Leaving the stem by itself

Olivier Bonami
Université de Paris, LLF, CNRS
olivier.bonami@u-paris.fr

Sacha Beniamine
Max Planck Institute for Evolutionary Anthropology
sacha_beniamine@eva.mpg.de

Final version of January 18, 2021

Abstract

Stem allomorphy plays a central role in the recent history of morphology, in no small part thanks to a research program initiated by Aronoff (1994). Yet, there is no agreed upon way of deciding whether some bit of form should be considered a proper part of a stem allomorph or an independent exponent. We explore the possibility of just doing away with the notion of stem allomorphy in inflection. We use computational methods to identify within each word a sequence of strings that do not take part in any alternation within that word's paradigm. We then discuss the relationship of such sequences to the classical notion of a stem, and argue that discontinous stems are both conceptually and empirically more satisfactory.

Keywords: stems, allomorphy, exponence, alignment, European Portuguese, English, French

1 Introduction

The notion of stem allomorphy plays a central role in the recent history of morphology, in no small part thanks to a research program initiated by Aronoff (1994) that embraces the systematic existence of organized sets of stem allomorphs as a way of making sense of patterns of morphological behavior. Yet, as Spencer (2012) notes, a fundamental problem for this approach is that there is no agreed upon way of deciding whether some bit of form should be considered a proper part of a stem allomorph (and hence not a morphological unit) or an independent exponent.

In this chapter we explore the possibility of just doing away with the notion of stem allomorphy in the context of inflection. We use computational methods from item-and-pattern morphology to identify within each word a (possibly discontinuous) sequence of strings that are inflectionally inert, in the sense that they do not take part in any morphophonological alternation within that word's paradigm. Such discontinuous sequences can be identified through simple methods and are indisputably devoid of morphosyntactic import of any kind. We argue that these 'discontinuous stems' are a useful addition to the morphologist's toolbox, and question whether classical collections of stem allomorphs still have a useful role to play. Specifically, we define an algorithmic procedure for inferring both discontinuous stems and sets of continuous stem allomorphs from raw paradigms, and assess their usefulness in capturing the implicative structure of the French, English and European Portuguese conjugation systems.

2 The quest for stem allomorphy

One of the basic questions facing morphology is what we call the 'Inflected Word Recognition Problem' (IWRP):

(1) What allows speakers to draw inferences from a word's form to its content?

The IWRP is modeled on Ackerman, Blevins & Malouf's (2009) 'Paradigm Cell Filling Problem' (2). Like the PCFP, it is a question about the structure of the morphological system (what allows the inferences to be drawn) rather than a question about learning and processing (which exact inferences are drawn, and how).

(2) What licenses reliable inferences about the inflected […] surface forms of a lexical item? (Ackerman et al., 2009, 54)

A speaker faces the IWRP each time they are exposed to an inflected word for the first time, a common situation (Bonami & Beniamine 2016). The shape of the solution is obvious: when seeing an unknown wordform *w*, speakers are able to identify phonological aspects of *w* that they have previously encountered associated with some content, and consider the possibility that these convey the same content in this new instance. The question faced by morphologists is, what are the aspects of the phonology of *w*. Importantly, the IWRP and the PCFP raise different issues for morphology. Although a classical, constructive approach to morphology typically addresses both issues by reference to a single segmentation, it may well be that aspects of forms that are crucial for inferring other forms are distinct from those aspects that are crucial for inferring content.

In this chapter we focus on a subproblem of the IWRP, namely identifying that part of a wordform that conveys the word's lexical meaning, as opposed to morphosyntactic and morphosemantic content. In Matthews's (1972) terms, this amounts to separating out a word's (inflectional) 'stem' from the 'exponents' combining with that stem. In the simplest systems, this is an easy task. As Corbett (2007) highlights, canonically, each lexeme is equipped with a unique stem that conveys solely lexical information, while each paradigm cell is characterized by a combination of exponents that conveys only inflectional information; the unique stem is hence the relevant part of the word. This is illustrated in the small slice of the European Portuguese conjugation system presented in Table 1: stems are the longest substrings that are constant across rows, exponents are the longest substrings that are constant across columns.

Table 1. Indicative present of three European Portuguese first conjugation verbs

	1sg	2sg	3sg	1pl	2pl	3pl
ficar	ˈfik*u*	ˈfik*ɐ∫*	ˈfik*ɐ*	fiˈk*ɐmu∫*	fiˈk*ai∫*	ˈfik*ɐ̃ũ*
entrar	ˈẽtr*u*	ˈẽtr*ɐ∫*	ˈẽtr*ɐ*	ẽˈtr*ɐmu∫*	ẽˈtr*ai∫*	ˈẽtr*ɐ̃ũ*
tentar	ˈtẽt*u*	ˈtẽt*ɐ∫*	ˈtẽt*ɐ*	tẽˈt*ɐmu∫*	tẽˈt*ai∫*	ˈtẽt*ɐ̃ũ*

Inflectional systems are rarely that simple though. First, consideration of the other conjugation classes of European Portuguese reveal the existence of substrings that are constant neither across cells nor across lexemes, highlighted by underlining in Table 2. The theoretical status of these theme vowels is a constant source of hesitation for the Romance morphologist: should they be considered to be part of the stem? If so, we have pervasive stem allomorphy, and stems express more than just lexical meaning: for instance, the use of the consonant-final allomorph signals that we are dealing with a 1SG form. Should they be considered to be stem-external exponents? If so, then some material outside of the stem contributes to expressing lexical identity, as theme vowels narrow down the set of candidate solutions for the identification of which lexeme was used. We will not be proposing a solution to this puzzle, but we argue that the puzzlement is caused by the assumption that words can be segmented into a part expressing lexical identity only and a part

Table 2. Indicative present of three European Portuguese fully regular verbs

	1sg	2sg	3sg	1pl	2pl	3pl
ficar	ˈfik*u*	ˈfikɐ*ʃ*	ˈfikɐ	fiˈkɐ*muʃ*	fiˈkai*ʃ*	ˈfikɐ̃ũ
viver	ˈviv*u*	ˈvivə*ʃ*	ˈvivə	viˈvɐ*muʃ*	viˈvɐi*ʃ*	ˈvivɐ̃ĩ
imprimir	ĩpˈɾim*u*	ĩpˈɾimə*ʃ*	ĩpˈɾimə	ĩpɾiˈmi*muʃ*	ĩpɾiˈmi*ʃ*	ĩpˈɾimɐ̃ĩ

expressing inflectional content only. Postulating stem allomorphs amounts to saving exponents from lexical value at the cost of imbuing stems with exponential value. We see no principled reason why this or the opposite solution should be preferable.[1]

A second complication is apparent when looking at the data in Table 3. European Portuguese exhibits stress-conditioned vowel alternations that affect what we call the pre-thematic vowel. While these alternations are almost categorically predictable from regular phonology, they do have

Table 3. Stress-conditioned vowel alternations in European Portuguese verbs

	1sg	2sg	3sg	1pl	2pl	3pl
chegar	ˈʃegu	ˈʃegɐ*ʃ*	ˈʃegɐ	ʃəˈgɐ*muʃ*	ʃəˈgai*ʃ*	ˈʃegɐ̃ũ
começar	kuˈmɛsu	kuˈmɛsɐ*ʃ*	kuˈmɛsɐ	kuməˈsɐ*muʃ*	kuməˈsai*ʃ*	kuˈmɛsɐ̃ũ
libertar	liˈbɛɾtu	liˈbɛɾtɐ*ʃ*	liˈbɛɾtɐ	libəɾˈtɐ*muʃ*	libəɾˈtai*ʃ*	liˈbɛɾtɐ̃ũ
pagar	ˈpagu	ˈpagɐ*ʃ*	ˈpagɐ	pɐˈgɐ*muʃ*	pɐˈgai*ʃ*	ˈpagɐ̃ũ
chamar	ˈʃɐmu	ˈʃɐmɐ*ʃ*	ˈʃɐmɐ	ʃɐˈmɐ*muʃ*	ʃɐˈmai*ʃ*	ˈʃɐmɐ̃ũ
retomar	rəˈtomu	rəˈtomɐ*ʃ*	rəˈtomɐ	rətuˈmɐ*muʃ*	rətuˈmai*ʃ*	rəˈtomɐ̃ũ
jogar	ˈʒɔgu	ˈʒɔgɐ*ʃ*	ˈʒɔgɐ	ʒuˈgɐ*muʃ*	ʒuˈgai*ʃ*	ˈʒɔgɐ̃ũ
mudar	ˈmudu	ˈmudɐ*ʃ*	ˈmudɐ	muˈdɐ*muʃ*	muˈdai*ʃ*	ˈmudɐ̃ũ

an impact on the IWRP. Consider *chegar*: within its paradigm, having a /ə/ in prethematic position is a partial indication of the fact that the form is 1PL or 2PL. Hence, while this /ə/ does contribute to the expression of lexical identity, it also has some exponential import; on the face of it, deciding whether it should be considered as part of the stem (which then entails that *chegar* has multiple stem allomorphs beginning in /ʃeg/ or /ʃəg/) or a separate unit should be just as contentious as deciding whether theme vowels are part of the stem.

Interestingly, discussions of this or similar data with morphologists over the years suggest to us that many would rather treat theme vowels as stem-external but prethematic vowels and similar segments as part of a stem alternant. We submit that such preferences are due to the use of the two heuristic principles in (3).

(3) a. Stem alternants should not be multiplied.

[1] Walther & Sagot (2011); Walther (2013) address the issue by comparing the description length of full implemented descriptions of a system relying on one or the other hypothesis, within a given formal framework for morphological description. This is in essence an operational implementation of the early generative grammar notion of an evaluation metric for alternative theories, and in that sense is a remarkable effort to model the heuristics that morphologists use to decide on a segmentation. However Walther & Sagot's empirical results show that variation in length among descriptions are small enough that the costs of a suboptimal choice are low, both for a speaker and for an analyst.

b. Stems should be continuous substrings of words.

Of course neither of these principles can be absolute: stem suppletion cannot be accomodated without relaxing (3a), and infixation is an immediate violation of (3b). But they still have some role to play in less extreme cases such as those found in the European Portuguese data under examination. More specifically, since theme vowels are no threat to (3b) nothing precludes one from adhering to (3a) and having shorter stems. In the case of prethematic vowels, however, adhering to (3a) would entail postulating discontinuous stems, hence it is more tempting to be guided by (3b) and have longer stems.

This example suggests that these principles play an important role in shaping the segmentations that are typically taken to be pre-theoretical decisions not worthy of detailed discussion, but that ultimately constrain our perception of the nature and extent of typological variation across inflection systems. In the absence of a principled way of arbitrating which of the two principles is more important, we run the risk of theorizing on shaky grounds.

From these observations we conclude with Spencer (2012) that there is no agreed-upon method for identifying which part of an inflected word is a stem, and that the heuristics used by morphologists in that area are neither systematic nor principled enough. This is unsatisfactory, given the prominent place taken by stem allomorphy in morphological theorizing in the wake of Aronoff (1994) — see, among many others, Maiden (1992); Brown (1998); Cameron-Faulkner & Carstairs-McCarthy (2000); Pirrelli & Battista (2000); Stump (2001); Bonami & Boyé (2002); Blevins (2003); Boyé & Cabredo Hofherr (2006); Montermini & Bonami (2013); Stump & Finkel (2013), and the papers collected in Bonami (2012).

One possible reaction to this situation is to forego segmentation entirely, and stop worrying about stems. Recent literature has highlighted how purely word-based methods can efficiently be deployed to address morphological problems, most prominently the PCFP (see among many others Ackerman et al. 2009; Ackerman & Malouf 2013; Blevins 2016; Bonami & Beniamine 2016; Sims & Parker 2016). If stem identification proves problematic, perhaps we should just dismiss the very notion of a stem.

Such a reaction isn't entirely satisfactory, however, at this point in the history of the field. First, since Carstairs-McCarthy's seminal work (Carstairs 1987; Carstairs-McCarthy 1994), stem allomorphy has been taken by many to be governed by constraints distinct from those governing the distribution of exponents. In fact, as Blevins (2016) notes, much of the literature about morphomic stem distributions is implicitly or explicitly dedicated to addressing predictability in paradigms.[2] Hence it is important to establish whether stem allomorphs have a dedicated role to play in addressing the PCFP. Second, we take the PCFP and the IWRP to be two complementary crucial questions that morphological theory should address; and while the PCFP might not require dealing with subword material, such a position is, at first sight, harder to defend for the IWRP, for which subword structure seems crucially relevant.

In the remainder of this chapter we attempt to further our understanding of these issues by exploring the consequences of the principles in (3). We examine two extreme ways of prioritizing the two principles. The first option is to take (3a) as absolute, and tolerate no stem allomorphy. As a consequence, only substrings that occur in all forms of a lexeme can be considered as being part of this stem, and stems will more often than not be discontinuous sequences of substrings of words. This we call the 'unique discontinuous stem hypothesis'. The second option is to take (3b) as absolute, and tolerate no discontinuity. As a consequence, stem allomorphy will be pervasive, as any variation in shape that is encapsulated within stem material has to also be stem material.

[2] See also Stump & Finkel (2013) on exponent-based vs. stem-based implicative relations. Note though that Bonami & Boyé (2014) explicitly argue on the basis of a detailed explicit comparison of stem-based vs. word-based explorations of French conjugation that segmentation decisions for stems actually get in the way of an understanding of implicative structure.

This we call the 'continuous stem sets hypothesis'.

Although they only constitute two extremes of a vast space of possible ways of applying the principles in (3b), these two options have the advantage that they are simple enough to be implemented in full and tested on realistically-sized dataset. In §3, we outline our implementation of the two hypotheses. In §4 we examine how useful the two hypotheses turn out to be in terms of capturing implicative structure. Section 5 draws some general conclusions.

3 Automatic inference of stems

3.1 Alignment

In this section we present an algorithmic method for inferring both single discontinuous stems and sets of continuous stem allomorphs from raw paradigms of surface forms. In order to find stem-like material in forms, we align together all the forms of each paradigm using a heuristic algorithm for multiple alignments. The goal of such an algorithm is illustrated in Table 4. Each cell of the table contains a single phoneme, each row represents a surface form of the paradigm, and each column represents matched material across forms.

Table 4. Aligned surface forms of *libertar*

Paradigm cell	Aligned form									
prs.ind.1sg	l	i	b	ɛ	ɾ	t	-	-	u	-
prs.ind.2sg	l	i	b	ɛ	ɾ	t	ɐ	-	-	ʃ
prs.ind.3sg	l	i	b	ɛ	ɾ	t	ɐ	-	-	-
prs.ind.1pl	l	i	b	ə	ɾ	t	ɐ	m	u	ʃ
prs.ind.2pl	l	i	b	ə	ɾ	t	a	-	i	ʃ
prs.ind.3pl	l	i	b	ɛ	ɾ	t	ɐ̃	-	ũ	-

The optimal pairwise alignment of two forms can be computed easily given a scoring scheme (Needleman & Wunsch 1970). This is the basis for systems which generate alternation patterns, such as Albright & Hayes's (2006) Minimum Generalization Learner and Beniamine's (2018) Qumin. The generalization to multiple sequences is much more complex, especially for large number of sequences. Fortunately, this problem was researched in depth in evolutionary biology (see Durbin 1998) for the purpose of aligning DNA and protein sequences. Several heuristic methods thus exist, which can find good solutions in reasonable time, often by performing repeated pairwise alignments. Some of these methods were adapted to sequences of phonemes by List (2014) in order to compare potential cognate words across languages. Our implementation was written specifically for the alignment of inflectional paradigm. The algorithm we rely on is called 'progressive alignments' (Feng & Doolittle 1987), and proceeds in three steps:

1. Find all pairwise alignments, obtaining scores for each pair of form.
2. Using the scores, perform hierarchical clustering. We obtain a binary guide tree which relates sequences.
3. Align sequences pairwise along the guide tree, first joining two sequences together, then joining sequences to alignments, or alignments to alignments.

Our adaptation to inflectional paradigms has two important properties, building on Beniamine (2018). First, our scoring scheme uses phonological similarity (Frisch 1997). This measure relies on natural classes and captures satisfactory morphological alternations in pairwise alignment

(Albright & Hayes 2006; Beniamine 2018). Rather than minimizing edit distances, our algorithm maximizes similarity scores. Since they are to be expected in inflection, insertions (alignments to gaps) are free, while substitution costs are proportional to phonological similarity. The consequences of this strategy can be seen in Table 4, where columns contain either entirely identical or similar phonological segments. Second, there is often more than one optimal alignment at each intermediate pairwise step. Most algorithms deal with this ambiguity by selecting a solution randomly. Instead, we keep up to 30 competing alignments at all time. We then choose among these hypotheses based on the generality of the resulting segmentation.

3.2 Unique discontinuous stems

From an alignment table, we can easily identify the sequence of substrings which are constant across the paradigm: they appear in columns of identical segments (shaded in Table 4). This sequence can be discontinuous, and may be of different length for different lexemes. We call the phonemes in these columns 'inflectionally inert'.

Sequences of inert segments that are always adjacent across the paradigm are collected into strings. The unique discontinuous stem is the sequence of these strings. The sequence is continuous if no cell in the paradigm has exponential material occurring between inert segments. It is fragmented if and only if at least one cell splits the stem in at least two parts. In our example, the unique discontinuous /lib-ɾ-/ is fragmented into two substrings /lib/ and /ɾ/ due to the vocalic alternation between /ə/ and /ɛ/.

It is worth noting that unique discontinuous stems go against much of our muscle memory on the representation of morphophonology. Suppletion is not captured at all. If traditional suppletive stems have nothing in common (think, e.g., of Latin *fero* and *tuli*) they will be inferred to have an empty common unique stem—a reasonable conclusion. If however they by chance have a common substring (e.g., Latin *tuli* and *latum*), that will be considered to be the unique common substring (here ⟨tu⟩). Hence the length of the unique stem is not an indicator of suppletion in the traditional sense. Another potentially surprising aspect is that all phenomena of nonperipheral inflection lead to the same outcome of discontinuities in the unique stem, irrespective of whether they are caused by, e.g., infixation, vowel alternations, stress shift, or root and pattern morphology. While this may seem counterintuitive, we do not think it is problematic: the causes of discontinuity may be variegated, they still are all discontinuities, which raise the same challenges for speakers trying to draw inferences from a wordform.

Table 5 gives some statistics on the prevalence of fragmented stems in English, French, and European Portuguese.[3] Lexemes with zero stem fragment are cases of suppletion, although, as discussed above, some cases of suppletion will result in one or more fragments. Unsurprisingly, fragmented stems are attested but infrequent in both English (due to forms such as *ring, rang*) and French (e.g., *mener* /məne/, *mène* /mɛn/), but much more prevalent in European Portuguese, where they make up one third of the lexicon; the vast majority of these are due to the stress-conditioned vowel alternations documented above.

3.3 Sets of continuous stems

We can now build on discontinuous stems to deduce a set of continuous stems. The idea here is very simple: for each word, we count as its continuous stem the shortest sequence that encompasses all substrings of the lexeme's discontinuous stem. The stem set associated with a lexeme is then the set of continuous stems of its inflected forms. Table 6 illustrates.

[3] Data is derived from CELEX for English (Baayen, Piepenbrock & Gulikers 1995), *Flexique* for French (Bonami, Caron & Plancq 2014), and the Coimbra pronuciation dictionary for European Portuguese (Veiga, Candeias & Perdigão 2012).

Table 5. Stem fragmentation in three systems: number of fragments in a discontinuous stem

	0	1	2	3
English	2	5890	172	0
French	4	5136	109	0
Portuguese	2	1255	738	1

Table 6. Inference of continuous stem set

Cell	Form	Disc. stem	Cont. stem
1sg	libɛʀtu	⟨lib, ʀt⟩	libɛʀt
2sg	libɛʀtɐʃ	⟨lib, ʀt⟩	libɛʀt
3sg	libɛʀtɐ	⟨lib, ʀt⟩	libɛʀt
1pl	libəʀtɐmuʃ	⟨lib, ʀt⟩	libəʀt
2pl	libəʀtaiʃ	⟨lib, ʀt⟩	libəʀt
3pl	libɛʀtɐ̃ũ	⟨lib, ʀt⟩	libɛʀt
Continuous stem set		{libɛʀt, libəʀt}	

Note that this method uniformly arbitrates for shorter and less numerous stems—notably, theme vowels are not taken to be parts of stems. Other than that, we submit that it matches closely the intuition outlined above and shared by many linguists that the postulation of stem alternants is justified by alternations occurring internal to what is otherwise lexical material.

Table 7 documents the prevalence of stem allomorphy in the three systems. The number of stems for French is much lower than assumed in Bonami & Boyé (2002) and related literature. This is due to the fact that Bonami & Boyé (2002) does not implement Principle (3a), and integrates theme vowels and other alternating material occurring marginally with respect to inert material into stem allomorphs. The classification matches closely more traditional accounts such as Swiggers & van den Eynde (1987). Again, Portuguese stands out compared to the other two languages, thanks to the prevalence of vowel alternations.

Table 7. Stem allomorphy in three systems: number of distinct stem allomorphs per lexeme

	0	1	2	3	4
English	2	5909	129	24	0
French	4	5136	108	1	0
Portuguese	2	1255	630	102	7

4 How useful are continuous stem allomorphs?

We now turn to the question we asked at the end of §2: how helpful are the two notions of stem under investigation in capturing the inferential structure of paradigms, as characterized by the

IWRP and the PCFP?

In terms of the IWRP, the answer is quite simple. Sets of continuous stems are by definition less useful than a unique discontinuous stem: the unique discontinuous stem identifies *exactly* that part of the word that has no exponential value, while stem allomorphs blur the distinction between exponential and nonexponential material. In this connection it is worth reflecting on the paradigm of a verb such as *libertar*. In a form such as 3SG /libɛɾtɐ/, the /ɛ/ has partial exponential value: it indicates that we are dealing either with a SG or a 3PL form. In this it is not different from the /ɐ/, which also gives partial information that we are dealing with a 2SG, 3SG or 1PL form. But both the /ɛ/ and the /a/ also give partial lexical information, as both narrow down the class of lexemes that we might be dealing with. This illustrates how continuous stem allomorphs are unhelpful with respect to the IWRP: insisting on contiguity stops us from seeing that /ɛ/ makes a contribution different from that of the neighbouring segments but similar to that of more peripheral segments.

If sets of stem allomorphs are counterproductive when addressing the IWRP, maybe they are helpful for the PCFP. We thus turn to the relationship of the two notions of stem to implicative structure.

A first way of addressing this is to consider the implicative structure of the system of exponents. Let us define holistically the exponential part of a wordform as whatever remains when the stem has been removed. Table 8 contrasts what the exponents look like in the present indicative of LIRBERTAR depending on whether one works from the unique discontinuous stem or from the set of continuous stems.

Table 8. Two notions of exponence exemplified

	1sg	2sg	3sg	1pl	2pl	3pl
Full words	liˈbɛɾtu	liˈbɛɾtɐʃ	liˈbɛɾtɐ	libəɾˈtɐmuʃ	libəɾˈtaiʃ	liˈbɛɾtɐ̃ũ
Unique disc. stem				⟨lib, ɾt⟩		
Exponents	_ɛ_u	_ɛ_ɐʃ	_ɛ_ɐ	_ə_ɐmuʃ	_ə_aiʃ	_ɛ_ɐ̃ũ
Continuous stems	libɛɾt	libɛɾt	libɛɾt	libəɾt	libəɾt	libɛɾt
Exponents	_u	_ɐʃ	_ɐ	_ɐmuʃ	aiʃ	_ɐ̃ũ

Consider now the problem of predicting the exponent in one cell from the exponent in another cell. As the table should make clear, predictability is affected in both directions by one's view on stems. On the one hand, exponents are harder to predict under the unique stem hypothesis, since they contain more material—in this example, prethematic vowels, which may be nontrivial to predict. On the other hand, exponents are better predictors under this view, since they are larger and hence convey more information. In the case at hand, exponents based on the discontinuous stem include the prethematic vowel in the predictor form, which is partially predictive of its counterpart on the predicted form (Bonami & Luís 2014).

Given these two observations, and depending on how correlated central and peripheral exponence are in a particular system, one might expect unique discontinuous stems to lead to easier, harder, or equally difficult prediction than those based on sets of continuous stems. To establish this, we adapt the methodology of Ackerman & Malouf (2013): tables of exponents per lexeme and paradigm cell are used to estimate probability distributions of exponents by paradigm cell, taking into account the type frequency of exponents. We then compute the conditional entropy of choosing an exponent in a predicted cell from the choice of exponent for a predictor cell, and the average conditional entropy across ordered pairs of distinct cells. Table 9 reports the average results for English, French, and European Portuguese conjugation.

For all three languages we observe a variation of less than 0.03 bits, leading to the conclusion

Table 9. Average conditional entropy of exponents under two conceptions of stems

	Continuous	Discontinuous
English	0.9917	0.9655
French	0.7109	0.7091
Portuguese	0.5782	0.6016

that there is no difference to speak of between the two strategies in terms of how much implicative structure they capture. For French and English, where stems are overwhelmingly continuous, this is unsurprising. For Portuguese however we might have expected a larger difference. Figure 1 allows us to explore this in more detail. This figure reports, for a distillation of the paradigm (Stump & Finkel 2013), the difference between conditional entropy for exponents based on continuous and discontinuous stems. A positive value means that the cell in row is more predictive of the cell in column if we reason with continuous stems than if we reason with discontinuous stems. We can see that there are sizeable differences in predictability in both directions, notably located

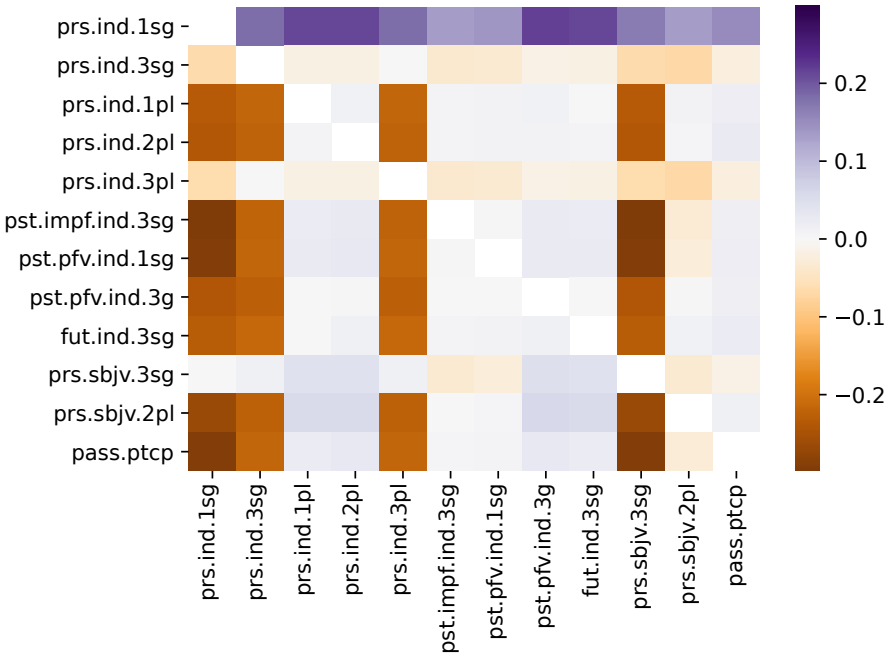

Figure 1. Heatmap of differences between conditional entropy of exponents based on continuous and dicontinuous stems, for a distillation of European Portuguese verbal paradigms

in the first two rows and two columns. However, the picture is almost symmetrical along the diagonal: where cell c is a better predictor of cell c' under continuous stems, cell c' tends to be a better predictor of cell c under discontinuous stems. This is easily explained by the distribution of prethematic vowels in Portuguese (Bonami & Luís, 2014). Negative values correspond to cases where the predicted cell has a stressed prethematic vowel that is hard to predict from a predictor cell with an unstressed prethematic vowel; positive values correspond to the opposite situation, where prediction of unstressed vowels is easy, but knowledge of the prethematic vowel helps predict the rest of the inflectional behavior.

Of course we would need to examine more languages to arrive at a general conclusion. However, based on the evidence examined in this chapter, although there are local predictive advantages to one or the other strategy for particular pairs of predictor and predictee, these balance each other

on average, so that the overall predictive value of exponents defined from continuous or discontinuous stems is equivalent.

We now turn to a different question on the predictive value of stems. The continuous stem set hypothesis induces a segmentation of words into two subparts that each have their own non-trivial implicative structure. The question then arises how the implicative structure of these stems compares to that of whole, unsegmented words. The presumption is that predicting stem allomorph from stem allomorph should be easier than predicting word from word: first, by focusing on the stem, we are abstracting away from the hard problem of predicting exponent variation across inflection classes, which has been the central focus of attention of the literature on predictability since Carstairs (1987) and Wurzel (1989). Second, stems are expected to give rise to little variation overall across the paradigm, which should make them easy to predict.

To assess whether this presumption is warranted, we computed 'implicative entropy' (Bonami & Beniamine 2016; Beniamine 2018) for paradigms of whole words and paradigms of continuous stems. Implicative entropy is the conditional entropy of the alternation pattern linking two forms given the phonological shape of the predictor form. Unlike conditional entropy computed from exponents, implicative entropy does not assume prior knowledge of a segmentation of the forms under consideration. Hence it is an adequate way of assessing the predictive power of whole words.

Table 10. Average implicative entropy for stem allomorphs and for whole words

	Continuous stem allomorphs	Words
French	0.0196	0.1844
English	0.0415	0.1739
Portuguese	0.1483	0.1670

Table 10 reports average implicative entropy across all pairs of cells for the three languages under consideration. In French and English, we do get the expected results that continuous stem allomorphs are more predictive of each other than words are. This was to be expected, given the low prevalence of stem allomorphy in these two languages. In Portuguese, however, it is barely harder to predict words from words than stems from stems. This state of affairs has a clearly identifiable cause. The main two sources of unpredictability in European Portuguese conjugation are theme vowel alternations and prethematic vowel alternations (Bonami & Luís 2014). While these two phenomena are mostly orthogonal and complementary, the statistical distribution of prethematic and theme vowels is not entirely independent. Thus, knowledge of one vowel is informative of the other, and it is easier to predict both at the same time from joint knowledge of the two. As a result, segmentation into continuous stems and exponents is unhelpful as it segregates two pieces of information that are more usefully brought together.

In this section we have established that a divide-and-conquer strategy, where stems and exponents are first cleanly separated, does not lead to better performance in addressing either the IWRP or the PCFP. We conclude that stem allomorphs play no major role in capturing the predictive structure of paradigms.

5 Conclusion

Although segmentation of words into subword units plays a central role in morphology, there is a lack of attention in the literature to the motivation of segmentation choices and the consequences of these choices for later theorizing. In this chapter we addressed only part of this general issue by concentrating on stems. We identified two simple, and coherent views on the nature of stems that are relatively easy to operationalize: unique discontinuous stems and sets of continuous stems.

We then presented an algorithmic method to infer both kinds of stems from raw paradigmatic data, and examined on that basis the usefulness of both conceptions in addressing what we take to be two central questions for morphology: predictability of forms from forms (the PCFP) and predictability of content from form (the IWRP). We concluded that the two concepts were of equally little use to address the PCFP, and that unique discontinuous stems were inherently more useful to addressing the IWRP. This leaves us questioning the usefulness of the concept of stem allomorphy for morphological theory.

There are different reasons to take this provocative conclusion with a bit of skepticism. First, the conclusion is dependent on particular choices when operationalizing segmentation principles: the devil is in the detail, and a different operationalization might lead to a different conclusion. For instance, note that our algorithm does not rely on any notion of optimization of the size of the lexicon, which typically plays a role in segmentation decisions. This is a conscious decision, motivated both by computational considerations and by literature disputing the usefulness of lexical optimization (see among many others Jackendoff 1975; Bochner 1993; Hay & Baayen 2005; Blevins 2006). However, it would be useful to see whether different operationalizations of segmentation principles, or the adoption of other principles, lead to different results.

Second, our reasoning here only applies to strictly synchronic aspects of morphology. Language change leads to situations where a unique continuous stem has multiple descendants that are related but not identical in form. As a way of describing that situation, 'stem allomorphy' is an essential concept in the description of morphological change. But importantly, such a concept of stem allomorphy is defined purely in terms of diachronic correspondences, not on anything that can be observed directly at a synchronic stage. Our conclusions hence lead to the question of whether a concept that is useful in the description of change is also useful in the description of the synchronic system.

Finally, some may find that we do not go far enough. Our point of departure was to address the IWRP, and more specifically to ask which aspects of a word's form provide information about lexical identity, as opposed to inflectional information. We went on to identify a unique discontinuous stem negatively, as the sequence that has no exponential value. But this cannot be the full story. As we noted before, in any system with multiple inflection classes, there will be material in a word that is partially indicative of lexical identity and partially indicative of morphosyntactic import. A full answer to the IWRP should not ignore this, and should ask instead, about each phonological segment in a word, what it contributes to narrowing down the content of the word. The unique discontinous stem hypothesis addresses only the easier part of that question—but at least it addresses it.

Acknowledgements

This paper stems from ideas which were presented at a workshop on New approaches to the typology of inflection systems, organized by the now defunct CNRS *Fédération Typologie et Universaux* in Paris, France in November 2018; as well as an invited talk funded by Stony Brook University at the 4th American International Morphology Meeting, in Stony Brook, NY, in May 2019. We thank the organizers and organizing institutions of both events for the opportunity to present our work, as well as the audiences for fruitful discussion. Interactions with Farrell Ackerman, Mark Aronoff, Matthew Carroll, Grev Corbett, Berthold Crysmann, Sebastian Fedden, Jean-Pierre Koenig, and Adina Williams were particularly valuable. We also thank Matías Guzmán Naranjo for numerous engaging conversations on the topic of paradigm alignments, and Ana R. Luís for teaching us most of what we know about European Portuguese conjugation.

We would not even be able to ask the questions we address in this chapter without Mark Aronoff's influence, both in the form of writings which have shaped our thinking on these is-

sues since the very beginning of our careers, and of friendly dialogue over the years. We can only hope that the present chapter will stimulate more of both.

This work was partially funded by a public grant overseen by the French National Research Agency (ANR) as part of the "Investissements d'Avenir" program (reference: ANR-10-LABX-0083).

References

Ackerman, Farrell, James P. Blevins & Robert Malouf. 2009. Parts and wholes: Implicative patterns in inflectional paradigms. In James P. Blevins & Juliette Blevins (eds.), *Analogy in grammar*, 54–82. Oxford: Oxford University Press.

Ackerman, Farrell & Robert Malouf. 2013. Morphological organization: The low conditional entropy conjecture. *Language* 89. 429–464.

Albright, Adam & Bruce Hayes. 2006. Modeling productivity with the gradual learning algorithm: The problem of accidentally exceptionless generalizations. In Gisbert Fanselow, Caroline Féry, Matthias Schlesewsky & Ralf Vogel (eds.), *Gradience in grammar: Generative perspectives*, 185–204. Oxford: Oxford University Press.

Aronoff, Mark. 1994. *Morphology by itself*. Cambridge: MIT Press.

Baayen, R. Harald, Richard Piepenbrock & Leon Gulikers. 1995. *Celex*. Philadelphia: Linguistic Data Consortium 2nd edn.

Beniamine, Sacha. 2018. *Typologie quantitative des systèmes de classes flexionnelles*. Paris: Université Paris Diderot dissertation.

Blevins, James. 2003. Stems and paradigms. *Language* 79. 737–767.

Blevins, James P. 2006. Word-based morphology. *Journal of Linguistics* 42. 531–573.

Blevins, James P. 2016. *Word and paradigm morphology*. Oxford: Oxford University Press.

Bochner, Harry. 1993. *Simplicity in generative morphology*. Berlin: Mouton de Gruyter.

Bonami, Olivier. 2012. Stems in inflection and lexeme formation. *Word Structure* 5(1). 1–6.

Bonami, Olivier & Sacha Beniamine. 2016. Joint predictiveness in inflectional paradigms. *Word Structure* 9(2). 156–182.

Bonami, Olivier & Gilles Boyé. 2002. Suppletion and stem dependency in inflectional morphology. In Franck Van Eynde, Lars Hellan & Dorothee Beerman (eds.), *The proceedings of the HPSG '01 conference*, 51–70. Stanford: CSLI Publications.

Bonami, Olivier & Gilles Boyé. 2014. De formes en thèmes. In Florence Villoing, Sarah Leroy & Sophie David (eds.), *Foisonnements morphologiques. etudes en hommage à françoise kerleroux*, 17–45. Presses Universitaires de Paris Ouest.

Bonami, Olivier, Gauthier Caron & Clément Plancq. 2014. Construction d'un lexique flexionnel phonétisé libre du français. In Franck Neveu, Peter Blumenthal, Linda Hriba, Annette Gerstenberg, Judith Meinschaefer & Sophie Prévost (eds.), *Actes du quatrième congrès mondial de linguistique française*, 2583–2596.

Bonami, Olivier & Ana R. Luís. 2014. Sur la morphologie implicative dans la conjugaison du portugais : une étude quantitative. In Jean-Léonard Léonard (ed.), *Morphologie flexionnelle et dialectologie romane. typologie(s) et modélisation(s).* (Mémoires de la Société de Linguistique de Paris 22), 111–151. Leuven: Peeters.

Boyé, Gilles & Patricia Cabredo Hofherr. 2006. The structure of allomorphy in spanish verbal inflection. In *Cuadernos de lingüística*, vol. 13, 9–24. Instituto Universitario Ortega y Gasset.

Brown, Dunstan. 1998. Stem lndexing and morphonological selection in the Russian verb: A network morphology account. In Ray Fabri, Albert Ortmann & Teresa Parodi (eds.), *Models of inflection*, 196–224. Niemeyer.

Cameron-Faulkner, Thea & Andrew Carstairs-McCarthy. 2000. Stem alternants as morphological signata: Evidence from blur avoidance in polish nouns. *Natural Language and Linguistic Theory* 18. 813–835.

Carstairs, Andrew. 1987. *Allomorphy in inflection*. London: Croom Helm.

Carstairs-McCarthy, Andrew. 1994. Inflection classes, gender, and the principle of contrast. *Language* 70. 737–788.

Corbett, Greville G. 2007. Canonical typology, suppletion and possible words. *Language* 83. 8–42.

Durbin, Richard. 1998. *Biological sequence analysis: Probabilistic models of proteins and nucleic acids*. Cambridge: Cambridge University Press.

Feng, Da-Fei & Russell F. Doolittle. 1987. Progressive sequence alignment as a prerequisitetto correct phylogenetic trees. *Journal of Molecular Evolution* 25(4). 351–360. `https://doi.org/10.1007/BF02603120`.

Frisch, Stefan. 1997. *Similarity and frequency in phonology*. Evanston, IL: Northwestern University dissertation.

Hay, Jennifer B. & R. Harald Baayen. 2005. Shifting paradigms: gradient structure in morphology. *TRENDS in Cognitive Science* 9. 342–348.

Jackendoff, Ray. 1975. Morphological and semantic regularities in the lexicon. *Language* 51. 639–671.

List, Johann-Mattis. 2014. *Sequence comparison in historical linguistics*. Düsseldorf: Düsseldorf University Press.

Maiden, Martin. 1992. Irregularity as a determinant of morphological change. *Journal of Linguistics* 28. 285–312.

Matthews, Peter Hugoe. 1972. *Inflectional morphology. A theoretical study based on aspects of Latin verb conjugation*. Cambridge: Cambridge University Press.

Montermini, Fabio & Olivier Bonami. 2013. Stem spaces and predictibility in verbal inflection. *Lingue e Linguaggio* 12. 171–190.

Needleman, Saul B. & Christian D. Wunsch. 1970. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of Molecular Biology* 48(3). 443 – 453. doi:https://doi.org/10.1016/0022-2836(70)90057-4.

Pirrelli, Vito & Marco Battista. 2000. The paradigmatic dimension of stem allomorphy in Italian verb inflection. *Rivista di Linguistica* 12.

Sims, Andrea D. & Jeff Parker. 2016. How inflection class systems work: On the informativity of implicative structure. *Word Structure* 9. 215–239.

Spencer, Andrew. 2012. Identifying stems. *Word Structure* 5. 88–108.

Stump, Gregory T. 2001. *Inflectional morphology. A theory of paradigm structure* (Cambridge Studies in Linguistics 93). Cambridge: Cambridge University Press.

Stump, Gregory T. & Raphael Finkel. 2013. *Morphological typology: From word to paradigm*. Cambridge: Cambridge University Press.

Swiggers, Pierre & Karel van den Eynde. 1987. La morphologie du verbe français. *ITL Review of Applied Linguistics* 77–78. 151–251.

Veiga, Arlindo Oliveira da, Sara Candeias & Fernando Perdigão. 2012. Generating a pronunciation dictionary for European Portuguese using a joint-sequence model with embedded stress assignment. *Journal of the Brazilian Computer Society* 19(2). 127–134.

Walther, Géraldine. 2013. *De la canonicité en morphologie: Perspective empirique, théorique et computationnelle*. Paris: Université Paris Diderot dissertation.

Walther, Géraldine & Benoît Sagot. 2011. Modélisation et implémentation de phénomènes flexionnels non-canoniques. *Traitement Automatique des Langues* 52(2). 91–122.

Wurzel, Wolfgang Ulrich. 1989. *Inflectional morphology and naturalness*. Dordrecht: Kluwer.