

Approaching Grammar: Towards an Empirical Linguistic Research Programme

Holger Keibel and Marc Kupietz

Institute for the German Language (IDS), Mannheim

Abstract

Any attempt to capture the apparent regularities and system-like properties of a natural language in a theory aiming for explanation is faced with a wide gap between the desired level of theoretical description and the phenomenologically accessible level of language use. To make progress to this end, the authors advocate an emergentist perspective on language according to which all rule-like and conventional aspects of a language are an epiphenomenon of language use. On this conceptual basis, they outline a research programme that attempts to mimic the processes that are driving this emergence, in order to incrementally arrive at similar generalisations as the language community.

1. Terms and definitions

We start with a couple of definitions. In this paper, a *theory* is taken to be a set of assertions about “something” with (roughly) the following properties: (i) the assertions are formulated precisely and unambiguously; (ii) they are justified rationally (not metaphysically); and (iii) they are logically consistent such that there are no obvious contradictions between them.

One can distinguish at least three kinds of theories: explanatory, descriptive, and normative theories. The subject of *explanatory theories* are usually objects in the *real* world. Such theories can be tested empirically, i.e., they are falsifiable. Propositions and hypotheses in an explanatory theory are connected by more general principles – much in contrast to (purely) descriptive theories. The motivation underlying explanatory theories is typically gain of knowledge and, of course, explanation.

The subject of *normative theories*, by contrast, are not objects in the real world but rather objects in an *ideal* world. This is why such theories are generally not falsifiable. Normative theories are usually motivated and justified by aesthetic and/or ethic considerations. Examples for normative theories as defined here can be found among theories of society and theories of science, in areas of law and in theories of individual languages (e.g., normative grammars).

Although normative and explanatory theories have quite different characteristics, they may nevertheless in part be influenced by the same factors. This is also true for the case of grammar as illustrated in Fig. 1.

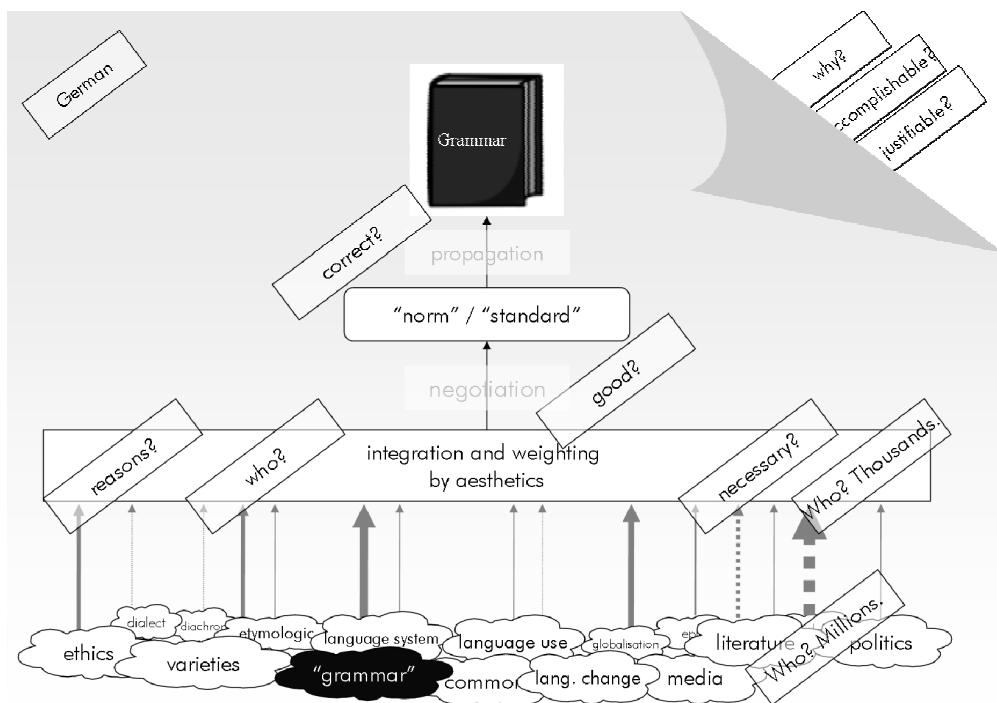


Figure 1: The process of constructing a normative grammar.

A normative grammar generally incorporates a number of different factors: apart from the probably most influential factors (marked by stronger arrows) such as *language system* and “*grammar*”, there are also factors such as *literature*, *media*, *politics* etc. Because these factors are so different in nature, they can only be integrated into a single theory via aesthetics. More precisely: Because the various factors are incommensurable, their assessment and integration can only be decided and justified aesthetically. This involves processes of negotiation which finally lead to a *norm* or *standard* which can then be communicated in the form of specific grammar books. To avoid misunderstandings from the outset, it should be stressed that although aesthetics plays a central role in the integration of the relevant factors, it has generally fairly little influence as a factor of its own. That is, we do *not* argue here that normative grammars primarily reflect the individual linguistic preferences of their authors. The process of developing an explanatory theory of grammar may also in part be positioned within Fig. 1. Such a theory is essentially influenced by factors such as *language use*, “*grammar*” or *language system*.

Note the crucial distinction between two readings of *grammar*: As Stetter (1999: 172) puts it, in the normative reading, a grammar is *written* (“*ge-schrieben*”) which inevitably involves an aesthetic line of argument. By contrast, in a descriptive or explanatory reading, grammar is *described* (“*be-schrieben*”) – aesthetic arguments would be neither appropriate nor necessary. We are convinced that the objectives underlying these two readings are both well-justified directions in linguistics. One cannot, however, reasonably pursue both kinds of objectives simultaneously.

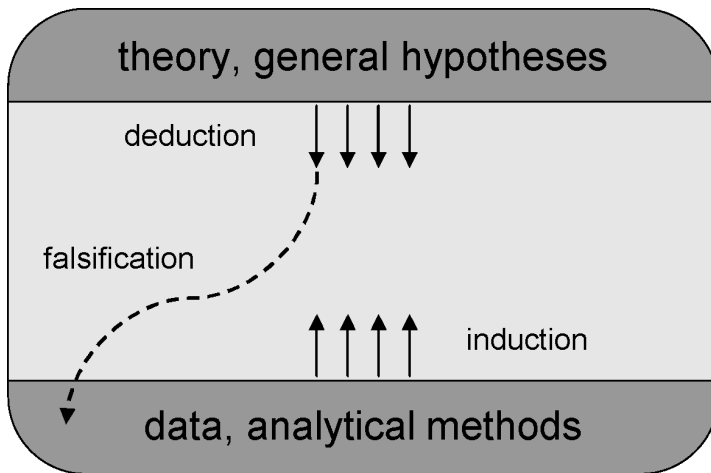


Figure 2: The explanatory gap.

The final term to be defined only concerns explanatory theories. Irrespective of the specific object of investigation (unless a trivial one), any attempt to construct an explanatory theory is faced with what one may call an *explanatory gap*, namely a gap between the relevant basis of empirical data with appropriate analytical methods on the one hand and a theoretical level with hypotheses about general principles on the other hand (cf. Fig. 2). For the case of an utterly complex object such as language, it is not to be expected that this gap will ever be closed in its full width. What one can expect, however, is to decrease the gap somewhat in various places. One approach to do so is to start at the lower end of the gap and generalise in an inductive and data-driven fashion from concrete observations to more abstract structures. Likewise, one may start at the upper end and derive, in a deductive and theory-driven fashion, specific predictions from general hypotheses. These predictions would then have to be tested empirically against the data – generally by means of falsification as a test procedure.

However, we believe that real progress with respect to explanatory theories is only possible if the pursuits on both ends go hand in hand. Only operating at the theoretical end without any contact to real data would not be a reasonable approach; and likewise, blindly digging in the data without any direction in mind will not lead very far. Unfortunately, corpus linguistics is often perceived either as operating exclusively at the theoretical level and confining itself to criticising Chomsky

(1965), or as operating only at the data level, developing a versatile body of methods but declining all scientific responsibility beyond these methods. This perception is without doubt not entirely accidental. In our view, however, it lies at the core of corpus linguistics – and of empirical linguistics in general – to address the explanatory gap as a whole.

2. Some fundamental decisions in empirical linguistics

The above considerations concerning the explanatory gap underscore that any promising approach aiming for explanatory theories is necessarily rooted in empirical research. However, empirical observation plays a central role not only in testing hypotheses but also from the outset in the course of developing a theory. Before elaborating on this second aspect in the next section, we discuss some fundamental assumptions and research strategies that we think should be avoided (2.1.) or chosen (2.2.) in any empirical-linguistic attempt to developing an explanatory theory.

2.1. Rejected assumptions

There are a number of credos that are fairly common in theoretical linguistics that should be avoided – not because these credos may have been proven to be wrong, but rather because they have not been proven to be true. This concerns (i) the assumption that language can be adequately captured as a formal system; (ii) the assumption that *decomposition* applies without any restrictions as a principle of explanation; (iii) defining a *complete theory* as a research objective; and (iv) selecting *linguistic competence* as the object of investigation. Below we explain briefly for each of these credos why we think they should be rejected, before outlining what other assumptions should take their place.

Language may be adequately captured as a formal system

Given what is currently known about natural languages, we see no proper justification for assuming that they may be captured as formal systems in any explanatorily adequate way. Moreover, the fact that language is a dynamic cultural artefact renders this assumption rather unlikely. One might argue that, if language itself is an artefact, its creator is surely not; thus it might be possible to adequately capture human beings as formal systems – for instance, in analogy to a symbolically programmed computer. If this secondary assumption is true, so the argument, then the (language) behaviour of human beings and language as a result of this behaviour should display some characteristics of a formal system, as well. However, looking at 50 years of AI research, we see no strong evidence supporting this secondary assumption either.

Of course, these two assumptions have not been proven wrong either, so it is possible that future research will uncover sufficient evidence supporting them. But such issues should be the result, rather than the starting point of empirical-linguistic research. Hence, it currently seems advisable not to make these assumptions. As a consequence, one should not assume a priori the

existence of abstract categories and formal rules – neither with respect to psychology nor with respect to the language system.

Decomposition is universally applicable as a principle of explanation

Closely related to the assumption just discussed is the assumption that utterances and sentences as well as any categories in language can always, and without loss of information, be decomposed hierarchically into more basic components. However, according to our knowledge, it is again an open research question whether and to what extent decomposition (e.g., into modules, constituents, categories) can be applied as an analytical-explanatory methodology. The factual existence of perfect hierarchical structures has not yet been demonstrated, neither with respect to the mind nor with respect to language. And as long as this is the case, we consider it advisable to avoid this assumption.

Completeness as a research objective

Some branches of theoretical linguistics pursue the long-term goal of formulating a complete theory for a given individual language. Within generativist research programmes, this goal even has constitutive status. However, when an individual natural language as viewed as a dynamic artefact which cannot necessarily be captured as a formal system, it does not appear to be a reasonable methodological decision to seek a single complete and coherent theory. First of all, this goal would most likely never be reached for there is always a way to violate any alleged rule without leaving the language (cf. section 2.2.). And second, setting out this goal is not even useful for it leads to bad strategies: When modifying an existing theory towards *more* completeness, it becomes increasingly complex such that its explanatory value will generally decrease. The point is that with an artefact such as language, one cannot expect to find simple and complete explanations. In turn, incompleteness is no argument against a theory. As an alternative objective, we therefore propose local models (cf. section 3.6.).

Linguistic competence as the object of investigation

Inasmuch as explanation is the goal, it is in principle not a viable approach to select the linguistic competence as the object of investigation, for theories of linguistic competence are by nature not explanatory but normative. This is because linguistic competence refers to an ideal speaker and as such is not part of the real world but rather a construct pertaining to an ideal world. In consequence, theories of linguistic competence may be assessed in aesthetic terms, but they are in principle not falsifiable. With a deliberate trace of polemics one might even argue that these theories are in most cases also motivated and justified primarily by aesthetic considerations, namely in that they are inspired by an ideal computer (i.e., the Turing machine) as their blueprint and in that they are shaped by principles of mathematical elegance. In a slightly different context, Stetter (1999: 172) has described such theories as the “products of a ›technical‹ fantasy” (original: “Produkte einer

›technischen‹ Phantasie”); translation ours).

2.2. Advocated assumptions

After having stated which kinds of assumptions we do not consider reasonable to make with respect to explanatory theories, we now outline an alternative approach. Our point of departure is a strictly usage-based perspective on language – such a perspective arises quite naturally, if not inevitably, for any empirical-linguistic approach. In any case, there does not seem to be any ontologically sound way to define the phenomenon *language* independently of language use.

In this respect, the term *language use* is somewhat misleading for it seems to suggest that speakers make use of an autonomous object named *language* which exists independently of speaking and writing. According to such a view, the language system (i.e., *langue*) would be essentially static: Although it may undergo long-term processes of language change, it is thought to have definite and fixed properties in each point in time. In fact, such a static nature of the language system is often (implicitly) assumed even for greater time intervals, and it is this assumption that establishes the contrast between synchronic and diachronic linguistics. However, this view of the language system as a static object that exists independently of language use does not withstand scrutiny – in what part of the real world could such an object with these properties exist?

Instead, the language system is a cultural artefact which is not directly accessible by phenomenological experience but only indirectly through factual events of language use (i.e., *parole*). And without an autonomous existence of the language system, the only viable conclusion is that language use is not only an epistemological window to the language system, but also its primary cause. In other words, the language system exists only as an emergent epiphenomenon of language use. Such an approach to language is, of course, anything but new, but also the more recent literature contains a growing body of similar emergentist positions, especially within the so-called *usage-based framework* (e.g., Hopper 1987, 1998; Bybee 1998; see also the papers in Bybee & Hopper 2001; Langacker 1987, 2000; Goldberg 2006; Elman et al. 1999; Tomasello 2003).

The influence between language use and the language system is in fact bidirectional. The individual language experience shapes the procedural language knowledge of each member of a given language community, that is, it shapes their cognitive routines for the production and processing of language material (see also Sinclair’s *idiom principle*, 1991, and Hoey’s concept of *lexical priming*, 2005). Because the language experience of different speakers is never identical, they may develop very different cognitive language routines. In particular, they may have very different grammars *in their heads*. Of course, this individual grammar of a speaker influences their language productions and thereby contributes, in turn, to the language experience of other individuals.¹

The language system – or rather: those parts of language that we perceive as system-like and

¹ Of course, normative grammars can also be part of an individual’s experience with language.

governed by regularities – can be characterised as conventions of the language community which may vary with respect to their stability and their distribution among members of the community. These conventions are negotiated in an ongoing process. And it is precisely this ongoing process that characterises the emergence of the language system. Potentially, this negotiation may take place in every single linguistic interaction such that new conventions may arise any time between any two speakers or in small groups. Of course, for most languages, the language community as a whole never gets together, instead it negotiates conventions across many such individual interactions. Keller (2006) characterised this phenomenon also as “coordination without a coordinator” (translation ours).

It must be stressed that these processes of negotiation do not only concern language change but the innermost core of language. Language is always fluid, always dynamic and adaptive. Considering the large number of individuals contributing to these processes, one may wonder why language displays such a high degree of conventionality in the first place. There are several potential causes for this fact, including the following ones: First, we are all members of the same species, equipped with similar brains and bodies. Second, we live in essentially the same world and are faced with very similar problems and challenges. Third, we are all generally interested in the success of our communication efforts, if not dependent on it. Therefore, speakers will tend to choose expressions with a good probability they will be understood roughly as intended. Fourth and closely related, for economic and other reasons, speakers tend to imitate – generally without being aware of it – language behaviour they have observed in similar communicative situations. Especially these last two factors contribute to the observed degree of stability of conventions, and also to the fact that the conventions do influence the language routines and, hence, the language productions of individual speakers. While the consequences of these last two factors are necessarily language-specific, the former two also contribute to language universals (e.g., cf. Tomasello 2003).²

In sum, one important consequence of taking an emergentist perspective on language is that grammatical rules in the sense of stable and authoritative objects do not exist in natural languages. Instead there are regularities and rule-like patterns in the form of emergent structural conventions whose degree of conventionality may vary considerably. These conventions are intrinsically instable, dynamic and context-dependent, and given a specific context, they can be *violated* in a meaningful way without leaving the language. When many such violations aggregate, these conventions may change.

² Media, schools, normative grammars and other factors also contribute to the formation and relative stability of conventions, but it is likely that they do so mainly via language use.

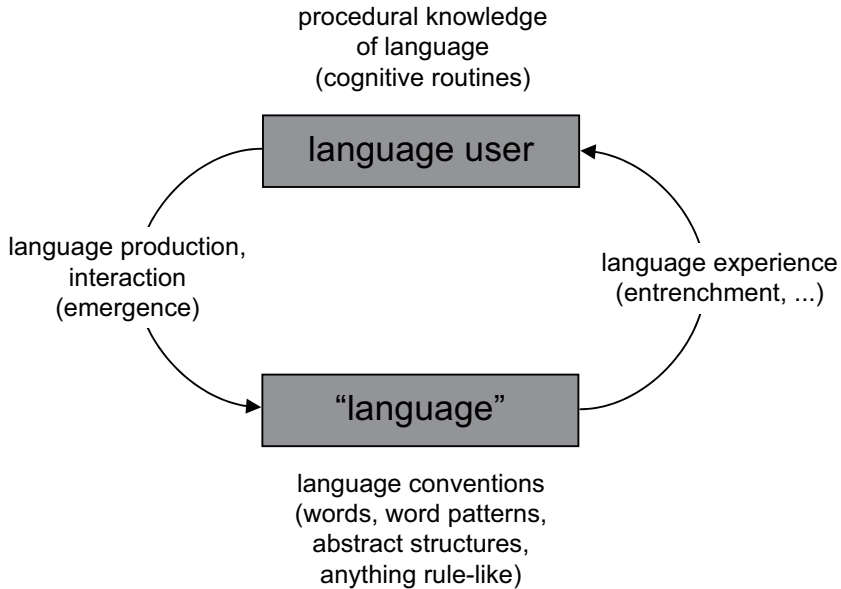


Figure 3: Mutual influence between "language" and individual language users.

Fig. 3 is a schematic summary of how the individual speakers' procedural knowledge of language influences "language" and vice versa, and how language structures emerge as a result of this mutual influence. Note that here and in what follows, "language" (in quotes) is intended to simultaneously refer to both the observable language use and the emergent language system. This mutual influence is key for the methodological access to the study of the phenomenon "language" as will be described in the next section.

3. An empirical-linguistic research programme

Given an emergentist perspective on language, one may conceive a number of different possible approaches aimed at explanatory theories which all comply with the requirements specified in the previous section. Of these possibilities we outline a specific empirical-linguistic research programme that distinguishes objects of investigation at three levels.

3.1. Object of investigation

At the central level – the level of "language" (with the two readings described above) – the specific object of investigation depends on the given question to be studied and also on the intended kind of theory, if applicable. Possible objects of investigation range from rather narrow language domains such as the language of a specific author, to more general domains such as doctor-patient communication, and further on to highly general domains such as *the* (written)

German language.

Because normally this central level cannot be accessed by direct observation, a second level comes into play. This level consists of a specific corpus of text or speech, that is, a collection of recorded acts of communication along with some relevant context (e.g., written texts, audio files, video data). The corpus has to match the nature of the given object of investigation at the central level, in the sense that it can be considered a valid sample of the statistical population defined by this object. At the same time, objects of investigation can also be formulated directly at the corpus level, when the research questions or hypotheses to be studied do not go beyond the corpus. In such a case, the corpus is both sample and population.

The third, upper level concerns the communicator (i.e., the speaker or hearer). It comprises not only psychological aspects of the communicator but also their immediate situational environment. We believe that explanatory theories are very unlikely, if not impossible, to be formed unless this psychological level is taken into account.

Hypotheses may be formulated at each of these three levels but the way by which such a hypothesis has to be tested crucially depends on the respective level. At the lower level (viz., the corpus), the situation is fairly simple: Each hypothesis can be unambiguously verified or falsified directly on the corpus. Any insights obtained at this level continues to be valid as long as the corpus remains unchanged. At the central level (“language”), the situation is somewhat more complicated: Any hypothesis would have to be tested on one or multiple appropriate corpora. An ultimate verification, however, is generally not possible in this case, hence falsification is the typical strategy of hypothesis testing at this level. An exception, of course, is the trivial case that the given language domain is entirely identical with the corpus. Testing hypotheses at the upper level probably lies outside the scope of empirical linguistics and rather falls into the discipline of experimental psychology.

3.2. Methodological access to “language”

How can meaningful hypotheses at the central level be generated and motivated in the first place? This question constitutes one of the key challenges in this research programme. Because “language” as a whole cannot be accessed directly, the principal direction of reasoning, of course, will be induction from the corpus (the sample) to “language” (the population). However, a blind and unrestricted induction from the data will in most cases not lead to hypotheses about “language” that are potentially explanatory. The basic idea to avoid this trap is to take advantage of the mutual influence between “language” and the level of communicators discussed above.

This methodological access to “language” may be summarised metaphorically like this: The goal is to trace the kinds of generalisations that speakers would *in principle* arrive at if their language experience were taken from the given corpus data. Note that this metaphor already incorporates both directions of the mutual influence: The generalisations are part of the individual speakers’ procedural knowledge of language (as a result of the influence of “language” on

individual knowledge), and as such they influence the individuals' language behaviour and thereby the emergent conventions (this is the opposite influence). Thus, by abstracting from individual speakers and observing directly the generalisations that *some* speaker would arrive at *in principle*, this emergent effect normally involving many individual speakers is already included in the metaphor. At first glance the concept of an *abstracted speaker* may be somewhat reminiscent of the concept of an *ideal speaker*, but there are in fact some fundamental differences between these two terms. Unlike the ideal speaker, the abstracted speaker is not the object of investigation but some practical construct. Strictly speaking, it does not even describe any speaker – neither real nor ideal – but rather the effects of language on the individual knowledge of language, averaged across many individual speakers.

In order to arrive at similar generalisations as the given language community, the abstracted speaker is modelled by means of known psychological facts and well-justified assumptions. These considerations suggest the following research strategy: (i) Collect a range of psychological premises about speakers (known facts and well-justified assumptions) that might play a role in generalisation processes; (ii) operationalise these premises in structure-detecting corpus-linguistic methods; (iii) systematically apply these methods to the given corpus; (iv) interpret the resulting generalisations as hypotheses about “language”. Note that the level of the communicator only serves as an intermediate link in this strategy, but no objects of investigation are defined at this level itself.

The strategy is justified in that psychological mechanisms that play a crucial role in generalisation processes are very likely to have a strong influence on the resulting generalisations, both in the individual and in “language”.³ Ideally, one should select psychological premises rather conservatively. An example is given in 3.3. Also after selecting any premises, one should proceed with caution and generalise from the corpus to “language” only in relatively small steps.

3.3. Similarity and preference relations

A first candidate for a very general psychological premise concerns the ability to recognise *similarities*. In fact, we make use of this ability so continuously and automatically that it seems more accurate to speak of an inability to ignore similarities. This ability is vital for human beings and all other higher organisms, for it plays a fundamental role in all learning processes. Because no two situations and contexts are ever fully identical, previous experience could never be generalised, i.e., never be utilised in future behaviour, if we were not sensitive to similarities. This holds in particular also for generalisations in language. Thus, aspects of similarity have a strong influence on any resulting generalisations in the individual and thus also on those in “language”. It is precisely this prediction that we hope to exploit in this approach.

³ One may object that this strategy partly shifts the epistemological problem from the level of “language” to the level of psychology, and this objection is fully legitimate. This is why it is important that any selected psychological premises are in principle falsifiable.

We start with a very general notion of similarity, referring to a gradual counter-concept to *identity*. Consequently, the concept is meant to comprise various aspects of similarity, ranging from basic sensory perceptions to highly abstract concepts, including analogies of all kinds, associatively connected entities and so forth. Moreover, two objects are thought to be either immediately similar by virtue of their properties, or similar more indirectly because they occur in similar contexts. The notion of context is also understood to be rather broad in scope. Note that the aspects of similarity listed above are not independent nor entirely distinct, and there are certainly also other aspects of relevance not mentioned here. However, it is an open research question which aspects of similarity are of particular importance with respect to “language”.

Any generalisations across similarity relations will be termed *preference relations* in the following. This label will become more obvious towards the end of this section. It should be stressed right away, however, that the term *preference* is not used here in its intentional reading but rather with a statistical sense. As we use a very broad notion of similarity, preference relations may in principle exist (and be uncovered) between any two objects of any kind. This includes language objects as well as language-external contexts, basic objects as well as inferred objects, concrete and abstract objects alike. When sets of objects are treated as complex objects, they may also enter preference relations. Moreover, also preference relations themselves can be treated as objects such that preference relations may even exist between preference relations. In this iterative fashion, increasingly complex structures of preference relations can be defined and uncovered.

In contrast to formal grammatical rules which may be characterised as hard categorical relations, preference relations are gradual and context-dependent relations which can, in principle, be quantified – for instance in the form of conditional probabilities. Nevertheless, the concept of hard categorical relations is, as an extreme case, included in the concept of preference relations. Grammatical rules are therefore just special types of preference relations. As for grammatical norms and their variants: In an emergentist perspective on language, there is no qualitative difference between the two notions. Both have the same status as preference relations which, however, may be preferred at different degrees. In many traditional approaches, variation is viewed as some kind of noise in usage data, overlaying the norm like some interfering signal. In an emergentist perspective, by contrast, variation is predicted and constitutes an integral constituent of “language”. Thus, norms in the traditional sense do not exist according to this perspective. One may, however, define a new notion of norm that is statistical by nature and depends on context.

In sum, apparent rules, norms and their variants can all be described as preference relations. They are special cases that stand out. Between them, however, there is a much greater number of other preference relations which are less salient and are therefore often missed or ignored by linguistic description. In simplifying words, all regularities and rule-like patterns in language consist of preference relations, and it is this pervasiveness by which the concept *preference relation* qualifies as the primary descriptive category in this research programme.

3.4. Preference relations: Examples

Preference relations are defined as generalisations across similarity relations. Different types of similarity relations give rise to different types of preference relations. For instance, *syntagmatic preference relations* may be conceived as generalisations across objects that are similar in time. A fairly straightforward operationalisation of *similarity in time* of any two objects would be in terms of *significant frequency of co-occurrence* of these objects (i.e., a high co-occurrence frequency that cannot be explained by mere chance). In this case, the resulting preference relations are potentially discontinuous, and the size (in terms of *number of objects* occurring in-between) of these discontinuities may vary. Syntagmatic preference relations may display different degrees of cohesiveness between their components, and this cohesiveness is quantified by formal cohesion measures. The following pattern provides a simple example for a syntagmatic preference relation.

- (1) sich [zu] vergewissern daß|ob|dass ... auch wirklich
“[to] assure oneself that|whether ... really”

This example contains one discontinuity (represented by the triple-dot mark) for which competent speakers will easily verify that the word distance may vary considerably across specific instantiations.

Paradigmatic preference relations constitute a second type. They may be characterised as generalisations across analogy relations between objects. It depends on the specific type of analogy relation how these relations may be operationalised. In any case, an adequate operationalisation is likely to be more elaborate in nature and more difficult to develop than the above operationalisation of similarity in time. Because a paradigmatic preference relation arises as a generalisation across objects that are analogous but typically not identical, there will also be some variation between these objects. The resulting paradigmatic preference relation will therefore comprise at least one abstract *slot* that captures this variation as a paradigmatic class of typical fillers for this slot. This class, however, is no predefined and general category (such as the word class *adjective* or the semantic feature *human*), instead it is in principle specific to this very slot in this very preference relation, much like the construction-specific categories in Croft’s *Radical Construction Grammar* (Croft, 2001).

A major challenge is to find a conceptualisation for the resulting paradigmatic preference relation that is adequate with respect to explanation. Such a conceptualisation will involve not only the underlying objects and the analogy relations between them but also a characterisation of the respective paradigmatic class(es).

3.5. Further preference relations and psychological premises

Syntagmatic and paradigmatic preference relations are already very powerful concepts because they may exist between objects of any kind. Nevertheless, other types of preference

relations may be defined by generalising across similarity relations other than analogy and similarity in time. Yet other types of generalisations are to be expected when more psychological premises are incorporated in addition to the single premise discussed so far (viz., the ability to recognise similarities).

For example, we currently explore and discuss a candidate premise according to which, in a nutshell, there are not only those cognitive mechanisms which respond to similarities in a spontaneous and associative fashion, but also other mechanisms which scan their input for apparent regularities and attempt to capture these regularities in the form of symbolic generalisations. The reality of such systematising mechanisms is obvious: Otherwise there would be no theories and no scientific modelling, and speakers would not be able to formulate ad-hoc hypotheses about grammatical regularities on the basis of their language experience. But the working hypothesis here is that humans make use of this capacity all the time, automatically and mostly without being aware of it. Whereas the spontaneous-associative processes can be attributed to general cognition, some of the systematising processes may be fine-tuned to language. The former type of processes strictly precedes the latter with respect to phylogenesis, ontogenesis and cognitive processing (in particular, cognitive language processing). Nevertheless, we postulate a tight and non-decomposable entanglement between both types of processes. Consequently, the hypothesis is that both types together contribute substantially to shaping the language behaviour of individual speakers and thereby the emergent conventions of the language community. We are therefore convinced that empirical linguistics will benefit considerably from studying the interaction between both types and their joint effects.

3.6. Local modals

In this research programme, the general procedure will be to start from observable language productions (recorded in the respective corpus) and generalise over them in an inductive and data-driven fashion such that increasingly complex and abstract preference relations can be uncovered from the data. It is important that this induction proceeds in small steps and that each step is motivated deductively by falsifiable psychological premises. The previous sections already gave some examples of this kind of stepwise induction. But how far can one get with this approach, and to what kind of model will these generalisation cascades eventually lead?

In section 2.1. it was argued that one should abandon the goal of complete language models and instead aim for local models. However, this is not necessarily a sad farewell, and it does not only apply to linguistics but also to other empirical disciplines such as physics which is concerned with similarly complex systems. This is why Paul Budnik states in a documentary on infinity:

Locality is arguably the most powerful simplifying assumption in all of physics. I suspect the physics community has been far too cavalier in abandoning this principle in favor of the elegance of their mathematical models. (Budnik 2007)

Thus, local modelling is not only a necessary constraint but, above all, a potentially valuable method for simplifying the object of investigation. A *local model* can be defined informally as a model with limited scope such that it is valid only under certain conditions. Local models may overlap, and in their overlapping areas, they may be redundant or – outside their scope of validity – inconsistent. Due to these possible inconsistencies, they cannot generally be combined to larger models. As an example, each preference relation together with its underlying objects can be described as a local model. Within the proposed research programme, the reverse statement is probably also applicable: Any local model (formulated in this programme) is a preference relation.

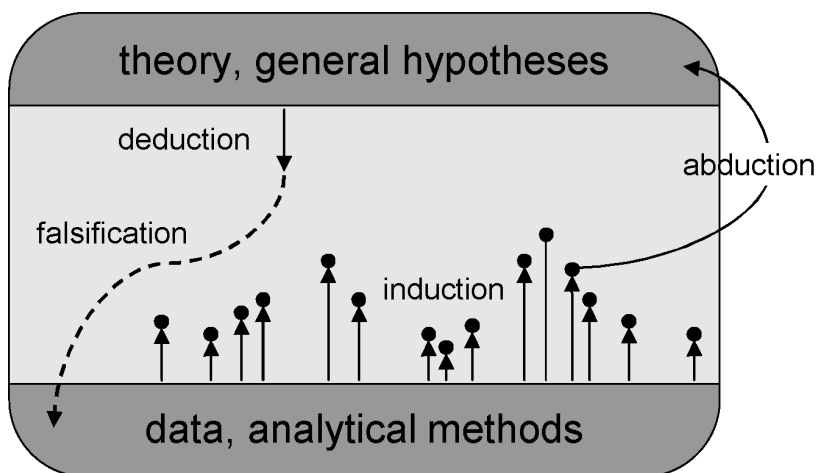


Figure 4: Local models (depicted as circles) and the explanatory gap.

Recall the explanatory gap that was discussed in section 1. Local models may be located in that gap as depicted in Fig. 4. They are the result of stepwise induction over corpus data, by means of inductive methods that are psychologically motivated. The resulting local models may be either considered the goal and primary result of this line of research, or they may be considered intermediate results which have to be explored in order to abductively arrive at new theoretical hypotheses on “language” which, of course, will in turn have to be tested empirically. It is mainly this second possibility by which the proposed programme can contribute to the development of explanatory theories.

4. Discussion

The scientific investigation of “language” – and in particular of its structural regularities – with the objective to make progress towards explanatory theories is not only a time-consuming enterprise but also a tremendous challenge, above all with respect to philosophy of science. How can one approach an object of investigation that is not directly accessible by phenomenological

experience? How can one come to insights that are supported by sufficient empirical evidence?

The programme outlined here proposes a strictly empirical strategy: Starting from an emergentist perspective on language and essentially free of any additional assumptions with respect to the object of investigation itself, it exploits the mutual influence between the “language” (in the sense of the language system and the observable language use) and the individual language users. Psychological premises are to be selected conservatively, and they must be well motivated and, at least in principle falsifiable. These premises are exploited to guide to development of specific inductive corpus-linguistic methods which are then applied in small inductive steps in order to come to similar generalisations as the language community. By being deeply rooted in psychology, the programme minimises the risk of a blind induction that any empirically driven approach is faced with. Thus, the local models resulting from this induction are sufficiently well-founded in order to abductively derive from them general hypotheses about “language”. These hypotheses are to be formulated in such a way that they can in principle be falsified. In this fashion, an explanatory theory of “language”, and particularly of grammar, may emerge successively. Since the psychological premises were made explicit, and the inductively generated local models as well as the hypotheses derived from them by abduction are driven by empirical facts, this emerging theory is empirically well-founded and falsifiable as a whole.

In sum, this programme provides a well-defined framework for the development of explanatory theories, and at the same time, its descriptive categories (viz., preference relations) are flexible enough to accommodate types of generalisations that were not anticipated at the outset. This conceptual flexibility is a key prerequisite for any data-driven approach. With these properties, the programme has good prospects to mature, in terms of Kuhn (1962), from a pre-paradigmatic stage to a full-grown research paradigm.

The programme will benefit from the fact that in the future significantly larger corpora will be available, and yet more powerful hard- and software to analyse them. Further stimulating impulses can be expected especially from future insights of experimental psycholinguistics, and we strongly hope that bringing together corpus data and experimental psychological data and analysing them from an interdisciplinary perspective will, in the long run, inspire the development of more adequate and more restricted models. The declared goal and guiding principle underlying the programme is to capture *language* empirically in an adequate way, and we are confident that in, say, thirty years, the discipline of empirical linguistics will be a good deal closer to this goal.

References

Budnik, Paul. (2007) *Mathematical Infinity and Human Destiny*, Video. Retrieved in March 2008 from <http://www.mtnmath.com/>

Bybee, Joan. (1998) “The emergent lexicon”, *Chicago Linguistics Society 34*: The panels, 421-435.

Bybee, Joan / Hopper, Paul (eds). (2001) *Frequency and the emergence of linguistic structure*,

Amsterdam.

- Chomsky, Noam. (1965) *Aspects of the Theory of Syntax*, Cambridge, Mass.
- Croft, William. (2001) *Radical Construction Grammar: Syntactic theory in typological perspective*, Oxford.
- Elman, Jeffrey L. / Bates, Elizabeth A. / Johnson, Mark H. / Karmiloff-Smith, Annette / Parisi, Domenico / Plunkett, Kim. (1996) *Rethinking Innateness: A Connectionist Perspective on Development*, Cambridge, Mass.
- Goldberg, Adele. (2006) *Constructions at work: The nature of generalization in language*. Oxford.
- Hoey, M. (2005) *Lexical Priming: A New Theory of Words and Language*, London.
- Hopper, Paul. (1987) "Emergent Grammar", *Berkeley Linguistic Society* 13, 139-157.
- Hopper, Paul J. (1998) "Emergent Grammar", *The New Psychology of Language: Cognitive and functional approaches to language structure*, Tomasello, Michael (ed.), Mahwah, New Jersey, 155-175.
- Keller, Rudi. (2006) "Pfade des Bedeutungswandels", *Von Intentionalität zur Bedeutung konventionalisierter Zeichen*, Proost, Kristel / Winkler, Edeltraud (eds.), Festschrift für Gisela Harras zum 65. Geburtstag. Tübingen, 337-358.
- Kuhn, Thomas S. (1962) *The Structure of Scientific Revolutions*, Chicago.
- Langacker, Ronald W. (1987) *Foundations of cognitive grammar* (Vol. 1: Theoretical Prerequisites), Stanford.
- Langacker, Ronald W. (2000) "A dynamic usage-based model", *Usage-based models of language*, Barlow, Michael / Kemmer, Suzanne (eds), Stanford, 1-63.
- Sinclair, John M. (1991) *Corpus, concordance, collocation*, Oxford.
- Stetter, Christian. (1999) *Schrift und Sprache*, Frankfurt a.M.
- Tomasello. (2003) *Constructing a language: A usage-based theory of language acquisition*, Cambridge, Mass.