*Peter Harder**

If Your Language Was a Car…
The Object(s) of Linguistic Research, or: Towards a Shared Geography of Linguistics

Abstract
The article suggests that there are underexplored possibilities for fruitful communication between formal and functionalist linguistics. A key issue is the question of exactly what each approach is aiming to capture about language. This is especially relevant for understanding the status of claims about autonomy. The role of distributional regularities and their precise relations with semantic motivation is argued to be a shared problem that could fruitfully be addressed from both sides of the divide – and the role of niche construction as a dimension of evolutionary theory is put forward as providing a new take on the innateness debate. Torben Thrane's work is discussed as an illustration example.

1. Introduction
This article is about the conflicting perspectives on linguistics taken by the research communities that are generally known as 'formal' and 'functional' linguistics – as represented by (respectively) Torben Thrane and myself. The aim is not to confront them, but to stand back in order to propose a meta-perspective that might make a constructive discussion between them easier. The background for this proposal is my recent work (Harder 2010) on language as a social phenomenon. As a spinoff from trying to get clearer about the nature of language-in-society as an object of description, some conclusions about the nature of language-in-mind have offered themselves, some of which may present new ways of addressing familiar bones of contention. The nature of this aim means that the article concludes in questions (albeit leading questions!) rather than assertions: If we adopt this approach, perhaps it will be easier to throw light on the same issue from our different perspectives?

2. Aims and objects of description
It is striking how many separate sub-issues Torben Thrane and I have been in agreement about over the years. They include the need for a clear distinction between statements about language use and about language structure; the instructional nature of linguistic meaning; the claim that linguistic representations stand on the shoulders of primary, pre-linguistic representations – and that structures based on functor-operator relations are central to the way linguistic meanings collaborate.

The reason why these partial agreements have not combined into a broader and more coherent agreement is to do with the nature of the object(s) of description that are associated with our overall approach. Differences between presupposed objects of description are obvious in the case of hyphenated disciplines such as sociolinguistics and psycholinguistics, but they exist also between people who simply understand each other as 'linguists'.

* *Peter Harder*
University of Copenhagen
Department of English, Germanic and Romance Studies
Njalsgade 128, 24.2
DK-2300 Copenhagen S
harder@hum.ku.dk

If we view the issue historically, the core object of description has changed several times. Classic traditional grammar arose in a universe centred on culturally focal written texts, and maintained the relationship as part of a broadly philological conception of the humanities, until the structualist revolution divided the paths. But the core object of description differed also between the different structuralists. Thus American structuralism was formed by the need to describe languages without a written tradition, and stayed close to the data – while European structuralism was inspired by rather theoretical notions and worked with the languages that were already familiar from traditional grammar. Generative grammar, while structural in orientation, was part of the cognitive revolution where the idea of formal simulation was central, and changed the focus to formal mechanisms assumed to underlie the distinction between grammatical and ungrammatical sentences. Later, semantics has come to play an increasing role within this cognitive generation of linguists, while maintaining 'language in the mind' as the core object. My own preoccupation has been to take the step to a societal perspective on language, which includes but is not limited to cognition.

Torben Thrane and I both became linguists during the generative pioneer age. The reason we ended up adopting different perspectives was not a propensity to adopt a particular faith as 'true believers', but simply that we gravitated towards positions in the linguistic landscape with different types of focal features (with different attractions). For that reason, I believe the differences of perspective may be symptomatic for other language persons.

Generative grammar had many attractive features as part of the cognitive revolution. The idea that there had to be underlying mental structures in order to make language use possible appeared uncontrovertible, once the black box had been declared open. The computer metaphor also appeared both intellectually and practically promising as a format for offering theoretical explanations that were both complex and falsifiable, compared with the traditional practice of offering competing verbal formulations with uncertain falsifiability criteria. Chomsky also offered linguistics a position as part of the vanguard, along with Marr's theory of vision, in the new drive to close the last frontier for science: the human mind. There was much to be fascinated with.

One point that was decisive for my own orientation was the issue of how far you can get by pursuing a descriptive strategy that takes language 'in reality' to be of the same kind as language in the formal mechanism – especially as a way to account for meaning. Formal simulation as a descriptive strategy tones the description in the direction of ever-subtler attunement to distributional distinctions, allowing a gradually more abstract hierarchical representation which in principle could end up with, cf. Pinker/Jackendoff (2005), up to seventeen abstract nodes in the clause representation. While empirical falsification was a strong point precisely in relation to distribution, however, the relation with meaning and understanding was not equally safeguarded – and that dimension was more appealing from my point of view.

Even so, the phenomena that trigger descriptive choices are if not the same, then to a considerable extent overlapping. Moreover, the two sides can gain from being pursued in tandem. One thing that painstaking distributional analysís can reveal (cf. Thrane 2003b on 'degree phrases') is that in a given case there may be no connection that can be attributed to the distribution-based generative hierarchy (in contrast to what previous models assumed): a combination of inherent semantic features and an instructional semantics matches the facts much better that abstract distributional categories.

This is an example of the fact that functionalists can also benefit from the discipline that is enforced by a formal model that generates potential counterexamples. Most functionalist models overgenerate without triggering much concern among the authors (as pointed out by Russell Tomlin on many occasions) – and this is a problem that needs addressing. Thrane (2003a) on Danish *sikke* is a contribution to the careful mapping out of differences in terms of distributional classes in English and Danish with obvious potential links to a function-based account. Let me mention in passing also Klinge's dissertation (2005) on NP-structure which shows a difference between Danish and English that is only visible if you are more careful about distributional classes than most functionalists. The question is how to conceive of an object of description that makes a collaborative approach more accessible.

3. Autonomy and context

The key issue is probably the autonomy issue and its reverse side, the relations between language and context. Whenever that debate flares up, a pattern of mutual misunderstanding reasserts itself. Roughly speaking, functionalists believe that generativistst regard language as pure disconnected structure, while generativists believe that functionalists try to derive structure directly from function – and both parties vehemently deny these absurd allegations.

Among aspects of the complicated truth are that both parties operate with structural as well as semantic properties, and both parties would like to make them fit as tightly as possible. The crucial difference is associated with the point of departure from which these issues are approached. One side starts out with ‘language as such’ and postpones the issue of how it functions, arguing basically from an analogy that just as it does not make sense to ask about the function of an organ such as a thumb or spleen before you have described that organ, so it does not make sense to inquire about the function of language except on the basis of a description of language (as such). The other side starts out with language-in-function and postpones the issue of structural properties, arguing basically that what really exists is linguistic interaction, and unless structure is described based on that, it is hypostatized and ontologically dubious. Basically, the two parties think each others’s project is doomed to fail.

Although Thrane basically aligns himself with the first option, his approach includes ’canonical’ functions, cf. Thrane (2009). While I basically take the second approach, I attach great importance to the distinction between structure and usage based on a notion of autonomy, albeit one based on a distinction between ‘partial’ or ‘weak’ autonomy on the one hand and ‘absolute’ or ‘strong’ autonomy on the other. Strong autonomy means that a domain follows only its own laws and is sealed off from outside influences. Weak autonomy means that a domain on the one hand shares properties and factors of influence with surrounding domains, but in addition it has properties that are special to the domain itself. The two senses are easy to conflate because the typical case is that autonomy is weak but the focus is so strongly on the autonomous properties that the others are forgotten.

If these two concepts are confused – and in my experience this is what typically happens in linguistic discussions – the result is an acrimonious but almost entirely meaningless discussion. In point of fact, everybody agrees that language is weakly autonomous in the sense that linguistic properties cannot be predicted exhaustively from extralinguistic properties. In contrast, I have yet to hear an argument explicitly (and compellingly) defending absolute autonomy. Even the most abstract generalizations, including the most promising candidates for universal status, can be understood as (partially) associated with functional features. Seeing them as ‘pure structure’ is a preferred stance, not an empirical necessity. What takes the place of an argument for (and against) absolute autonomy is usually an argument about particular linguistic categories, such as ‘subject’ or ‘gender’, where one side is trying to maximize motivation and the other side is trying to minimize it – or, from the other end of the spectrum, one side is trying to minimize the role of structural categories while the other side is trying to maximize it. The point where understanding inevitably fails is where generativists take the step from asserting the existence of structural properties to attributing them to a special compartment in the theory, while functionalists take the step of placing structural properties in relation to functional properties.

My point in this article is that this is essentially a matter of descriptive preference and cannot be settled by the evidence. Just as you can choose to go for ‘pure structure’, you can choose to start at the structural end – or at the functional end. What can (at least in principle) be settled by empirical evidence is the question of motivation: How strong is the correlation between functional potential and structural categorization, for each of the structural categories we are interested in? The debate on autonomous structure takes other forms as well, but this is well suited as an example of how the unproductive polarization can be avoided: to look for the aspects of the problem that can be linked with empirical predictions and leave the rest to internal meetings between true believers. As pointed out by Croft (1995) in an argument against extreme reliance on functional motivation, there is no way to avoid arbitrariness – on the other hand, there is also no way to avoid motivation (why else would speakers sometimes prefer one way of speaking to another?).

My own general strategy is to view the task of describing language as such as inherently bound up with the description of language in function: Properties specific to the domain of language can only be identified as part of a process where the other properties are also in play. Roughly speaking, the decision to treat a linguistic category as wholly and absolutely autonomous of all external factors could only be taken once you had actually tested out whether there was a significant correlation with any of the likely candidates for motivating factors. And for all partially autonomous cases, a precise description is one that captures both the correlations between the structural categories and motivating factors and also describes cases where motivation fails to predict the linguistic facts.

This approach is conceived within an evolutionary perspective (see also the next section). Evolution, however, involves both a set of biological processes and a set of cultural processes. While generative grammar has been most concerned with the biological dimension (because of the innateness issue), there must have been cultural evolution also: an ’African Eve’ (= the first individual to carry the language gene) cannot have opened conversation by saying ’whose banana did Joe steal [gap]? Linguistic communication must have evolved on the basis of presupposed, non-linguistic resources (cognitive and interactive), which therefore have a share in determining what a grammatical sentence comes to look like.

This does not rule out a choice of focus on the fully formed, stable, linguistíc infrastructure in the individual, as the prioritized object of description – thus continuing the generative agenda. However, it means that this interesting object came with presupposed sociocultural foundations, and the larger object of description that includes those foundations is part of the general scientific endeavour. This is where I see (although possibly we will have to wait until Chomsky retires definitively!) a possibility of converting a religious war into a division of labour. It is clearly one of the tasks of linguistics to develop a model of the infrastructure that underlies the ability to produce the kinds of sentences that enter into an individual’s language. Equally clearly, it is a task of linguistics to describe the kinds of sentences that are found in actual linguistic interaction in a given community. The central problem is probably going to be to get people to see what kind of sense it makes to start from the other (~’wrong’) end.

4. E-language and evolutionary theory: populations, niches and culture

From the formal or generative side, a major stumbling block will be the question of what precisely E-language is. It is by no means obvious how to solve the problem of taming E-language as a manageable object of description – all types of functionalists have to produce their own types of abstractions in order to get a grip on the raw and multifarious sprawl of linguistic communication. The point in this connection is that even from a generative point of view, it needs to be done. The reason is the scenario I argued for above: without a cultural component, the genetic component would not be able to give rise to sentences of the kind that constitute the input to linguistic description. If we ignore that part of it, looking only for properties that can be put down to genetically shaped infrastructure, linguistics will be in the position of looking for the key under the streetlight because that is the only place where there is (what we recognize as) light.

Cultural evolution has the same type of causal dynamics as biological evolution – but not the same mechanisms. All evolutionary dynamics depends on two levels: individual and population (cf. Croft’s evolutionary theory of language change (2000), based on Hull (1988). Reproduction (of utterances and DNA) belong at the individual level, while selection and proliferation are population-level phenomena – which in a linguistic context means a social, community-level phenomenon. To the extent properties of individuals reflect selection pressures, this is a social, population-level feature, not a purely individual-level feature. Analogously, wearing the fashionable colours of the season is not the result purely of internal wiring.

This is bound up with the process whereby linguistic expressions gain 'currency' (another social phenomenon). The comparison with fashion may suggest that this is merely a surface phenomenon – but on reflection, it will be clear that a language does not exist in a community unless it has 'currency'. Hence, according to a functional approach, having community-wide 'currency' is not an external, but possibly the most foundational feature of language. The existence of a linguistic feature in the community therefore has to be understood in relation to the factors that contribute to making it 'current' – i.e. the process of evolution, including cultural evolution. This observation is related to Dobzhansky's claim that nothing in biology makes sense except in the light of evolution. Neither biological organs nor linguistic expressions can be understood except as constituents of a larger whole that includes individuals and populations existing in a dynamic equilibrium. It is still okay to describe a single animal or organ – but this is a choice you make, which requires you to carve it out from a context that is not in a foundational sense *external* to the object you are describing.

This entails that biological species and expressions are 'lineages', cf. Croft (2000). They exist in a panchronic space, from which both synchronic and diachronic descriptions are abstractions. Such lineages (like the horse lineage from *eohippus* to *equus*) are partly the results of selection pressures, partly also the *sources* of selection pressures – prey and raptor species 'co-evolve'. Lineages therefore influence each other – and thus they also influence also the world they live in.

This last process is called 'niche construction'. A classic example is the beaver: beavers are adapted to a landscape with ponds – built by beavers, cf. Deacon (1997). Lactose tolerance is an example from human evolution: in a culture that keeps cattle, there is a selective advantage in being able to digest milk But because human beings also adapt to other minds, evolutionary dynamics can run on other things than DNA. Language change is a form of cultural evolution, like changes in institutional practices (such as the legal system): at any given time there is an evolving cultural niche to which new members adapt (even while they change it).

The combination of niche construction and cultural evolution provides a scientifically respectable home for language as something that exists in the individual's environment rather than as something that exists only inside the individual. In my terminology, the language system (= 'langue') is part of the *niche* – while *competency* (please note the '-y'!) is what arises in the individual's mind as a result of the process of adaptation. There is of course an innate point of departure for this process, but I am going to leave that issue aside here. What matters for the purposes of this article is the internalized competency that represents a finished state of adaptation in the individual. As will be evident in the subsequent section, I believe this corresponds to I-language as described by Torben Thrane.

In order to get a complete picture of linguistic facts, we therefore have to distinguish several related, but different objects of description – all of which are part of the whole evolutionary process. Niche elements function as 'affordances' – as things that can be understood by the individual (if he has the key!), and which yield a selective advantage to those who do. 'Competency'-elements *constitute* the key: they function as tools for the individual, and in combination they constitute a quasi-organ.

The larger landscape that I have tried to depict in this section thus includes objects of description inside individuals (competencies), and it also comprises a cultural niche (including other people's linguistic interaction) to which individuals adapt. I am going to argue that this provides a framework that allows for the empirical pursuit of both generative and functionalt interests – provided some adjustments are made in the generative picture. (Otherwise we'll just have to continue the quarrel!)

5. Towards a shared geography of linguistics?

Torben Thrane's object of description is basically the Chomskyan I-language, but it includes a strong element of functionality – pure formal structure is not enough. This is why I think his linguistic project would be able to accommodate to the linguistic landscape I have tried to outline. As illustrations, I have selected some passages dealing with core theoretical issues:

> Skal man nu give en internt-systematisk beskrivelse af, og forklaring på, bilens indretning ud fra eksterne, kommuter-sociale kriterier? Eller skal man opfatte dens strukturelle indretning ud fra hvad der er teknisk og naturvidenskabeligt mulige løsninger på?(...) det er principielt muligt at forklare biler både ud fra det kommuter-sociale synspunkt og ud fra det teknisk-naturvidenskabelige. Men jeg vil som min personlige holdning gå ud fra at den sidste er den fornuftigste hvis man vil sige noget om bilens strukturelle indretning. (Thrane 2008: 140)
>
> (My translation: Should we now offer an internally systematic description of, and explanation for, the construction of the car based on external 'commuter-social' criteria? Or should we try to understand its structural setup based on what is technically and scientifically feasible? (…) it is possible in principle to explain cars both from the commuter-social point of view and the technical-scientific point of view. But as my personal preference I assume that that last option is the most sensible one if you want to say something about the car's structural construction) (Thrane 2008: 140)
>
> *Sprogevnens initialtilstand er genetisk bestemt og definerer under navnet 'Universal Grammar' det sproglige fællesgods for mennesker. I løbet af individets opvækst, og under miljømæssig påvirkning, gennemløber sprogevnen en række tilstandsændringer for at munde ud i en såkaldt stabil tilstand. Stabile tilstande er I-sprog. I-sprog er med andre ord det der ligger bag et individs evne til at producere og forstå sproglige udtryk. (Thrane 2004)*
>
> (My translation: The initial point of the language acquisition process is genetically determined and (under the name of 'Universal Grammar') defines the shared linguistic equipment of human beings. During the maturation of the individual, and under impact from the environment, the language ability undergoes a series of changes, eventuating in a so-called stable state. Stable states are I-languages. I-languages, in other words, are what underlies an individual's ability to produce and understand linguistic expressions.) (Thrane 2004).

In [Chomsky's] programme, HLF [= the human language faculty] is assumed to be neutral with respect to production and reception of utterances, on the grounds that it would be 'unnatural' to assume that we were endowed with two radically distinct faculties, one for production and one for reception and understanding. This may appear to be a reasonable view, but the fact remains that production and reception are radically distinct processes. It is therefore only a reasonable view at a level of abstraction at which HLF is regarded as an insulated system, closed off relative to other cognitive systems, a system that can be shown to be relevantly engaged in both production and reception, but which at the same time can be shown to remain immune from those factors that make production and reception different processes, a sort of "brain in a vat".

> *This view is* a priori*. Rejecting it, or rather suspending judgement on it to be confirmed or falsified by empirical evidence, this paper will discuss some of the key issues in Chomsky's programme with a view to determining whether it is equipped to answer what I take to be a more legitimate* a priori *question:* How can we explain the manifest human ability to <u>understand</u> linguistic utterances as vehicles for the expression of thoughts and ideas? (Thrane 2009: 2)

The stable mind-internal end state described in the last quote above appears to me to be equivalent to what I call 'competency'. For reasons described above, this object of description has many of those properties that are sometimes mistakenly assumed to go with strong autonomy. Distributional regularities are a key part of the structural complexity of language in both its I and E manifestations; individual 'competency' has an innate foundation, but in its end state it reflects adaptation to linguistic 'affordances' in the community – as also pointed out by Thrane as quoted above.

I discreetly add that a complete theory thus requires a description of the language that exists outside the individual and exerts selection pressure on acquisition, and the distinction between grammatical and ungrammatical utterances thus cannot be captured merely by reference to mind-internal constructs – although you can choose to *focus* on internal aspects.

From this point of view, functionalists and formalists have a shared duty to figure out how to clear up the mixture of isomorphism and divergence between semantic and distributional properties. This requires both parties to resist the temptation to go for complete isomorphism based on their own favourite view of what is ‘basic’.

6. Final issues and questions

Although my orientation is strongly towards a shared wider object of description that can accommodate core concerns of both generativists and functionalists, it would be misleading to suggest that one can avoid theoretical confrontation in some points. The two major issues, as far as I can see, involve the way linguistic infrastructure is conceived. I hope to drive a wedge between Chomsky and Torben Thrane, leaving the latter and myself on the same side of the fence.

First, there is the question of the level of concretion of the object of description that constitutes the infrastructure. Once more, the issue is linked to the status of ‘pure structure’ as a preferred stance rather than as part of the ontology, i.e. part of the assumed object of description. If my language is a car (cf. the first quote above), it would appear to be necessary to see it as including both the reception and the production apparatus (analogous to forward and reverse gears). A neutral description may be a preferred stance, but it cannot be part of the object of decription that is given before the linguist starts his work (cf. extract 3). Differently put, if language is an organ, it must constitute whatever is necesssary to deliver the goods, not just something that can in a technical sense generate ‘all-and-only the grammatical sentences of the language’. (Just as the liver, with Newmeyer’s analogy (1998), has to be described in terms of its actual capacity to process input from the organism, not just in terms of structural abstractions). Similarly, the I-language/competency that linguists describe is the fully adapted ability to speak as a member of the community, not just the abstract general features.

If that is so, it also appears to me that what one describes is a ‘can-do capacity’ – not a ‘tacit knowledge’ that is claimed to differ from ability (cf. Chomsky 1980: 59). There is a type of process that is generally recognized to work in a manner that is sealed off from intentions and mental content, namely procedural ability, and it is plausible to assume that a fully developed and stable language ability has a strong procedural component. In that sense it is also analogous to the walking ability – an illustration that has been invoked to illustrate the status of the innate language ability. But the level of purely structural abstractions (that constitute tacit knowledge but not ability) again appears to be ontologically epiphenomenal.

What this article asks, therefore, is whether this is not compatible with a pursuit of the generative agenda in linguistics – in which case I think collaboration is feasible. Although these questions should perhaps not be raised as the first item on the agenda when a shared geography is to be hammered out, I think answering them will be necessary at some point in the process – otherwise the polarization will return!

References

Chomsky, Noam 1980: *Rules and Representations*. Oxford: Blackwell.

Croft, William 1995: Autonomy and functionalist linguistics. In *Language* 71, 490-532

Croft, William 2000: *Explaining Language Change. An Evolutionary Approach.* London: Longman.

Deacon, Terrence 1997: *The Symbolic Species. The co-evolution of language and the human brain.* Harmondsworth: Penguin.

Harder, Peter 2010: *Meaning in Mind and Society.* Berlin: Mouton de Gruyter.

Hull, David 1988: *Science as a process: an evolutionary account of the social and conceptual development of science.* Chicago: University of Chicago Press.

Klinge, Alex 2005: *The Structure of English Nominals. A Study of Correlations between Form, Distribution and Function*. Copenhagen: Copenhagen Business School.

Newmeyer, Frederick 1998: *Language Form and Language Function*. Cambridge, MA: MIT Press.

Pinker, Steven/Jackendoff, Ray 2005: The Faculty of Language: What's Special about it? In *Cognition* 95, 201-236.

Thrane, Torben 1999: Forklaringstyper i lingvistikken. In *Funktionelle Fodnoter: Formel sprogteori: Nogle udfordringer til den funktionelle tilgang,* 66-81.

Thrane, Torben 2003a: Breakstructures. In *Journal of Linguistics*, 39(2), 327-372.

Thrane, Torben 2003b: Gradsudtryk. Konferencen: 9. møde om udforskningen af dansk sprog, 10. oktober 2002–11. oktober 2002, 285-297. Aarhus: Aarhus Universitet.

Thrane, Torben 2004: Hvorfor er sproget så svært at forstå når det er så let at forstå? In *Tidsskrift for Sprogforskning* 2, 37-70.

Thrane, Torben 2008: Hvis funktionalisme er svaret – hvad er så spørgsmålet? In *NyS* 36, 120-150.

Thrane, Torben 2009: The Language Faculty – Mind or Brain? In *Tidsskrift for Sprogforskning* 7, 1-42.