic data. Bond defines the requirements for the base: it should be minimal in words, and maximally inclusive to avoid being language-specific. At the same time, it should be specific enough for a linguist using it to be able to recognise whether it makes sense to look for the phenomenon in a particular language at all. There is no requirement for the base to contain a functional explanation of the phenomenon, or to give reference to the linguistic form. On the basis of the existing works which use the approach, Bond distinguishes two types of CT: ‘exploratory CT’, where the domain of investigation is defined through the base, and criteria are established on the basis of the observed cross-linguistic variation and ‘retrospective CT’, which apply the method to resolve issues arising from long-standing debates. The chapter represents the first exposition of the methodology for defining the starting point in CT: the base for linguistic negation is defined, and 18 criteria for canonical negation are established. This allows the author to generate a

² It is worth noting Haspelmath's (2007: 128) point in defence of the notion that semantics constitutes substance: "experience shows that people can understand each other across linguistic boundaries with some efforts." Ultimately, a shared understanding between linguists of the approximate meanings of particular linguistic categories is also an example of this general ability to translate concepts.

full description of the canonical instance of linguistic negation without necessarily observing it in reality. The base for linguistic negation defines the domain as maximally inclusive, so that all syntactic levels are involved: constituent, clausal and bi-clausal. Derivational negation is also part of the domain. Another important distinction made in this chapter concerns the fact that linguistic negation appears in expressions that are not treatable in terms of truth-values (such as imperatives) and therefore cannot be defined in the same way as negation in logic.

‘Canonical morphosyntactic features’ by **Greville G. Corbett** tackles the very general problem of features. Criteria for canonical morphosyntactic features are based on two overarching principles: “a feature (I) has robust formal marking and (II) is constrained by simple rules of syntax”. Ten criteria have been proposed, falling under these general principles. However, if we had fully canonical features, meeting all these criteria, these features would be indistinguishable one from another (and an appeal to semantics does not solve our problem). The chapter therefore examines the interaction of canonical morphosyntactic features with canonical parts of speech. The latter are defined as those where semantic, syntactic and morphological behaviours line up perfectly. Four criteria are proposed for the interaction with morphosyntactic features: exclusiveness, exhaustiveness, open/closed class and compositionality. Investigation of possible deviations from the canon (understood as a logical endpoint where all criteria converge), has an important result of ordering the features according to decreasing canonicity, where number comes as the most canonical, followed by gender, person, (respect), case and (definiteness), where brackets indicate less clear instances of features.

The chapter on the typology of quotations by **Nick Evans** argues that postulation of the canonical point should further our understanding of *observable* or

recognisable phenomena. Canonical points which represent something not instantiated in any language are allowed, but only as long as they serve as points of reference for meaningful typological comparison. Evans applies the canonical approach to the grammatical means used by languages to mark reported speech. He gives definitions of canonical direct and canonical indirect speech, conforming to the “canonical base” as defined by Bond, in that they are semantic, maximally inclusive and not language-specific. The chapter discusses various deviations from these criteria. First, it investigates languages where all quotations are basically direct, but there are deviations from this in terms of stylistic effects and “second person magnetism” (usage of the second person for the primary hearer irrespective of what had actually been said in the speech reported). Next, the deviations from canonical indirect speech are presented, in the same order: first stylistic devices (such as expletives), then the grammatical means: the absence of expected tense shift, the retention in the imperatives and vocatives, spatial deixis, etc. One can call all these ‘semi-direct’ or ‘semi-indirect’ speech. The canonical approach is advantageous in that it describes the deviations along different dimensions (style, person, tense, mood, politeness etc). The chapter also suggests that there is another canonical point: biperspectival speech, with dedicated grammatical means to express the perspective of the original and reported speakers simultaneously. There are no canonical instances that we know of, yet logophors represent one point in this theoretical space. In this respect, the chapter presents a clearly described canonical space for something that has not been attested, yet is possible and worth looking for.

The chapter ‘Unpacking finiteness’ by **Irina Nikolaeva** establishes criteria which define the canonical point of finiteness. Finiteness has been viewed by typologists as a “scalar ‘meta’-notion” characterised by interdependent and

hierarchically organized properties. Nikolaeva argues that the implications and notion of scale should be abandoned, since no implicational relations suggested for finiteness parameters so far have been without exceptions. She claims finiteness to be a clausal notion which can be characterised by independent criteria belonging to different linguistic components: morphology, syntax and semantics. The suggested criteria are of the form ‘tense marking > no tense marking’, which reads as ‘overt tense marking of finite clauses is more canonical than no tense marking’. Seven morphological criteria are suggested, involving agreement, marking of tense, mood, politeness and evidentiality, switch-reference and nominative subjects. The syntactic criteria deal with dependency of the clause, subject licensing and morphosyntactic expression of information structure. The semantic criteria concern assertion, independent temporal anchoring and information structuring. Each criterion is illustrated with examples from typologically diverse languages representing more and less finite clauses.

All the chapters discussed so far used the canonical approach to clarify the notions used in typology or to state the usefulness of a notion like finiteness for cross-linguistic comparison. The chapter on clitics by **Andrew Spencer** and **Ana Luís** uses the canonical approach to demonstrate that it can be used to define a linguistic phenomenon which has been claimed to fall between affixes and function words. Clitics, as the paper claims are affixes in form and function words in distribution properties. The authors formulate criteria for canonical words and canonical affixes. The criteria are divided into three groups: those dealing with forms (phonology and broadly understood morphology), those dealing with distribution (syntax) and those dealing with content.

Anna Siewerska and **Dik Bakker** tackle another methodological issue of canonical typology: how to define whether a particular characteristic belongs to the

canonical core. They investigate passive constructions in 279 languages to find out whether the expressability of an agent is canonical for passives, taking the definitions of a prototypical passive as their departure point. The question of the overt agent is the only point where frequency-based and exemplar-based approaches to prototypes disagree. The authors' answer within the canonical approach is twofold: a certain property is taken to be more canonical than another if that property "distinguishes the given construction from canonical realizations of other constructions", and since the overt agent distinguishes passive from impersonal, inverse, stative, anticausative and such like, its presence is taken to be part of a passive canon. The chapter explores the relationship between the expressability of agents and other canonical features of passive (both formal and semantic). The frequency of the overt agent in the passive is taken to be defining: languages which allow (though not necessarily require) overt agents are more frequent (contrary to what the prototype approach states), and therefore passive constructions with implicit agents claimed to be more canonical. The form of the subject is also important: personal passives are more canonical than the impersonal ones, and, as they more often allow the overt agent, it is also a reason to consider passives with overt agents more canonical. Other factors influencing the canonicity of the overt agent include the obligatoriness of its expression, semantic transitivity of the construction, animacy of the agent, its semantic role, and formal realization. The chapter concludes that an expressible non-obligatory semantically agentive lexical agent realized as an oblique argument is canonical.

Martin Everaert suggests the canonical criteria for reflexivization. The languages of the world employ different strategies to encode the identity between two arguments of a clause (a binder and a bindee), the suggested criteria rank them along four dimensions: the expression of the binder, the expression of the bindee, the

morphosyntactic encoding and the domain. The chapter states that reflexivization is “morphosyntactically distinct from a pronominal coreference relation” and it is “important to distinguish the notion ‘anaphor’, or ‘anaphoric dependency’ from reflexive and reflexivization”, and the criteria help to make the necessary distinctions. The second issue that the criteria resolve concerns the multiplicity of strategies: in general, languages have more than one means to encode reflexivization (nominal reflexives, inalienable possession anaphors, null-reflexives and locally bound pronouns), and it is important to be able to recognize the canonical core. The canonical reflexive relation is between two arguments rather than non-arguments, the binder is canonically a definite NP of third person. The bindee is canonically an object and the reflexivity is marked on it (rather than on the predicate). The morphosyntactic encoding of the reflexive construction is canonically different from the encoding of the identity between non-arguments. The domain of the relation is canonically a predicate rather than a sentence, and a simple sentence rather than a complex one.

In their chapter on possession-modification scale **Irina Nikolaeva** and **Andrew Spencer** employ the canonical typology approach to explain the hierarchy established by themselves previously: the strategies for encoding dependants of the noun respect the monotonicity of the following scale: attributive modifier < modification-by-noun < alienable possession < inalienable possession. Defining the criteria for the canonical ideal of the construction described by this scale, the authors seek to get to the origins of the scale (which was just a result of empirical observation). In this respect, the chapter exemplifies a *retrospective* canonical typology as defined by Bond (this volume). The link between modification and possession is easily spotted: many languages encode constructions like *tall girl* and

Mary's daughter in the same or very similar way. The authors establish semantic, formal and syntactic criteria for canonical modification and possession, based on which canonical modification construction is defined as a construction headed by a noun and containing an adjective which denotes some gradable property, and the canonical possession construction is the one denoting inalienable possession and expressed by two nouns. Modification by noun and alienable possession constructions are viewed as deviations from the canon. The chapter shows that languages respect this division into canonical and non-canonical constructions and never use the same strategies to encode, say two opposite canonical points in the same way without encoding the deviations from the canon in exactly the same way. But that would be predicted by the hierarchy above anyway. The interesting result is the uncovering of a very frequent strategy where two canonical points are encoded differently, but in the same way as the deviations close to them: one strategy for the modifier and modification-by-noun and another, different strategy for alienable and inalienable possession, so establishing a canonical space deepens our understanding of the hierarchy.

The chapter concluding the volume, by **Scott Farrar**, applies the canonical approach to linguistic ontologies. Linguistics has collected so much data by now that (inter)operability and formalization are of increasing importance. To get to this point we need the data to be in a standard format, and the terminology used for data annotation to be universal. The ultimate goal is to operate on the level of *e-linguistics* where computers are used as “*primary* means to publish, search and visualise the descriptive data”. Ill-defined data types and the absence of an infrastructure are a barrier to e-linguistics. The chapter suggests that the implementation of the General Ontology for Linguistic Description (GOLD) together with Communities of Practice

Extensions (COPEs) will resolve these issues. The chapter explains how GOLD is structured, what type of entities it contains and what type of predicates to express relations between these entities. GOLD is supposed to contain only expressions of the highest level of abstraction, which should be applicable to any language and are not possible to change once the ontology is built. Examples of such expression are given: (1) “a verb is a part of speech; (2) “a verb can assign case” etc. COPEs are of a lower level of abstraction, they contain the examples where the level of detail can go right down to one language. For example, if there is a very unusual grammatical feature that is only attested in two languages, and even in these languages it is slightly different, then the labels of the feature will contain the name of the language. In this respect, ontologies are very close to the canonical approach. The chapter presents an example of GOLD implementation and concludes by summarising the parallels between canonical typology and principles of building an ontology.