

Remarks on lexical knowledge

Edwin Williams

Cognitive Science Laboratory, Princeton University, 221 Nassau Street, Princeton, NJ 08540, USA

The abstractness of lexical knowledge, and its independence from the ‘words’ of a language, are generally underestimated. Abstract idioms demonstrate that the lexicon is the repository not only of fixed expressions, but of abstract language-particular phrase- and word-patterns as well. Paradigms are language-particular patternings dimensioned and structured in language-particular ways, with language-particular patterns of syncretism. The learning burden of these two abstract but learned features of lexical knowledge falls outside of what current thinking about language learning would allow.

1. Introduction

I will survey two aspects of lexical knowledge which I feel pose special problems for learning. In both cases the remarks will tend to emphasize the extensiveness abstractness, and at the same time the language-particularity of lexical knowledge, and will consequently magnify the learning problem. One conclusion that could be drawn from these observations is that intricate structures can be learned, and this learning is not adequately modelled by either parameter setting or list learning.

It is useful to distinguish two notions of lexicon, one Bloomfieldian and the other grammatical. The Bloomfieldian lexicon is the repository of all of a language’s idiosyncracies. The grammatical lexicon is the linguist’s theory of the category of linguistic object we call a ‘word’. These are quite different conceptions, but have come to be identified in modern times. The result is a picture of grammar in which we have a clean streamlined syntax, where syntax is the theory of phrase, and a notion of word that entails that words are idiosyncratic and irregular at heart, with partial regularities expressed as ‘redundancy rules’ whose name implies low level generalizations where exception is the rule.

I think that this picture is wrong. A more correct picture is I think given in Disciullo and Williams (1986). In that view, both the word formation and the

syntactic system are ‘clean’ streamlined systems, independent of the lexicon. The lexicon is the repository of forms about which something special must be learned. This certainly includes, for example, all the monomorphemic words, but it also includes composed units of both syntax and word formation, composed units whose properties do not all follow from the rules which compose them. This would include a great deal of words, but also a great deal of phrases. In fact, I now think there are more lexical phrases than there are lexical words, but this remains a speculation.

In addition, I think that lexical knowledge includes knowledge of complex abstract structures that cannot be arrived at through parameter setting, and which must be learned from the data in a strong sense. I will discuss two of these: ‘abstract idioms’ in section 2, and paradigms in section 3. If my view of these is correct, the rich structure each exhibits results not from a rich innate predefined linguistic structure, but rather from a richly structured learning strategy.

If nothing else, then, the remarks that follow draw attention to the learning problems that might arise if first, more of linguistic knowledge is lexical than has been thought, and second, if acquired lexical knowledge is more abstract and structured in language-particular ways than has been thought. I think that the sort of structure that is found in each case reveals the hand of the learning strategy.

2. Idioms

The phrases in the lexicon are called ‘idioms’. We generally think of idioms as ‘frozen’ expressions. I will use the term *idiom* to refer to any defined unit whose definition does not predict all of its properties. This will include the phrase *kick the bucket*, a phrase, whose idiomatic meaning is *die*, but also *transmission*, which unpredictably means ‘such and such car part’. We don’t think of *transmission* as an idiom, but it will be useful and I believe correct to include it.

It has been commonly assumed that idioms are well-formed structures, but I think lately this view has been questioned. To say that an idiom is well formed is to say simply that it conforms to the rules of formation for the type that it is. Of course, the idiom does not conform to all of the rules, else it would not be an idiom. What rules do idioms obey? It is useful to consider the rules of ‘form’ apart from the rules of interpretation. We find idioms violating both types of rules.

The idioms that do not obey the rules of interpretation are deviant in meaning. The rules of argument structure are not obeyed in some cases, and in others, the rules of reference are not obeyed. We might refer to both these types of idioms as ‘semantic’ idioms.

It is perhaps misleading for me to say that there are idioms which do not obey the rules of form, for it is an overwhelming fact that idioms obey the basic rules of form in a language. For example, all idioms in English obey the ‘head-initial’ setting of the head position parameter. On the other hand, there are unpredictable, language-particular exploitations of the formal possibilities in a language; we might call these idioms ‘formal’ idioms.

A widely shared set of assumptions about idioms is the following:

- (1) *usual assumptions*:
 - idioms are listed
 - idioms are well-formed phrases
 - idioms have empty parts:
 - The cat has Xs tongue
 - relation of idioms to syntax:
 - insert the idiom, then fill in the parts

So, for example, *The cat has got X’s tongue* is an idiomatic phrase, fully consistent with the laws of English syntax, which is inserted in a phrase marker for an S position; further substitution of a referential NP for the position of X will yield an English sentence.

These assumptions, though common, lead to some surprising conclusions, or at least some surprising conclusions, or at least some surprising questions. In particular, they lead to the view that a great deal of what we have called ‘rule’ in syntax might really be ‘idiom’. To give an example, consider the fact that embedded questions in English must begin with a WH word. It is well known that this is not to be described by making WH movement obligatory in English. What kind of information is it then that [Wh-phrase S] is a (potential) embedded question in English? It COULD be that this form is an idiom:

- (2) [[Wh-phrase] S]_s: ‘embedded question’

Under this conception, (2) could be listed in the lexicon, awaiting insertion into some embedded S position, where it is an embedded question (as the slogan after the colon above indicates). Further substitution of some wh-

phrase for the wh-phrase position in (2) and of an S for the S position in (2) will complete the sentence, exactly parallel to *The cat has got X's tongue*.

One does not ordinarily think of this as idiomatic information, even though it is language particular; rather, one thinks of this feature of English as one of the 'parametric' possibilities. (2) is not the ordinary way to encode this information.

Example (2) may seem sufficiently different from *The cat has got your tongue* that it would never be misidentified as an idiom in the same sense. However, I intend to supply enough examples of cases intermediate between the two that I cannot see how it cannot be questioned whether (2) is the correct description or not.

2.1. Idioms are instances of well-formed structures

The view that I am presenting here would have large numbers of objects of a sometimes quite abstract character stored in the idiom lexicon of language. This view might then be seen as converging with several recent proposals that there are 'constructions' in language, a backlash against the principles and parameters model. There is, though, a fundamental difference between the proposal here and those proposals – specifically, I hold that idioms are well formed, and that the rules of well-formedness are simple 'parameterized' rules.

2.1.1. Filmore's examples

For example: C. Filmore (p.c.) has suggested to me that while most prepositions in English are pre-positions, some are not, and consequently the notion that there is a uniform head position in English is not correct; the most one can say is that there are a number of constructions in which the head is leftmost, and some others in which it is not. His candidates for postposition in English are *notwithstanding* and *ago*:

(3a) John notwithstanding, we will go there tomorrow.

(3b) John left 3 days ago.

I believe that neither of these is a preposition, and that the notion that English is head initial can be maintained in a strong form.

Notwithstanding can be assimilated to the following constructions:

(4a) John aside, ...

That noted, ...

which are clearly small clause constructions of some kind. The difference between *notwithstanding* and *aside* is that *aside* can appear in other small clause constructions, whereas *notwithstanding* is restricted to the context indicated:

- (5a) put that aside,
 (5b) *notwithstanding.

We may express this restriction on *notwithstanding* by not listing it in the lexicon on its own, but only as a part of the following ‘idiom’:

- (6) [[NP notwithstanding]_{sc} S]_s: ‘even with NP, S’

Importantly, (6) is an instance of a structure which is well formed in English independently, namely the structure of (4):

- (7) [NP AP] S > >
 [NP notwithstanding] S: ‘even’

I will use the double carat sign ‘> >’ to mean ‘has a well-formed instance’.

I think *ago* as well has a better analysis than the postpositional one. Consider:

- (8a) long ago
 (8b) 5 minutes in the past
 (8c) a few days before the party

It appears that time prepositions in general can take some sort of extent specification; however, this specification in general precedes the time preposition, as we might expect specifiers to. (8c, d) shows clearly that the extent specification is not a part of the complement structure, as the prepositions *in* and *before* have complements to their right. Now, *ago* differs from these in two ways. First, it cannot have a complement to the right. This however shows nothing except that *ago* is intransitive. Secondly, in the case of *ago*, the extent specification is obligatory. Here, if we think that extent specification is ordinarily outside of the subcategorizational reach of a head, we might appeal to an idiom to capture this exceptional feature of *ago*:

- (9) 'assign complements right, specs left' > >
 [(Extent-spec) P (NP)]_{pp} > >
 [Extent-spec ago]_{pp}

Here, the idiom is a well-formed subinstance of prepositional phrase, which itself is a 'subinstance' of all the structures that fall under the parametric specializations of English.

In sum then, although the notion of idiom we are employing is quite general, it is at the same time quite restrictive, in that all idioms conform to more abstract (though perhaps parameterized language particular) principles.

2.1.2. What rules do idioms obey?

Idioms basically obey the (rules determining the) phrase structure of the language, but not the rules of reference; so for example, *kick the bucket* is a well-formed VP, but no reference to any bucket is made. These two observations set some outer bounds. I think a more interesting probe is theta structure, where it turns out in fact that idioms differ in the extent of their conformity.

It is well known that idioms differ in their syntactic 'transparency' – 'kick the bucket' does not passivize, whereas 'cross that bridge when we come to it' does. Fiengo (1974) connected this difference to a further difference, namely, a difference in whether the object can be quantified:

- (10a) *They kicked the buckets (as an idiom)
 (10b) We will cross those bridges when we get to them.

In (b), 'bridges' refers to 'problems', and (b) says that we are talking about more than one problem. In (a) on the other hand, it is impossible to refer to more than one death.

Is this an arbitrary difference? I do not think so; I think it follows from the fact that the meaning of 'kick the bucket' is intransitive. This intransitivity of meaning means that there can not be a correspondence between the syntactic and semantic argument structures, since syntax is transitive here. In (b) on the other hand, the verb *cross* can be assigned a transitive theta structure, and one of its theta roles can be assigned to *bridges*. So what is wrong with the meaning of (b) is simply the reference, not the theta structure: *cross* refers to *solve* and *bridges* refers to *problems*. But in (a), it is not just reference, but theta structure itself, which is idiosyncratic.

If this is so, then (b) has a transitive theta structure, and its passivizability is no surprise, if passive is an operation on theta structures, as I believe it is.

And the fact (a) is not passivizable is no surprise either, since there is no theta relation between the direct object and the rest of the sentence.

If this account is correct, then theta structure is not necessarily respected in idioms, though it may be. A child learning an idiom will aggressively assign it as much structure as possible according to its rules. If it can assign it a semantic argument structure based on its syntactic argument structure, it will, but if not, then the idiom will lack argument structure.

2.2. Abstract idioms

2.2.1. Idioms with instances

If we begin with the idea that an idiom is a phrase, and that it can contain empty parts, then we immediately face the question, how empty can an idiom be? Can it, for example, be mostly empty? The question arises sharply for an example like the English noun *pants*. This noun, as is well known, is an 'arbitrary' plural; a shirt, for example, of the same genus topologically (at least when the fly of the pants is unzipped) is singular. This is a trivial piece of idiomatic, that is unpredictable, information about this noun, that it must be plural. Strikingly, though, this is not true of this noun alone, but in fact of every word that has the meaning that pants has: something worn on the legs in such and such a way:

(11) pants jeans shorts cutoffs culottes bermudas ...

New, made-up terms for lower trunk wear must conform as well. One exception, bathing-suit, is an exception only in that it does not specifically refer to lower-trunk wear, but rather means whatever one wears to bathe in; it is an accidental fact of current fashion that this refers to lower-trunk wear.

Now, what sort of information is this? It is information about a general restriction on form that follows from meaning; specifically, if a noun is going to have such and such a meaning, then it must be plural;

(12) Ns ←: 'lower trunk wear'

I have drawn the arrow from right to left to mean, if an item is going to have the meaning on the right, then it **MUST** have the form on the left. This is a different sort of idiomatic information from knowing that *kick the bucket* CAN mean *die*, and so the different notation.

Another example like *pants* is *fish*. Fish names, with some exceptions, are all unmarked plurals or have that as an option: trout, bass, perch, bream, yellowtail, mahimahi. The exceptions are not really fish, by and large: whale, guppy, minnow. Other animal families are untouched by this idiosyncrasy: bee, wasp, ant. As far as I can tell, this is an unpredictable, but very general fact about English, and counts as ‘idiomatic’ information about the language.

A further case of the same sort is the language-particular patterns of lexicalization identified by Talmy (1985). He found that languages systematically differed in the kinds of verb meanings they allowed. For example, English allows verbs of motion to indicate a means of motion. *Float* can be used as a directional verb, but at the same time, it indicates the manner of motion: *John floated under the bridge* can mean that John moved under the bridge by floating. Spanish and French lack entirely verbs of this kind. *Flotter* in French and *flotar* in Spanish (float) can mean only to float stationarily, and the restriction is apparently a hard and fast one. Similarly, verbs of posture (sit, kneel, lie, etc.) differ systematically from language to language, in whether the stative, inchoative, or causative is the basic or underived form; English, Japanese, and Spanish systematically differ in this choice – see Talmy (1985) for details.

Again, we have language-particular variation of a quite general sort. Again, the information is ‘idiomatic’, but the question remains, how to represent it. We might again represent it as an idiom ‘with a hole in it’:

- (13a) [inchoative]_V ←: posture verb (Japanese)
 where V is atomic.
- (13b) [motion verb]_V ←: manner (English)

Now, the representation here I think is not so important as the question of what this sort of information is, and especially how it is acquired. I think that language-particular patterns of the kind that have just been discussed – *pants* in English, the motion verbs in Romance, the posture verbs, etc. – fall outside of the ‘parametric’ core, and yet, they are quite general, and basically exceptionless. This means some sort of general induction, the kind that is meant to extract ‘lexical redundancy rules’, must be capable of acquiring from the data the ‘exceptionlessness’ of the rule. Some substitute for negative evidence, such as counting, and statisticking, is required. And some limitation on the space of searches must hold, in order for the induction to remain in the realm of possibility.

The most surprising cases are the cases that achieve exceptionlessness. The 'pants' idiom is exceptionless, in that any noun that means the right thing must participate in the idiom. The present participle is another case: there are no present participles that do not end in *-ing* whereas past tense forms, as is well known, are quite varied. The learner learns more than that there are no exceptions; he learns that there can be no exceptions. We will discuss such cases further in section 4. For the moment, we note that the formalism proposed suffices to express exceptionlessness, though it at the same time hides the learning problem implied by them.

2.2.2. Formal idioms: Exploited and unexploited avenues

The *pants* phenomenon just examined is I think more widespread. Formally, the grammar permits singular nouns to have the meaning 'such and such a type of legwear', but this avenue is unexploited in English, thanks to (12). By 'unexploited avenues' I mean possibilities that the formal system would seem to allow, but which it does not use. In identifying such cases one always risks missing the formal explanation for the missing possibility, though a case like *pants* I think clearly shows that there will not always be one.

As one example, consider compound terms in English and French. Both languages have means of putting together words from further words, or kind-denoting terms. In French, the syntactic system is exploited; so, one has compound terms of the following kind:

- (14) VP >> V NP >> V N >> *essuie-glace* 'wipe window', 'wind-
shield-wipers'
VP >> V PP >> V P N >> *laisse-pour-compte* 'left for count',
'abandoned one'

Here, the double carat means 'has as an instance'. So, compound terms in French are instances of syntactic constructions. As word level items they have their own limitations (e.g., no referential material may occur in them, and so, for example, no determiners are allowed), but they are nevertheless well-formed syntactic objects.

English, on the other hand, exploits a different system to form its compound terms:

- (15) [X Y]_Y >> [N N]_N

The system exploited here is the affixation system in the lexicon, which is right-headed. Ordinarily, Y is a suffix, forming the head of a word. English lets Y be a full noun, giving us compounds.

Importantly, both languages have both resources – English has the same (left-headed) syntax as French, and French has the same right-headed affixation system as English; however, they each exploit a different one of these for their compound terms. I assume that this is ‘idiomatic’ – that is, language-particular, but perhaps not ‘parametric’.

2.3. Syntactic idioms

English embedded questions must begin with a *wh*-phrase. This is not to say that *Wh*-movement is obligatory in English, as it clearly is not; not only do matrix questions not necessarily undergo *Wh*-movement, but even a *wh*-word in an embedded question need not move, when, for example, another has moved, or the complementizer is already *wh*, like *whether*:

(16) Who wonders whether George saw who?

The correct generalization is as stated: a *Wh*-word must appear at the beginning of an embedded Question. What sort of information is that? We might describe it as an idiom, in the sense developed herein:

(17) [*wh*-phrase] S ←: embedded question

That is, a sentence with a *Wh*-phrase at the beginning is a question, and nothing else is. The arrow goes backwards, as any embedded question must have this form.

Idiom (17) is a good candidate for a ‘parameter’, in that there is a small number of ways that question words can be dealt with: (1) moved to front; (2) moved to verb (as in Hungarian); or (3) left in situ. However, I think there is good evidence that idioms just like (17) must be countenanced, ones that are not reducible to parameters of variation.

One case is the ‘amount’ relatives. These have the following form:

(18a) [*wh*-phrase S] > >
 [what N' S]: Little (amount relatives)

(18b) I gave him what food I had.

(19) I give him what I had.

(18b) has an implication that (19) does not have, namely, that there was little food in question.

Now, where does this implication come from? It does not come from *what*, which does not have this implication in general, not even in free relatives, except in the context in (18b). Furthermore, it does not inhere in free relatives in general. In fact, it occurs only in the structure in (18), it is idiosyncratic to that structure. Assuming that there is no parameter to set here, then this is a learned fact about this structure.

What is interesting is how formally similar (18) is to (17). The only difference is that the semantics of (18) is very particular, and therefore plausibly idiomatic, whereas the semantics of (17) is very general. But the formal means may be the same in the two cases: a feature of meaning and form are connected in an idiomatic entry in the lexicon. If a learner can induce (18), it would seem that (19) would be accessible to the same mechanism.

A related sort of case arises from Subject Aux Inversion in English; the following is an idiom of English:

(20) [V S]_s ←: matrix yes/no question

This is comparable to (17) – an obligatory idiom (that is, of the *pants* variety) has the effect of forcing syntactic rules to apply. More interesting are the cases of inversion which receive a conditional interpretation:

(21a) Had I been there, this would not have happened

(21b) S > > [V S]_s: conditional

> > [had...]_s

> > [were...]_{VP}

The rule of inversion gives a large number of forms which are ungrammatical in this context:

(22) *Could I write poetry, I would not be a linguist.

In fact, inversion in the conditional context works only for the auxiliaries *had* and *were*. What sort of information is this? Importantly, the cases allowed in the construction are a subset of the cases allowed in general; hence, what is learned is that not all the formally allowed possibilities are realized. We will adopt the following convention for representing this situation:

(23) *Instance principle*:

If a form to which a meaning is assigned has listed subinstances, then those subinstances are exhaustive.

2.4. *Idiom families*

At one end of the language-particular information that a learner must acquire are the completely fixed expressions; at the other end are the broad typological parameters. I have suggested that there are intermediate ‘abstract’ idioms – *pants* and amount relatives, for example, which link these endpoints with a graded continuum. In compensation for this more complicated situation, and the more complex learning problem that it poses, I have suggested that each level of ‘abstractness’ must conform to the level above it; thus we do not have a wholesale theory of ‘constructions’, but still a broadly parametric model. This says that a construction – say, passive – must conform to the typological pattern that is determined by some parameter settings, but leaves open the possibility that not all features of the passive construction will be determined by this conformation.

As further evidence of this view, I will discuss here some idiom families, that is, language-particular idiom patterns. Each idiom pattern has a number of idioms as instances. The principle reason for recognizing the existence of the idiom family is that some languages will have idioms of a certain sort, and others will lack them altogether, apparently in a way not related to the parametric typology of the languages, though of course one could always be wrong about that for any particular case.

The most interesting sort of case I know of was pointed out to me by Martin Everaert (p.c.); the idioms are of the form:

- (24) $N_x P N_x$
side by side

The two Ns are meant to be identical tokens of the same noun, as in *side by side*. French, English, and Dutch have these idioms, Japanese lacks them entirely. For this reason, we would want to call (24) itself an ‘abstract idiom’ of the sort discussed in previous sections.

Idiom (24) probably has as its most immediate instantiations not actual idioms, but further, more concrete, idiom families, one for each P that participates:

- (25) $N_x P N_x > >$
 N_x by N_x
 N_x to N_x
 N_x on N_x
 N_x for N_x
 N_x from N_x

Some of the many examples of instances in these subfamilies are:

- (26) N to N
 $> >$ hour minute second year
window door station
head hand cheek
 N by N
 $> >$ minute, hour, etc.;
layer house cave
 N for N
 $> >$ dollar for dollar
 N from N
 $> >$ limb from limb
 N on N
 $> >$ layer on layer

If these are idioms, then we expect them to be instances of well-formed structures in the language. One is at first hard-pressed to identify what structures these are well-formed instances of, but I think the most likely candidate is the compound prepositional phrases, as *from ... to*:

- (27) John went [from a to b]

The compound prepositions generally require both prepositions to be overtly present, but this is not always so, as in the following:

- (28) I do not know what he is doing [one day to the next].

Thus the compound prepositions are somewhat like conjunctions, where we find both simple (*and*) and compound (*both ... and*) varieties. It so happens that the idiom family so productive in English is a subspecies of compound prepositional phrases where the first preposition is not present. In French, on

the other hand, both prepositions must be present: *de part en part* ('limb from limb') (V. Deprez, p.c.; T. Hoekstra has pointed out to me the existence of *heure par heure* and *côte à côte*).

The various subcases of (24) do not have a common element of meaning. So, for example, *cheek to cheek* refers to the pressing together of two cheeks, as in dancing; but *minute to minute* and *day by day* refers to a series of days in sequence. Even the instances sharing a common proposition do not have a completely shared element of meaning, as the two instances with *to* just cited show. Therefore, these forms are not compositional, despite clear patterns in the meaning.

On the other hand, from the fact that they are so prevalent in one language, and absent altogether in another, we know that they are present as a group in some sense. It seems unlikely that there is a parameter for this property alone; perhaps it follows from some other parameters, though it is hard to see how.

An alternative is that the structure in (25)–(27) is induced from the data of the language. How could this happen? Suppose that at a certain point in the course of acquisition, some number of forms with the shape N_xPN_x have been learned. As a class, they conform to the shape of leftheaded compound prepositional phrases, and so do not fall outside of the language altogether. On the other hand, their properties are not entirely explicable in terms of the general principles of the grammar; in particular, the use of bare singular count Ns as the objects of prepositions is not a general feature of prepositional phrases in English. So these remain idiomatic; however, they are idiomatic as a class, not as individuals.

Another family of idioms is illustrated in the following:

- (29) John hunts bear
 snares rabbit
 traps monkey
- (a) *John hunts book
 *I am going to grade paper
- (b) *John counts monkey
- (c) *hunt sleepy elephant

What is special about this case is the use of the bare singular as object. Normally, this is not allowed for English count nouns, but is allowed here. So, we have an abstract idiom, of the following form:

- (30) [V N]_{VP}
 + HUNT + ANIMAL

The form in (30) is a special case of the general form of VP (and so, for example, is 'V N' and not 'N V'), and thus conforms to our overall claim that abstract idioms are always instances of more general patterns of the language. The limits of this idiom are somewhat roughly indicated by (b) and (c) – the verb must be a verb of hunting, and the object must be an animal. (d) shows a further restriction – not only are determiners excluded, but adjectives as well (unless 'sleepy' is a kind of elephant).

One might conclude from this the observation that this construction was in fact a lexically compound verb – to 'hunt monkey' then would be syntactically intransitive. I doubt this, since English in general disallows compound verbs, and particularly disallows left-headed compound verbs; but even if the conjecture were correct, then the problem posed by these examples is not solved, but simply delivered to the lexicon, with fee still unpaid.

I assume that it is not at all predictable that English would have this pattern; and in fact, inspection of (b) and (c) might lead one to not expect this pattern. I conclude therefore, that learning (30) entails generalizing over examples like those in (29), and the limits of the generalization must follow in some way from the actual mechanism of generalization.

The existence of these idioms of intermediate abstractness argues that learning language does not reduce to (a) learning parameter settings, and (b) learning the properties of particular lexical items. Rather, there are structures between these two extremes, what I have called abstract idioms, which can only be learned as language-particular generalizations of particular forms.

3. Paradigms

A paradigm is a multidimensional array of linguistic forms – for example, a verb conjugation, or a Latin noun declension. A paradigm is not just a convenient way to display linguistic information; rather, it is a basic form of linguistic knowledge. It interests here because it is highly language-particular, and at the same time, quite abstract in structure. Paradigmatic structure is also pervasive.

Example (a) below is a slice of the Latin verbal conjugation:

(31) *Latin:*

(a) + – finite	+ – indicative	(b) amo	amamus
+ – passive	+ – perfective	amas	amatis
pres/imperf/fut		amat	amant
+ – plural	1/2/3person		

The Latin verbal conjugation is 8-dimensional, with the dimensions listed on the left.

It is not possible to say in a general way how many dimensions a paradigm will have, nor how many points on a given dimension, nor what the dimensions will ‘mean’, or what syntactic or semantic categories the points on a dimension will be taken as signifying.

There are several levels of abstraction involved in paradigms. At the lowest level, we have word-paradigms, such as in (31b). At a slightly higher level of abstraction, we have paradigms in the traditional sense, roughly speaking, sets of endings:

(32)

-o	-mus
-as	-atis
-t	-ant

At a slightly greater level of abstraction, one might regard a paradigm as a set of rules which, when applied to a stem, derive a word paradigm, by, for example, adding endings. So there is a rule for forming the past tense, a rule for forming the third person present tense, etc. In fact, though, I believe that a paradigm is even more abstract than that – a paradigm is a patterning which is more abstract than any set of forms, any set of endings, or any set of rules for filling the slots in a paradigm.

That the paradigm is a real object, and not the epiphenomenal product of various rules, is shown by the phenomena of blocking, syncretism, suppletion, and paradigm defectiveness, as argued in Halle (1973).

To consider only the first of these: if there are two rules for filling a slot in a paradigm, only one may be used; thus, we have *bit*, not *bited*, and in general, only a single past tense form for a given verb, despite multiple ways to form past tenses. This reveals that there is a target slot to fill, which is independent of the rules for filling it, and that slot is given by the paradigm.

3.1. Extensiveness

We tend to think of paradigms as a means of displaying inflectional information about the parts of speech. But in fact, the notion of paradigm is much broader than this. To begin with, paradigms must include syntactic items, or phrases, as well as words. To see this, we need only examine a slice of the Latin verb paradigm:

(33)	<i>active</i>	<i>passive</i>	
	present	amo	amor
	perfect	amavi	amatus sum

Here the forms are all 1st singular, present-perfective crossed with active-passive. One corner of this square contains a phrase, *amatus sum*, while the other three corners contain words. This shows that phrases form an inextricable part of paradigmatic information. If we removed the perfective passive form, we would have destroyed the symmetry of the paradigm, which is otherwise perfectly symmetrical.

We can see the same thing in an English paradigm, the comparative paradigm:

(34)	<i>A</i>	<i>COMP</i>	<i>SUPER</i>
	long	longer	longest
	compact	more compact	most compact
	good	better	best

The rule is, if an adjective is mono- (or nearly) syllabic, then form the comparative with *-er*; if not, then the comparative and superlative are formed phrasally. The existence of this paradigm is what permits us to speak of 'the comparative' of an adjective, even though there are two ways of forming comparatives. Many languages lack any way at all to form the comparative; English has two ways, one morphological, the other syntactic.

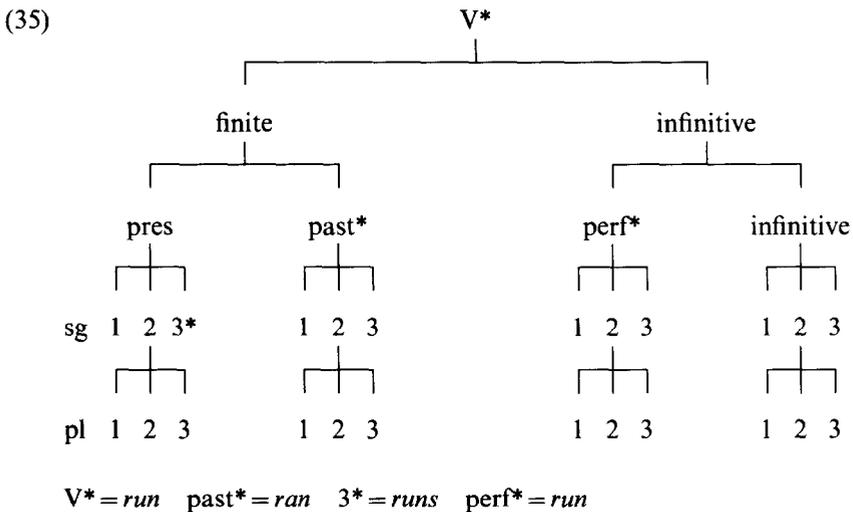
Paradigms include not just inflectional dimensions, but what have been called 'derivational' processes as well. I am sure that there is no distinction between derivational and inflectional morphology, but if there is, then paradigms are found in both morphologies.

3.2. The abstractness of paradigm structure

In general, there is not a one-to-one relation between the slots in a paradigm and the rules for creating the forms which fill the paradigm. In one circumstance, there are more slots than rules, and therefore, several slots will be filled with the same form – this is the phenomenon of syncretism. The pattern of syncretism that is found reveals the method of assigning forms to slots. The method consists of the following:

- (1) Hierarchalize the dimensions of the paradigm.
Assign forms to nodes in the hierarchy.

For example, the English verbal paradigm is 4-dimensional (person X number X present-past X finite), and there are 4 distinct forms. We may hierarchize the tree and assign the forms to it as follows:

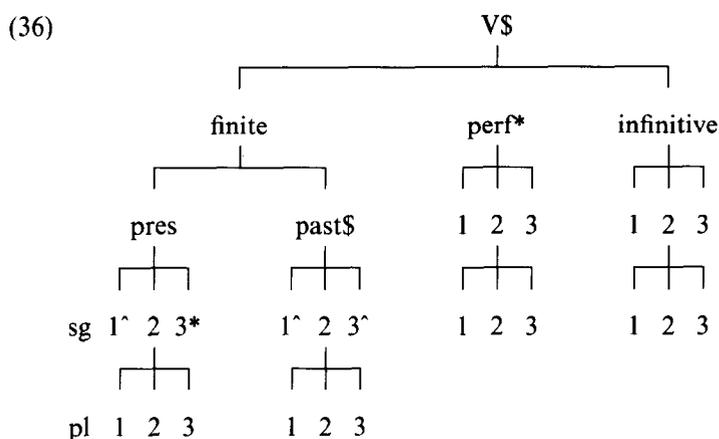


The terminal nodes are the actual cells of the paradigm. The starred nodes are the nodes to which actual forms are assigned. By convention, a cell is filled by the nearest specified form above it. The identity of the forms and the points in the tree at which they are mapped are given in (35). The assignment shown is the most economical, as each form is assigned once.

We might call the starred nodes ‘entry points’ – these are the points at which concrete forms are specified. The tree along with the starred nodes we

might call the 'pattern of syncretism'. This is a pattern which is independent of the rules for creating the forms in the pattern. It is a part of the formal structure of the paradigm. Strikingly, the pattern of syncretism holds generally in a category, not just for particular verbs, as we will see.

If we look at the pattern of syncretism for a variety of English verbs, a startling pattern emerges:



Entry points:

\$ modals

\$* regular; go-went

\$*^ be

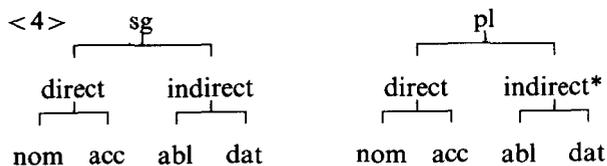
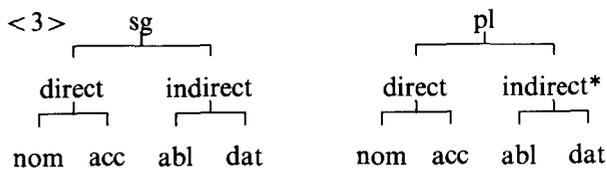
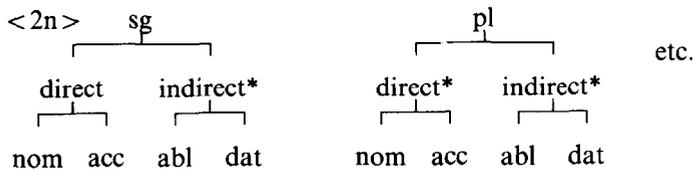
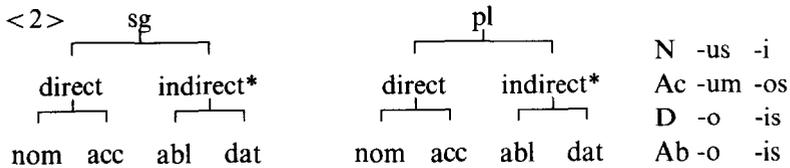
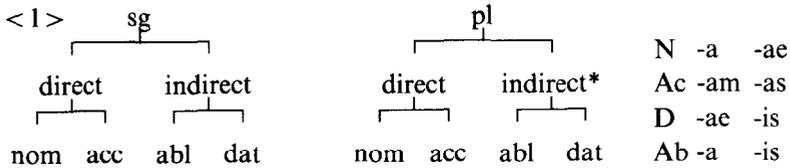
have, says, does – irregular in 3prsg: *haves

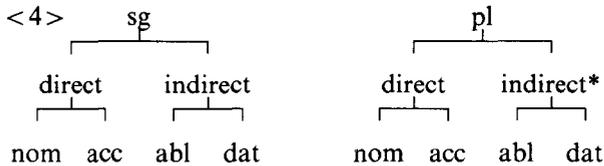
Here, marked on the same tree, are the entry points for several categories of verbs, including irregular verbs. As can be seen, the sets of entry points form a nested set; the verb *be* shows the most distinctions, and consequently has the most entry points, but all other verbs, including all irregulars, have some subset of the entry points of *be*. It is far from obvious that such a relation should exist – if a verb is going to be irregular, why should it not be irregular in having a different pattern of syncretism, a different set of entry points? But this does not happen, even irregular verbs respect the pattern of the language as a whole. In fact, even suppletive verbs, the limiting case of irregularity, respects the pattern of syncretism; the verb *go* has *went* as its past tense form. Things could have been different: *went* could have been the third past plural form, with *goed* (or something else) for all the other forms; but then, *go-went* would have violated the language-wide pattern of syncretism.

Hence, the pattern of syncretism is a quite abstract structure, standing above particular words, particular rules, particular suppletive relationships.

We can see this further in the noun declensions of Latin. Latin has 5 declensions, each with its own set of endings (we ignore here the genitive and the vocative):

Latin declension structure:





principle syncretisms:

1st decl: pl. indirect (-is)

2nd decl: pl. indirect (is), sg. indirect (-o)

neuter: above + direct sg. (-um), pl. direct (-a)

3rd decl.: pl. indirect(-ibus)

4th decl.: pl. direct (-es) pl. indirect(-ibus)

neuter: sg. (direct = indirect) (- \emptyset), pl. indirect, pl. direct

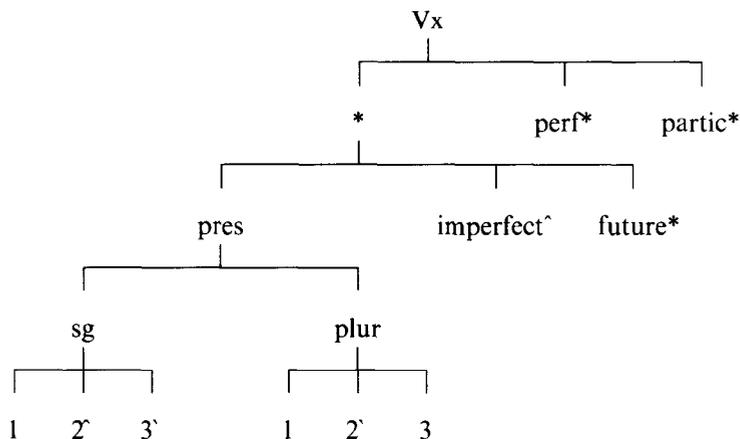
Here, the nominative and accusative have been grouped together as 'direct', and the ablative and dative as 'indirect'. The reason for this is that this reflects the patterns of syncretism: nominative and accusative fall together sometimes, and dative and ablative do as well.

A striking thread that runs through the entire set of declensions is the indirect plural syncretism. In the singular, there is an indirect entry point for 2nd and 4th neuter; the plural syncretism holds across the board. Importantly, this generalization is independent of the rules for forming the indirect forms, for in fact there are two different rules for that: in the first and second, *-is* is affixed; whereas in the 3rd and 5th, *-ibus* is affixed. Hence, the pattern is more abstract than the rules or affixes. I would suggest that a pattern is abstracted that applies to all the declensions, essentially the tree structure in $\langle 1 \rangle$. Other declensions will make further syncretisms, but this one will hold for all.

A general conclusion we may draw is that when there are multiple related paradigms, there will be one instantiated paradigm, and all others will have its syncretic structure, and perhaps some more. But no other related paradigm will have a contrary syncretic structure, making distinctions where that one does not. We will call that one paradigm the basic paradigm.

For the Latin nominal declension, the first declension is the basic paradigm. For English verbs, the verb *to be* is the basic paradigm.

Let us now consider Latin verbs. Every finite Latin verb form has a different form, so there is no syncretism at this level. But at the level of stem, there is a paradigm structure with a limited number of entry points. Below is a chart of the stem forms for various classes of Latin verbs:

(37) *Latin verb stem suppletion:**Entry points for stem:*

- *^` esse, edo
- *^` volo
- *^ regularized volo
- *^ do
- * regulars, including fero feri tuli latus
- x amo

The asterisked positions represent the entry points for the regular verbs of all four conjugations. This alone is striking, for again, there are different rules in the different conjugations for yielding the forms at these entry points. The first and second conjugation form the future by suffixing ‘-b-’, whereas the third suffixes nothing, but switches the conjugational class of the basic stem. Nevertheless, that the future is an entry point is common to all the conjugations.

The verb *esse* has the most entry points, and every other verb uses some subset of those entry points. Thus the conjugation of *esse* is the basic conjugation, to which all the others are related.

As a final example, we consider Anderson’s (1984) description of Georgian verb conjugation. The system is quite complex; Anderson’s account uses blocks of rules, both conjunctively and disjunctively ordered. It is my contention that such a system will fail to capture the most abstract patterning of a paradigm, which, as we have seen, is generally independent of affixes or

rules. Some indication that this is so can be derived from the following table of affixes (or, to use Anderson's account more directly, of rules for adding affixes) for present and past:

(38) <i>present</i>	<i>past</i>
-en 3pl subj	-es 3pl subj
-t pl	-t pl
-s 3rd subject	-s 3rd subject

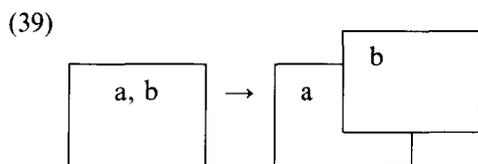
The rules are quite similar, and Anderson's remark 'This *-es* rule has the same status as the *-en* rule in this block, and also takes precedence of the *-s* and the *-t* rules ...' (1984: 8) shows I believe that an abstract structure is being replicated in different parts of the paradigm, thus underscoring the independence of that structure from actual rules or affixes.

3.3. Learning paradigm structure

We ask at this point, why are these patterns of syncretism in language? Why is a pattern of syncretism replicated across different modes of realizing paradigm cells?

I speculate that it is the acquisition of paradigm structure that is responsible for this arrangement.

Pinker has demonstrated how the blocking principle will give rise to the development of paradigm structure in the language learner. The basic idea is that whenever the language learner has been forced to posit two items to fill a single cell, he is then motivated to split the paradigm (really, to double it) so as to avoid violating the blocking principle:



The paradigm now has a new dimension, and a whole new set of cells; the language learner now must learn the 'significance' of the dimension, and fill in the rest of the cells.

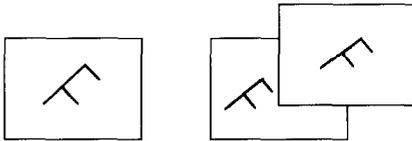
I think that this algorithm for building paradigms, combined with the notion of paradigm structure I have just outlined, will predict some of the patterns we have observed.

There are two sorts of splitting than can take place – intra- and inter-paradigm splitting. Given a present tense paradigm for some verb, ‘see’, for example, the recognition that ‘saw’, like ‘see’, can be used for the 1st person singular, leads to a postulation of a past tense plane, in which to locate ‘saw’. This is intra-paradigm splitting. The other kind of splitting might be called ‘declensional’ splitting: given a set of endings, say the endings for the 1st declension in Latin (*-a, -am, -ae, -a, etc.*), the recognition that the ending *-us* can signify nominative singular, just as *-a* does, triggers the splitting of the nouns into (at least) two declensions. The difference between this and the first case is that in the first case, a given word will have forms in every cell of the new and old plane, but in the second, a word will have forms in only a single declension. We may nevertheless consider the splitting to be formally the same in the two cases, and this is supported by the observation that syncretism patterns the same in the two cases.

An example of a ‘declensional’ split is the comparative paradigm in English; the paradigm is a linear three-point paradigm: adjective, comparative, superlative. But there are two modes of forming members, as we have seen – one for monosyllabics and simple disyllabics, and another for everything else. A telling point which shows that we are dealing with a declensional split here is that the same criterion that is used to determine whether the comparative is *A-er* or ‘more *A*’ is used to determine whether the superlative is *A-est* or ‘most *A*’. This criterion is not therefore a part of the rules themselves, but is rather a general criterion of membership in the two declensions; much as ‘feminine’ is a criterion for membership in the Latin 1st declension.

Now, suppose that a learner has learned a piece of paradigm structure, and has learned not only the labels for the dimensions, but has also learned the positioning of the entry points, which I have called the pattern of syncretism. Suppose further that when the paradigm is split, that this abstract pattern of syncretism is replicated along with the cells themselves:

(40)



If this is done, then we will expect to see patterns of syncretism recurring. The following prediction is made: whatever paradigm is learned first will embody

the most distinctions. This is because cells which are designated identical by the pattern of syncretism of the first pattern will remain identical in latter versions. There may be fewer distinctions in later folds, but not more.

Thus for example when the learner learns the first declension, with its indirect plural syncretism, and then learns that there is a second declension, that same indirect plural syncretism will show up in the second declension as well, copied as a part of the abstract structure of the first declension.

4. Learning words

4.1. *Learning morphemes*

The notion of an abstract paradigm, the blocking principle, and paradigm splitting may account for how a paradigm is elaborated, but what accounts for the identification of a (potential) paradigm in the first place? One ordinarily thinks of verbal paradigms, realizing person, number, tense reference, etc., as the typical paradigm, but in fact languages have novel paradigms that it is unlikely could be specifically anticipated in UG.

Suppose that one component of the learner is a device that uses extra cycles in the child's computational life to track down statistical correlations among various properties of its thus far stored linguistic units. What the set of properties is will not detain us here. The child learning English, for example, will discover a correlation between words ending in *-y* and adjectivehood:

(41a) fishy, lumpy, lucky, speedy, etc.

(41b) dainty, pretty, happy, etc.

Many adjectives do not end in *-y*, and many words ending in *-y* are not adjectives, but the probability that a word is an adjective increases once one knows that it ends in *-y*.

Note that this is true for a large class of adjectives where the *-y* does not serve as a suffix (the second group). Even excluding the cases where *-y* is an affix, there is a correlation.

A battery of such correlations could serve the next step: to identify morphemes, and assign them properties. In the case of $-y \rightarrow A$, there are sufficient cases to warrant postulating a suffix with the category adjective, and assign it wherever possible – that is, wherever an independent stem exists.

Even where the analysis into morphemes does not hold, the information of the correlation has been noted; it seems unlikely that this information is forgotten once the morpheme has been established.

Another example in English of a no-go morpheme is the suffix *-tude*; most of the 50 or so members do not admit an analysis into morphemes:

(42) altitude attitude platitude, etc.

And yet, this suffix so strongly marks nounhood that there is not a single verb or adjective with this ending.

The number of examples it takes to establish a correlation between two properties is quite small. For example, the English noun-deriving suffix *-al* has fewer than 20 instances; and yet, the correlation between suffix and category seems firmly established for all speakers, as well as the restriction that the stem must be stress final, a separate learned correlation:

(43) betrothal, avowal, approval, removal ...

The confidence of these identifications is surprising given the small number of cases involved.

So, there are two levels of analysis so far; first, the identification of correlating properties, and second, the postulation of analysis that arises from this. Units will be postulated that account for the correlation.

There is a further level of analysis, I believe. In some cases, it is determined by the learner not just that two properties correlate, but that one of the properties is uniquely correlated with the other property. We might suppose that children are built to look for this especially.

Several examples have been examined in this paper. One is the idiom Ns: 'wear on the legs'; the property of meaning 'wear them on your legs' is uniquely correlated with nouns that end in 's'. Another is the progressive, which is uniquely correlated with V + ing.

I believe that the uniqueness has special salience. It is what we might consider the 'ideal' case, and so the first sought and most readily accepted. We know, for example, that the past tense forms of verbs in English are irregular, and some are in fact suppletive. Given this, we are not surprised so much to learn new verbs with irregular past tenses. However, I think we would resist entirely learning a new verb with a suppletive progressive, a progressive that did not end in *-ing*.

4.2. Learning subsystems

In order to get a paradigm going, one must notice parallel correlations. For example, in English, there is a Latinate paradigm, of the following shape:

(44)	<i>verb</i>	<i>adjective</i>	<i>de-verbal</i>	<i>de-adjectival</i>
			<i>noun</i>	<i>noun</i>
	conclude	conclusive	conclusion	conclusivity

The strength of this paradigm is great enough that when a coined word meets the criterion for membership, it immediately participates in the entire paradigm:

(45) subduct: subductive: subduction: subductivity

How do we know that we are dealing with a paradigm here, and not simply with several distinguishable morphological rules? Because the correlations amongst the rules striking: for example, 80% of the words that take *-ive* to form adjectives take *-ion* to form nouns; this is as compared to the base rate of 6% of nouns that take *-ion*. Surely speakers notice this sort of correlation, and use it to construct a paradigm like (44), which can then be exercised on new forms.

The criteria for membership in this paradigm are somewhat rough, but include at least the following:

(46) word = prefix-stem
 prefix = {trans, sub, super, de, dis, con, etc.}
 stem = CVC(C).

The force of these criteria can be seen in the fact that 26% of the nouns that begin with *trans-* form nominalizations in *-ion*, against the base rate of 6% for nouns in general.

Incidentally, this sort of thing is not limited to the English Latinate vocabulary. 39% of the verbs beginning with the causative prefix *en-* (*encroach*, *endear*, *encase*, *ensnare*) form their nominalization in *-ment*, as opposed to the background rate of less than 5%.

We ordinarily think of *-ion* as an 'unproductive' suffix, compared, say, to the suffix *-ness*. However, as we have seen, in the domain of certain classes of stems, *-ion* approaches complete productivity, and even becomes applicable to new forms. The other side of the coin is that *-ness* is actually not so

productive over the entire set of stems in English – it is of course only in the adjectives that it shines. So *-ness* and *-ion* are equally productive, we might say, over their own classes – the class for *-ness* is adjectives, and the class for *-ion* is Latinate verbs. Given this, we might wonder, why isn't every affix completely productive within some arbitrarily drawn subclass of the lexicon, say the class of things to which it does attach? I think the answer is that subclasses cannot be arbitrary. The Latinate class is productive for *-ion* precisely because it can be identified independently of the occurrences of *-ion*: it can be identified as the class to which *-ive* attaches, or perhaps it can be identified in some way along the lines of (46). In either case, we must attribute to the language learner the ability and the inclination to compare subclasses, and look for high matches. When a high match is found, then a dimension of a paradigm has been identified.

References

- Anderson, S., 1984. On representation in morphology. *NLLT* 2, 157–218.
- Disciullo, A.M. and E. Williams, 1986. On the definition of word. Cambridge, MA: MIT Press.
- Fiengo, R., 1974. Semantic conditions on surface structure. MIT dissertation.
- Halle, M., 1973. Prolegomena to a theory of word formation. *Linguistic Inquiry* 4, 3–16.
- Talmy, L., 1985. Patterns of lexicalization. In: T. Shopen (ed.), *Language typology and language description*. Cambridge, MA: Cambridge University Press.
- Williams, E., 1981. On the notions 'lexically related' and 'head of a word'. *Linguistic Inquiry* 12, 245–274.