The emergentist program

William O'Grady

Department of Linguistics, University of Hawai'i, Honolulu, HI 96822, USA

Abstract

This paper presents an overview of the diverse body of thought known as 'emergentism', as it is manifested in the study of language. It is suggested that a unifying theme of emergentist work is the rejection of inborn grammatical principles (Universal Grammar) in favor of more general mechanisms of cognition and learning. A variety of emergentist proposals are outlined and the general thrust of this work is compared to ideas that have recently been put forward within the Minimalist Program.
© 2007 Elsevier B.V. All rights reserved.

1. Introduction

In recent years, a growing number of proposals about language have been put forward under the rubric of 'emergentism'. Although it is widely understood that emergentist approaches stand in opposition to nativist theories of development, especially those associated with Universal Grammar, there is perhaps less of a consensus as to exactly what emergentism itself claims about the workings of language.

Answering this question is made difficult by the fact that the term 'emergent' is used in different, sometimes contradictory ways that vary from discipline to discipline and from writer to writer. As Stephan (1997:305–306) observes, 'an immense variety of emergentist doctrines' can be identified, and the resulting theories have 'many faces'.

I will begin by outlining the fundamental tenets of emergentism as it is practiced in linguistics (section 2). Sections 3 and 4 briefly survey a number of emergentist proposals for grammatical analysis, drawing on ideas that have been put forward in the literature. Section 5 considers emergentist treatments of language acquisition. After some brief concluding remarks,

E-mail address: ogrady@hawaii.edu.

Appendix A considers a possible relationship between emergentism and recent versions of the Minimalist Program.

2. The emergentist tradition

The roots of emergentism can be traced to the work of John Stuart Mill (1930 [1843]), who proposed (Book III, Ch. 6) that a system can have properties that amount to more than the sum of its parts. As Mill observes, obvious examples of this come from chemistry.

> The chemical combination of two substances produces, as is well known, a third substance with properties different from those of either of the two substances separately, or both of them taken together. Not a trace of the properties of hydrogen or oxygen is observable in those of their compound, water.

Mill's insight is also relevant to the study of so-called 'Complex Systems' – ranging from atoms to flocks of birds in flight to the weather – whose dynamic non-linear behavior involves many interacting and interconnected parts.[1] The following two examples are frequently mentioned for their illustrative value. (The first is from Ellis, 1998:643 and the second from Bates et al., 1998:590; see also Bates and Goodman, 1999:32. MacWhinney, 2002 offers additional examples.)

- The queues at a red light on a multi-lane street are roughly equal in length. This cannot be explained by reference to the street, the cars, traffic laws, or the desire of drivers to create symmetrical queues. Rather it follows from the *interaction* of drivers who share the goal of getting to wherever they are going as quickly as possible.
- The hexagonal shape of the cells in a beehive does not follow from the properties of wax, of honey, or the geometric preferences of individual bees. Rather, hexagons arise from the practice of packing circles together in a way that minimizes wasted space.

What, though, would an emergentist theory of *language* look like?

As best I can determine, proponents of emergentism within linguistics are in general committed to the following thesis.

(1)   *The emergentist thesis for language*
          The phenomena of language are best explained by reference to more basic
          non-linguistic (i.e., 'non-grammatical') factors and their interaction—physiology,
          perception, processing, working memory, pragmatics, social interaction, properties
          of the input, the learning mechanisms, and so on.

Three additional points may help clarify the orientation and goals of contemporary emergentist work on language.

2.1. Properties versus principles

Although linguistic emergentism denies the existence of certain types of grammatical *principles*, it does not deny the existence of grammatical *properties*—some a matter of

---

[1] A system is dynamic if it is subject to constant flux; it is non-linear if effects are disproportionate to causes (as when the movement of a skier causes an avalanche).

stipulation (perhaps adjective-noun order, as observed by Dryer, 1992:95–96) and some the result of deep-seated principles (various aspects of agreement and binding, for instance). The issue has to do with how these and countless other linguistic phenomena should be analyzed and whether the core properties of language (whatever these turn out to be) can be shown to follow from the operation and interaction of non-grammatical factors.

2.2. Connectionist architecture

A good deal of emergentist work within linguistics adopts the techniques of 'connectionism', an approach to the study of the mind that seeks to model learning and cognition in terms of networks of neuron-like units whose relationship to each other is typically graded and probabalistic (e.g., Elman, 1999; Christiansen and Chater, 2001; Palmer-Brown et al., 2002).

Some varieties of connectionism reject the existence of the sorts of symbolic representations (including syntactic structure) that have played a central role in explanatory work on human language. (For a critique of this sort of 'eliminativist' program, see Marcus, 1998, 2001; Gregg, 2003. Smolensky, 1999; Steedman, 1999 discuss ways to reconcile symbolic approaches to language with connectionism.) However, the eliminativist position is far from universally accepted within emergentism, and symbolic representations of one form or another are evident in the work of many emergentists (e.g., Goldberg, 1999; Tomasello, 2003; O'Grady, 2001a, 2005).[2]

Regardless of one's views on the symbolist/eliminativist controversy, one thing is clear—connectionist modeling provides a useful way to test various predictions about language acquisition, processing, change, and evolution (see MacWhinney, 2002; Elman, 2002, 2005b for discussion). I will return to this issue in my discussion of language acquisition in section 5.

2.3. Emergentism and nativism

Emergentism is not opposed to nativism per se—the fact that the brain is innately structured in various ways is simply not a matter of dispute. However, emergentists are unanimous in their rejection of the idea that there are innate *linguistic* constraints on the computational system for language—a central tenet of grammatical nativism (i.e., Universal Grammar).

A somewhat less clearcut issue involves the status of the 'building blocks' of language—the features, categories, and contrasts that are manipulated and encoded in the course of forming and interpreting sentences. According to the widely held nativist view, at least some of these elements are linguistic in character, including the vocabulary of notions and relations that are vital to the formulation of grammatical principles (N, V, c-command, government, barrier, phase, and so forth).

In contrast, the emergentist thesis (see (1) above) takes as its starting point the idea that the linguistic system is constructed from non-linguistic parts. An appealing tag line for this idea comes from Bates and MacWhinney (1988:147): language, they say, is a 'new machine built out of old parts'. Unfortunately, it is not always clear what qualifies as an 'old part'. Obviously, a notion such as 'governing category' cannot qualify—it is without a doubt inherently grammatical

---

[2] O'Grady uses tree-like representations, while Goldberg employs the notation associated with Construction Grammar. As noted by O'Grady (2005:13n), it is worthwhile keeping open the possibility that representations might be 'symbolic approximations' in the sense of Smolensky (1999:594)—that is, abstract, higher-level descriptions that approximate the patterns of neuronal activation that connectionist approaches seek to model.

in character. But what about 'definiteness'? Or telicity? Or plurality? I know of no satisfactory discussion of this matter in either the UG or emergentist literature.

3. Emergentist approaches to grammatical phenomena

Although most work on linguistic emergentism focuses on language acquisition (see section 5), some research has addressed the traditional concerns of grammatical analysis.

The earliest work of this type, succinctly reviewed by McClelland and Patterson (2002), called into question the existence of morphological rules and representations, even for phenomena such as regular past tense inflection. Instead, it was suggested, a connectionist network learns the relationship between the phonological forms of stems and those of past tense forms (*run~ran*, *walk~walked*, etc.), gradually establishing associations ('connections') of different strengths and levels of generality between the two sets of elements. On this view, still highly controversial (see, e.g., Pinker, 1999), the so-called 'regular past tense' simply reflects the most general and strongest of these associations.

Other work has called into question the existence of morphemes as categorical units, proposing that word structure is gradient and sensitive to statistical regularities in the form-meaning relationships between words. Hay (2003) suggests that the relative frequency of a word compared to its base is particularly important in this regard: morpheme-like boundaries are most likely to emerge when the base is more frequent than the word containing it. Thus, native speakers take *inadequate* (whose base is more frequent than the whole word) to be structurally more complex than *inaudible* (which is more frequent than its base). Hay and Baayen (2005) present a overview of emergentist work on morphology.

There have also been attempts to develop an emergentist approach to the lexicon. For instance, Bybee (1998) has proposed that the lexicon emerges from the way in which the brain responds to and stores linguistic experience—by creating units whose strength and productivity is determined largely by frequency of occurrence. Some of these units correspond to words, but many are phrases and other larger units of organization, including abstract constructions.

Elman (2005a) presents a more radical view, proposing that lexical knowledge is implicit in the effects that words have on the mind's internal states, as represented in the activation patterns created by a simple recurrent network (see section 5).[3] A word's meaning (and its syntactic category too), he suggests, emerge from the contexts in which it is used rather than from an a priori vocabulary of linguistic primitives.

Finally, it is also possible to identify several strands of emergentist work on syntax. Differing views have been put forward by MacWhinney (2005) and O'Grady (2001a, 2005), both of whom address a series of phenomena that lie at the heart of contemporary syntactic analysis—the design of phrase structure, co-reference, agreement, the phonology-syntax interface, and constraints on long-distance dependencies. Whereas MacWhinney seeks to explain these phenomena in terms of pragmatics, O'Grady argues that they are best understood in terms of the operation of a linear, efficiency-driven processor.

Yet another idea, pursued in work on Construction Grammar, seeks to reduce syntax to stored pairings of form and function ('constructions'). Some of this work has a strong emergentist orientation (e.g., Goldberg, 1999; Tomasello, 2003), but its ultimate status is unclear. Indeed, two

---

[3] An SRN is a neural network architecture that processes patterns of sequentially ordered elements, producing outputs that are a function of both the current output and the SRN's own prior internal state. SRNs are especially good at noticing local co-occurrence relationships—given that we have the word X, what's the likelihood that the next word will be Y?

leading proponents of Construction Grammar have expressed disagreement (in a co-authored paper, in fact) over whether their framework includes a commitment to nativism (Goldberg and Jackendoff, 2004, 563).

In the interests of offering a more detailed example of what an emergentist analysis of a syntactic phenomenon looks like, I will briefly summarize a proposal that I have made for agreement in English. This analysis is not necessarily typical of emergentist work on syntax in general (the field is in fact highly diverse), but I believe that it does help illustrate something of the spirit and potential of the emergentist approach.

4. Agreement in English—a case study

In its simplest instantiation, the English system of agreement requires that a verb agree in person and number with its subject argument. Hence (2) is acceptable, but (3) is not. (I use Roman numerals and upper case for nominal features, and arabic numerals and lower case for verbal features.)

(2)   Third person singular subject, third person singular verb form:
        Milk is on the table.
        IIISG 3sg

(3)   Third person plural subject, third person singular verb form:
        ***Cookies** is on the table.
         IIIPL   3sg

The descriptive generalization can be summarized as follows.

(4)   Verbs agree with their subjects.

From an emergentist perspective, of course, this sort of generalization is problematic for a variety of reasons. For one thing, it is hard to know what to make of the notion 'subject', which appears to be a domain-specific grammatical notion (although see Morris et al., 2000). And, just as importantly, it is unclear why English should manifest subject-verb agreement rather than, say, verb-object agreement. Certainly, we could *stipulate* that it has subject-verb agreement, as grammatical analyses almost invariably do, but that is not the same thing as deriving the phenomenon from properties of its parts and their interaction with each other.

Interestingly, the generalization in (4) encounters a second difficulty—it is descriptively inaccurate. As (5) shows, there are patterns in which the verb agrees with a nominal other than its subject.

(5)   a.   There is   **milk** on the table.
               3sg IIISG

       b.   There are **cookies** on the table.
               3Pl   IIIPL

In fact, there are even patterns in which the verb agrees with just the first conjunct of a non-subject nominal, as Sobin (1997:324) documents.

(6)   There is [**milk** and water] on the table.
                3sg IIISɢ

How can facts such as these be accommodated by an emergentist approach? The analysis that I have put forward (O'Grady, 2005:90ff) presupposes a sentence-building system that consists of two 'parts'.

4.1. The sentence-building system

The first part is a lexicon that provides a list of 'formatives' (words and morphemes) and their associated properties. For the purposes of illustration, I assume a very conventional sort of lexicon, without taking a position on whether the information it contains might be derived in the manner suggested by Bybee and/or Elman (see above). Thus, as illustrated in (7), I assume that *exist* is a verb that has a single argument dependency (requirement), in the form of a nominal to the left. In contrast, *ignore* manifests two such dependencies, one involving a nominal to the left and other involving a nominal to the right.

(7)   a.   *exist*: V, <N>   (e.g., *Problems exist.*)
                          ←

       b.   *ignore*: V, <N N>   (e.g., *Mary ignored Sam.*)
                          ← →

Inflected forms such as *exists*, *ignores*, and so forth carry additional information—they introduce a third person singular feature that must be matched at some point with a similar feature on some other element in the sentence. (Let us refer to this as an 'agreement dependency'.)

(8)   a.   *exists*: V, <N>
                     3sg ←

       b.   *ignores*: V, <N N>
                     3sg ← →

Like argument dependencies, agreement dependencies are properties of individual lexical items. Crucially, however, the lexicon provides no information about how agreement dependencies are to be resolved. This problem is left to the second part of the sentence-building system—the computational system.

In the theory I propose, the computational system is indistinguishable from a processor—it operates in a linear manner, it combines elements, and it checks to make sure that lexical requirements are being satisfied. However, unlike familiar linguistic processors, the computational system that I have in mind is not constrained by grammatical principles. Rather, it is subject to a single efficiency-related imperative that is independent of language—it must minimize the burden on working memory. In the case of sentence building, it does this by complying with the following corollary requirement.

(9)   *The Efficiency Requirement*
         Resolve dependencies at the first opportunity.

The computational system/processor is constrained in this way regardless of whether it is used for sentence production (in which case it acts on words drawn from the speaker's own lexicon) or for

comprehension (in which case it operates on words presented to it by another speaker). The effects of efficiency are thus felt in both speaking and understanding.

4.2. How agreement works

The simplest manifestation of agreement occurs in sentences such as *Hope exists*, in which there is a single noun and a single verb. The processor brings the two together, creating an opportunity to resolve both the verb's argument dependency (not represented here) and its agreement dependency. Given the Efficiency Requirement, there is no choice but to immediately resolve both dependencies, thereby ensuring that the verb agrees with its 'subject' in this case.[4] (A check mark is used to indicate that a dependency has been resolved.)

(10)   [Hope exists]
              IIISɢ 3sg$\sqrt{}$

Matters are slightly more complicated in the case of a pattern such as *There is water (on the floor)*, in which the verb has the following lexical properties.

(11)   *is*: V, <*there* N>
              3sg

Here the processor, working from left to right, first brings together *there* and *is*, leading to resolution of the verb's first argument dependency. However, because *there* lacks number features, no opportunity arises to resolve the verb's agreement dependency.

(12)   [There is]
                  3sg

In the next step, the computational system combines *is* with the nominal to its right, creating an opportunity not only to resolve the verb's second argument dependency but also to take care of the agreement dependency. Consistent with the Efficiency Requirement, both opportunities are immediately exploited, thereby creating a pattern in which the verb agrees with a non-subject nominal.

(13)   [There [is   **milk**]]
                  3sg$\sqrt{}$ IIISɢ

An even more radical result occurs in the case of patterns such as *There is milk and water on the table*. Sentence formation proceeds in the manner summarized below, with the processor operating in a strictly linear manner.

(14)   a.   Combination of *there* and *is*:
              [There is]
                  3sg

---

[4] I use the term 'subject' as a descriptive convenience only; a more accurate term would simply be the verb's 'first argument'.

       b.   Combination of *is* and *milk*; resolution of the agreement dependency:
                [There [is   **milk**]]
                    3sg√ IIISɢ

       c.   Combination of *milk* and *and*:
                [There [is   [milk and]]]
                    3sg√ IIISɢ

       d.   Combination of *and* and *water*, creating a plural coordinate NP:
                [There [is   [milk [and water]]]]
                    3sg√ IIISɢ   IIISɢ

The key step here is the second one, in which the verb combines with just the first conjunct of the coordinate phrase, the nominal *milk*, creating the opportunity to resolve the agreement dependency.

    The end result is the phenomenon known as 'partial agreement'—the verb agrees with a subpart of one of its arguments. As expected, this phenomenon is only possible when the coordinate NP follows the verb. Where it appears to the left, and is therefore fully formed before the verb is encountered, partial agreement is impossible.

(15)   [Milk and water] are/*is on the floor.

This points to the importance of a dynamic, linear approach to sentence building that forces the computational system to build the sentence in a 'realistic' manner—one word a time from start to finish. More static approaches, which check agreement in fully formed sentential structures, are less well equipped to capture the left-to-right asymmetry that arises in the case of partial agreement.

    In sum, the agreement facts are not in the lexicon—the lexicon does no more than list the dependencies that must be resolved; it says nothing about *how* they are resolved. And the agreement facts are not in the computational system—it is a simple efficiency-driven processor that says nothing at all about language per se. Rather, the agreement facts arise from the interaction of lexical properties (agreement dependencies) with the properties of a parsimonious processor. The end result is a system with the curious – but emergent – properties associated with English verbal agreement.

4.3. A role for pragmatics

    Pragmatic considerations also provide fertile ground for emergentist explanations of language. In fact, MacWhinney (2005) argues that the pragmatic notion of perspective (and the way it is tracked, shifted, and shared) provides the key to understanding the properties of a wide range of syntactic phenomena. Interestingly, and consistent with the general thrust of MacWhinney's work, there is a significant tradition within formal linguistics of seeking to understand various phenomena in just this way. The study of pronominal conference is a case in point.

    It is universally recognized that pragmatic factors play a central role in determining the reference of pronominal expressions, as the following striking contrast from Winograd (1972) helps illustrate.

(16) a.   The city council denied the demonstrators the permit because they advocated violence.

b.   The city council denied the demonstrators the permit because they feared violence.

As noted by Kehler (2002) and Wolf et al. (2004), listeners interpret the pronouns in such sentences in a way that makes the discourse coherent. Hence *they* is taken to refer to *the demonstrators* in (16a) but to *the city council* in (16b), consistent with general assumptions about groups that might advocate or fear violence.

It has even been suggested that pragmatic considerations may do at least some of the work usually attributed to grammatical principles. Consider in this regard the familiar contrast illustrated in (17).

(17)   John$_i$ overestimates himself$_i$/*him$_i$.

The best-known grammatical account of this contrast involves binding theory—reflexive pronouns require a 'local' antecedent (Principle A), whereas plain pronouns reject such an antecedent (Principle B). O'Grady (2005:29ff) shows that Principle A effects follow straight-forwardly from the operation of a linear, efficiency driven processor, and that there is therefore no need for a grammatical principle to account for the interpretation of reflexive pronouns. But what of Principle B? In particular, why can plain pronouns not take a local antecedent in sentences such as (17)?

As a number of scholars have noted (e.g., Reinhart, 1983:166; Levinson, 1987:410), Principle B effects can be explained in terms of informativeness.

(18)   *The Informativeness Principle*
        Be as informative as possible; avoid ambiguity.

As Reinhart and Levinson observe, a reflexive pronoun is more informative than a plain pronoun. Because a reflexive requires a local antecedent, its referent is straightforwardly determinable – in (17), it can only be John. In contrast, the referent of the plain pronoun is subject to no such constraint – it could be anyone. It therefore makes sense to employ the reflexive pronoun where it is permitted, reserving the less specific plain pronoun for cases where the reflexive is not allowed.

On this view then, no grammatical principle prevents *him* from referring to John in sentences such as (17). Speakers simply avoid using the pronoun in this way for pragmatic reasons—where permitted, choice of the reflexive pronoun eliminates potential ambiguity and is therefore more informative.[5]

If this is right, then the constraint on the interpretation of plain pronouns is not in the lexicon—the entry for *him* does not say that it cannot have a local antecedent. Nor is it in the computational system—Principle B simply does not exist. Rather, the familiar facts arise from the interaction of pragmatic considerations (in particular, informativeness) with the ABSENCE of constraints on the interpretation of plain pronouns.

---

[5] This is not my position, however. As explained in O'Grady (2005:38ff), I attribute the unacceptability of a plain pronoun in patterns such as (17) to considerations of computational efficiency.

5. Language acquisition

The roots of the emergentist approach to language acquisition can be traced to Jean Piaget (e.g., Piaget, 1954), whose theory of development emphasized the interaction between experience and general cognitive mechanisms (Bates et al., 1998:590; MacWhinney, 2002). It is generally acknowledged, of course, that Piaget's work on language was flawed in a variety of ways. Among other things, it left open the precise nature of the learning mechanisms and their operation, and it said absolutely nothing about the sorts of complex phenomena that are central to the concerns of modern linguistic theory.

Nonetheless, much contemporary emergentist research remains committed to the idea that language acquisition can be reduced to the use of simple learning mechanisms to extract statistical regularities present in ordinary linguistic input. This is sometimes referred to as a 'usage-based' based view of development—knowledge of language is created and strengthened in response to opportunities to interpret and/or form utterances in the course of communication.

A particularly strong version of this view, put forward in somewhat different forms by various people (including me), holds that there is no special language acquisition device. Rather, the cognitive mechanism underlying acquisition is simply a processor—in fact, the very same processor that is responsible for interpreting and forming sentences in real time during actual language use. In other words, language emerges in children through the repeated processing of sentences that they encounter in the speech of others.

Of course, this raises the question of what the processor used by children looks like and how it operates. The emphasis in most emergentist work has been on a single property of this processor—its ability to identify and calculate distributional contingencies ('given x, there is a particular likelihood of having y'). This is particularly obvious in the case of connectionist modeling involving Simple Recurrent Networks (e.g. Elman, 1993, 2002, 2005b; Chang et al., 2006), which are specifically designed to take note of local co-occurrence relationships ('transitional probabilities').

A focus on frequency and distributional contingencies is also the hallmark of much non-computational work within the emergentist framework. For example, Tomasello (2003:327) suggests that the process whereby language structure emerges from language use 'depends crucially on the type and token frequency with which certain structures appear in the input'. And Ellis (2006a:8) puts forward the view that language learning is an 'intuitive statistical learning problem, one that involves the associative learning of representations that reflect the probability of occurrence of form-function mappings'.

There is no disputing the fact that processors are sensitive to frequency and to distributional contingencies. In fact, it is well established that even infants have a remarkable ability to compute transitional probabilities for phonemes (given the phoneme X, what is the likelihood that the next phoneme will be Y?) and to use this information to draw inferences about the likelihood of syllable and word boundaries (e.g., Jusczyk et al., 1999; Newport and Aslin, 2000). Moreover, probabilistic approaches to processing and acquisition in general are now well established and appear to have real contributions to make (see, e.g., Chater and Manning, 2006 for a review).

Although I too hold that the processor is the engine that drives language acquisition (that is a central thesis of O'Grady, 2005), I believe that achieving the explanatory goals of linguistics will require reference to more just transitional probabilities. The phenomenon of agreement offers a useful example for illustrative purposes.

Lewis and Elman (2001) trained a simple recurrent network (SRN) on sentence types such as the following, tokens of which were presented with a relative frequency that matched that found in parental speech to children.

(19)   a.   The women smoke cigarettes.
          b.   The boy is funny.
          c.   Can you read this?
          d.   Is she friendly?
          e.   The girl with the cat is nice.
          f.   The toy with the wheels is mine.
          g.   The man who came late is your uncle.
          h.   I like the teacher who speaks Spanish.

The network's success was judged by its ability to predict, at each point in the sentence, what word(s) might come next, based on the words that had already been encountered—a classic test of the ability to learn distributional contingencies. For instance, after hearing the words *the men* at the beginning of a sentence, would the network predict the next word to be *is* or *are*?

(20)   The men {is/are} . . .

More crucially, after hearing the words *the students with long hair*, would it predict *smokes* or *smoke*?

(21)   The students with long hair {smokes/smoke} . . .

Being able to make the right predictions in these cases essentially demonstrates mastery of the agreement facts. And, indeed, Lewis & Elman report that the network came to make the right predictions, 'learning' that the verb agrees with the head of its subject argument in this case, not simply with the nearest preceding noun (e.g., *hair* in (21)).

This is indeed impressive, but something is missing, I believe. In particular, the existence of the very statistical regularities in the input exploited by the SRN begs the question of why those regularities – rather than various others – are there in the first place. In other words, why does language have the particular properties that it does? Why, for example, are there languages (such as English) in which a verb routinely agrees with its subject rather than its direct object, but no language in which a verb agrees with its direct object rather than its subject (e.g., Croft, 2003:147)?[6]

Networks provide no answer to this sort of question. In fact, if presented with data such as the following, in which the verb agrees with its direct object rather than its subject, an SRN would no doubt 'learn' just this sort of pattern, even though it is not found in any known human language.

(22)   a.   Singular subject, plural verb form, plural direct object:
               The woman see**Ø two men**.

          b.   Plural subject, singular verb form, singular direct object:
               The women see**s a man**.

---

[6] I deliberately set to the side the case of ergative languages, for which there are ongoing disputes concerning the identity of the subject argument.

There is clearly something wrong here. Humans do not just learn language; they shape it. Moreover, these two facts are surely related in some fundamental way, which is why hypotheses about how linguistic systems are acquired need to be embedded within a more comprehensive theory of why those systems (and therefore the input) have the particular properties that they do. Gregg (2003:69) makes a similar point, noting the burden on emergentism to provide a 'property theory of linguistic competence that can rival [UG-based theories]'.

The prospects for an emergentist response to this challenge are in fact excellent. As we saw earlier, for instance, there is a plausible emergentist account (based on processing considerations) for why a verb should agree with its subject rather than its direct object in English—as first argument, the subject has the earlier opportunity to resolve the verb's agreement dependencies.

Crucially, this account fits well with what we know about the general ease and rapidity of language acquisition. In the case of agreement, for instance, it cannot hurt that an efficiency-driven processor directs learners to look initially to the verb's first argument (its 'subject') for an opportunity to resolve agreement dependencies. Or that, in the absence of agreement features on the first argument, as happens in *There is milk on the table*, it directs attention to the verb's second argument. These are precisely the sorts of facts about agreement that need to be mastered, and it cannot be a coincidence that they exactly mirror what the processor is predisposed to do in the first place.

As this simple illustration suggests, positing the right type of processor not only contributes to an account of why language has the particular properties that it does, it can significantly lighten the burden on language learners, reducing their need to rely on statistical regularities in the input to figure out how agreement (and other phenomena) work. This is a desirable result, even for input-based approaches to learning. As Lewis and Elman (2001:369) themselves acknowledge, 'determining the potential worth of . . . stochastic information is difficult'. There is clearly much that has to be gleaned from experience (including, most obviously, the form that agreement takes in a particular language); nothing is lost by easing the demand on the inductive mechanisms where possible.

Of course, not everything in language is driven by processing considerations. Some things are arguably motivated by other sorts of factors (pragmatics, for instance), and some things (perhaps adjective–noun order, as mentioned earlier) are simply arbitrary. In such cases, the processor has to learn the right routine ('look to the right for a noun') based on nothing more than exposure to particular patterns in the input. This in not a problem—in fact, we know that processors, even very simple processors, can do this very well, as the SRN work shows.

The more interesting problems involve phenomena whose properties are in some way principled—the very phenomena that have been at the heart of linguistic theory for the past several decades (binding, control, agreement, and so forth). And here we see that the explanation lies not in the input, but in the nature of the processor that analyzes the input. That processor is surprisingly simple, but it does more than just track frequency, create associations, and compute distributional contingencies. An equally important property is the mandate to minimize the burden on working memory—a design feature with far-reaching consequences both for the structure of language and for the manner in which it is acquired, as the extensive discussion in O'Grady (2005) attempts to demonstrate.

Second language learning presents an even more complicated situation, thanks to the relevance of additional factors—prior knowledge of another language, type of instruction, possible critical period effects, and individual differences in motivation, among others. Proposals have already been put forward for phenomena as disparate as grammatical morphology (Ellis, 2006b) and quantifier scope (O'Grady, 2007), but it must be acknowledged that work on an

emergentist theory of SLA is still in its infancy. It is perhaps not too much to hope that the papers in this volume, and the discussion that they stimulate, will help clarify the prospects and problems associated with this research program.

6. Concluding remarks

Sentences exhibit grammatical properties that are without a doubt worthy of study, both for their inherent interest and for what they reveal about the mind. However, grammatical *rules* and *principles* have no such reality or relevance. The mechanisms making up the language system are at once less obvious and simpler than orthodox grammatical analyses would have us believe.

There is currently no comprehensive emergentist theory of language or its acquisition, but there is a well-defined emergentist program for the investigation of those phenomena. That program is based on the simple thesis that the core properties of language are best understood with reference to more fundamental non-linguistic (i.e., 'non-grammatical') factors and their interaction. The viability of this idea can and must be measured against its success in confronting the classic problems of linguistic analysis—figuring out how language works and how it is acquired.

The standards for success in this regard are high. (In fact, I am in complete agreement with Gregg's, 2003:53 admonition to emergentists: 'if UG is the best explanation, so be it'.) Traditional forms of grammatical analysis have made substantial descriptive advances that support highly interesting hypotheses about both the nature of language and the manner in which it is acquired. The promise of the emergentist program lies in its potential to surpass these achievements by tracing the roots of language to more general mechanisms of cognition and learning, whose precise character and importance are the focus of current work.

Acknowledgements

Appendix A. Is the Minimalist Program emergentist?

The notion of Universal Grammar, the cornerstone of nativist approaches to the language faculty, has undergone considerable evolution since its inception (O'Grady, 2001b:623ff; Chomsky, 2005:9). In its first instantiations, Universal Grammar contained principles whose grammatical character was beyond dispute. Early proposals – Ross's 'island constraints', for instance – were stated in terms of familiar grammatical notions such as NP, clause, and the like.

(i) *The Complex NP Constraint* (Ross, 1967:127)
No element in a sentence dominated by a noun phrase with a lexical head noun can be moved out of the noun phrase.

During the Government and Binding era, principles became more abstract and the grammatical notions on which they drew (barrier, government, governing category, and the like) correspondingly more exotic, although still clearly faculty-specific in character.

(ii) *The Empty Category Principle* (Chomsky, 1981:250)
An empty category must be properly governed.

A further significant change occurred with the advent of the Minimalist Program, with its emphasis on constraints involving locality and economy, notions that invite a more general computational interpretation.

(iii)   *The Minimal Link Condition* (Chomsky, 1995:264)
          Make moves as short as possible.

Although care has always been taken to distinguish the grammar from the processor (e.g., Chomsky, 1995:167), UG continues to evolve away from the autonomous grammatical system that it was once taken to be—to the point where Chomsky (2005:8) observes:

> There is no longer a conceptual barrier to the hope that the UG might be reduced to a much simpler form, and that the basic properties of the computational systems of language might have a principled explanation instead of being stipulated in terms of a highly restrictive language-specific format for grammars.

An 'explanation of the properties of language is principled', Chomsky goes on to note (p. 10), 'insofar as it can be reduced to properties of the interface system and general considerations of computational efficiency and the like'—considerations that may well 'apply beyond language' (p. 9).

So far, so good. To the extent that the proposed innate endowment for language no longer makes reference to specifically linguistic notions and mechanisms, it becomes more compatible with the emergentist program. However, there is reason to think that things may not be quite as they appear.

A good deal of the advance described by Chomsky is based on a far-reaching reconceptualization of what constitutes the language faculty. On the one hand, there is *the faculty of language in the broad sense* (FLB), which includes 'all of the many mechanisms involved in speech and language, regardless of their overlap with other cognitive domains or with other species' (Fitch et al., 2005:180–181). On the other hand, there is *the faculty of language in the narrow sense* (FLN), consisting of 'the subset of mechanisms that is both unique to humans, and to language' (p. 181).

Chomsky's proposal, first stated in Hauser et al. (2002:1573), is that FLN can be reduced to 'mechanisms of recursion as they appear in narrow syntax and the mapping to the interfaces'. A clarification in Fitch et al. (2005:183) revises the hypothesis to 'at a minimum, . . . FLN *includes* the capacity of recursion' [emphasis added]. However, the hypothesis is later restated in its strong form (p. 189): 'FLN is restricted to a simple but powerful recursive mapping capability by definition, unique to humans and unique to the language faculty' (see also p. 202).

This is potentially problematic, since it appears to ignore the many core linguistic phenomena that are not recursive, including phonology, morphology, case, agreement, binding, quantification, and so forth, as Pinker and Jackendoff (2005:216) note. Fitch et al. (2005:182) insist that this is not so (see also Chomsky et al., 2005), but they nonetheless refrain from putting forward an FLN-based account of these phenomena.

In any event, even if we focus on the role of the recursive system in structure building, there are problems. One difficulty has to do with precisely how we are to interpret the notion of computational efficiency that the system is supposed to manifest.

Take for instance the proposal that the derivation proceeds in 'phases', with vPs and CPs receiving a phonological and semantic interpretation immediately upon being formed, after which they are 'forgotten' and become inaccessible to later operations. Chomsky (2005:17)

suggests that this 'minimize[s] computational load', but it is far from clear how this claim can be interpreted or evaluated psycholinguistically, given that the minimalist computational system that he employs builds structure from *right to left*—the reverse of what actual processors do.

Another difficulty stems from the fact that 'natural conditions of efficient computation' to which minimalists appeal are highly theory-internal in their conception, status, and consequences. Two such conditions are summarized by Chomsky et al. (2005:3):

(iv)   Syntactic objects that have been merged must be preserved without modification (thus 'eliminating extra computation').

(v)   Phonological computation must be minimized.

As Chomsky et al. (2005) note, adherence to these conditions yields the 'copy theory of movement', which forms the following representation (ignoring phases) for the sentence *What should Mary pick*?. The presence of the italicized 'copies' is attributed to the first principle of efficient computation; the fact that they are not pronounced is attributed to the second.

(vi)   [$_{CP}$ What should [$_{TP}$ Mary *should* [$_{vP}$ *Mary* pick *what*]]]

Such representations may well follow from computational considerations by minimalist logic, but it is clear that both the logic and the computations have a parochial character not revealed in Chomsky's more general remarks on the explanatory role of efficiency.[7]

Other sorts of problems arise when we scrutinize attempts to formulate the principles of computational efficiency required to make the minimalist view of the language faculty work. As Newmeyer (2003:588) observes:

> . . . as many distinct UG principles are being proposed today as were proposed twenty years ago. I would go so far as to claim that no paper has ever been published within the general rubric of the minimalist program that does not propose some new UG principle or make some new stipulation (however well motivated empirically) about grammatical operations that does not follow from the bare structure of the [Minimalist Program].

Take for example the central idea that that principles of efficient computation govern the feature checking operations that drive structure building in the Minimalist Program. Chomsky reports (2005:10) that Frampton and Gutmann (1999), among others, have achieved 'important results' in the application of these principles to language. These results include the following observations (Frampton and Gutman, p. 7).

> The first step [in featuring checking] is locating a head Y which X recognizes. [A head X recognizes a head Y if a removable feature +f of X can be removed by a feature set *v*, at least some part of which is contained in F(Y).] If X has no removable features, this step will necessarily fail. If X does have at least one removable feature, there are two options. One option . . . is taking a Y directly from the lexicon. More precisely, a copy Y of a purely lexical item is generated such that X recognizes Y. Because Y does not enter the syntax via

---

selection, this option will be limited to expletive Y. The second (much more common) option is finding a Y in (X,*a*). The search is top down. If a Y which X recognizes is found, the search does not go more deeply down into *a*. The major subtlety is specifying exactly what it means to say that Z is deeper in *a* than Y is, so that if X recognizes Y, then Y blocks access to finding a Z (which X recognizes) in the top down search of *a*. This is the question of equidistance (Chomsky, 1995:298). We will sidestep this issue here, since the intervention effects which are relevant in this paper are not delicate.

Here again, the parochial character of the computational system is revealed, calling into question its supposed simplicity and suggesting that it has retained its special grammatical character, although perhaps in a more abstract form.

For these reasons, it appears that the contemporary minimalist program does not count as emergentist.[8]

References

Bates, E., Goodman, J., 1999. On the emergence of grammar from the lexicon. In: MacWhinney, B. (Ed.), The Emergence of Language. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 29–79.

Bates, E., MacWhinney, B., 1988. What is functionalism? Papers and Reports on Child Language Development 27, 137–152.

Bates, E., Elman, J., Johnson, M., Karmiloff-Smith, A., Parisi, D., Plunkett, K., 1998. Innateness and emergentism. In: Bechtel, W., Graham, G. (Eds.), A Companion to Cognitive Science. Blackwell, Oxford, pp. 590–601.

Bybee, J., 1998. The emergent lexicon. In: Proceedings of the 34th Regional Meeting of the Chicago Linguistic Society: The Panels, Chicago Linguistic Society, Chicago, pp. 421–435.

Chang, F., Dell, G., Bock, K., 2006. Becoming syntactic. Psychological Review 113, 234–272.

Chater, N., Manning, C., 2006. Probabilistic models of language processing and acquisition. Trends in Cognitive Sciences 10, 335–344.

Chomsky, N., 1981. Lectures on Government and Binding. Foris, Dordrecht.

Chomsky, N., 1995. The Minimalist Program. MIT Press, Cambridge, MA.

Chomsky, N., 2005. Three factors in language design. Linguistic Inquiry 36, 1–22.

Chomsky, N., Hauser, M., Fitch, W.T., 2005. The minimalist program. Unpublished appendix to Fitch et al. 2005, available at www.wjh.harvard.edu/~mnkylab.

Christiansen, M., Chater, N., 2001. Connectionist psycholinguistics: Capturing the empirical data. Trends in Cognitive Science 5, 82–88.

Croft, W., 2003. Typology and Universals, second ed. Cambridge University Press, Cambridge.

Dryer, M., 1992. The Greenbergian word order correlations. Language 68, 81–138.

Ellis, N., 1998. Emergentism, connectionism and language learning. Language Learning 48, 631–664.

Ellis, N., 2006a. Language acquisition as rational contigency learning. Applied Linguistics 27, 1–24.

Ellis, N., 2006b. Selective attention and transfer phenomena in L2 acquisition: Contingency, cue competition, salience, interference, overshadowing, blocking, and perceptual learning. Applied Linguistics 27, 164–194.

Elman, J., 1993. Learning and development in neural networks: The importance of starting small. Cognition 48, 71–99.

Elman, J., 1999. The emergence of language: A conspiracy theory. In: MacWhinney, B. (Ed.), The Emergence of Language. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 1–27.

Elman, J., 2002. Generalization from sparse input. In: Proceedings of the 38th Regional Meeting of the Chicago Linguistic Society, Chicago Linguistic Society, Chicago, pp. 175–200.

Elman, J., 2005a. An alternative view of the mental lexicon. Trends in Cognitive Science 9, 342–348.

Elman, J., 2005b. Computational approaches to language acquisition. In: Brown, K. (Ed.), Encyclopedia of Language and Linguistics. second ed. Elsevier, Oxford.

---

[8] This is not to say that Chomsky does not adopt an emergentist stance on other issues. To the contrary, he has endorsed the proposition that 'things mental, indeed minds, are emergent properties of brains', observing that this is close to a 'truism' (2002:74). The point that I have been considering relates to the question of whether the properties of the brain's 'language organ' (p. 64) are themselves emergent or whether they must be stipulated by faculty-specific principles.

Fitch, W.T., Hauser, M., Chomsky, N., 2005. The evolution of the language faculty: Clarifications and implications. Cognition 97, 179–210.

Frampton, J., Gutmann, S., 1999. Cyclic computation, a computationally efficient minimalist syntax. Syntax 2, 1–27.

Goldberg, A., 1999. The emergence of the semantics of argument structure constructions. In: MacWhinney, B. (Ed.), The Emergence of Language. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 197–212.

Goldberg, A., Jackendoff, R., 2004. The English resultative as a family of constructions. Language 80, 532–568.

Gregg, K.R., 2003. The state of emergentism in second language acquisition. Second Language Research 19, 42–75.

Hauser, M., Fitch, W.T., Chomsky, N., 2002. The language faculty: What it is, who has it, and how did it evolve? Science 298, 1569–1579.

Hay, J., 2003. Causes and Consequences of Word Structure. Routledge, New York.

Hay, J., Baayen, R.H., 2005. Shifting paradigms: Gradient structure in morphology. Trends in Cognitive Science 9, 342–348.

Jusczyk, P., Houston, D., Newsome, M., 1999. The beginnings of word segmentation in English-learning infants. Cognitive Psychology 39, 159–207.

Kehler, A., 2002. Coherence, Reference, and the Theory of Grammar. Center for the Study of Language and Information, Stanford University, Palo Alto, CA.

Levinson, S., 1987. Pragmatics and the grammar of anaphora: A partial pragmatic reduction of binding and control phenomena. Journal of Linguistics 23, 379–434.

Lewis, J., Elman, J., 2001. Learnability and the statistical structure of language: Poverty of stimulus arguments revisited. In: Proceedings of the 26th Annual Boston University Conference on Language Development, Cascadilla Press, Somerville, MA, pp. 359–370.

MacWhinney, B., 2002. Language emergence. In: Burmeister, P., Piske, T., Rohde, A. (Eds.), An Integrated View of Language Development: Papers in Honor of Henning Wode. Wissenshaftliche Verlag, Trier, pp. 17–42.

MacWhinney, B., 2005. The emergence of grammar from perspective. In: Pecher, D., Zwaan, R. (Eds.), Grounding, Cognition. The Role of Perception and Action in Memory, Language and Thinking. Cambridge University Press, Cambridge, pp. 198–223.

Marcus, G., 1998. Rethinking eliminative connectionism. Cognitive Psychology 37, 243–282.

Marcus, G., 2001. The Algebraic Mind. MIT Press, Cambridge, MA.

McClelland, J., Patterson, K., 2002. Rules or connections in past-tense inflection: What does the evidence rule out? Trends in Cognitive Science 6, 465–472.

Mill, J.S., 1930 [1843]. A System of Logic Ratiocinative and Inductive. Longmans, Green and Co., London.

Morris, W., Cottrell, G., Elman, J., 2000. A connectionist simulation of the empirical acquisition of grammatical relations. In: Wermter, S., Sun, R. (Eds.), Hybrid Neural Systems. Springer-Verlag, Heidelberg.

Newmeyer, F., 2003. Review of On nature and language, by N. Chomsky. Language 79, 583–599.

Newmeyer, F., 2005. Commentary on Denis Bouchard. Exaption and linguistic explanation. Lingua 115, 1667–1671.

Newport, E., Aslin, R., 2000. Innately constrained learning: Blending old and new approaches to language acquisition. In: Proceedings of the 24th Annual Boston University Conference on Language Development, Cascadilla Press, Somerville, MA, pp. 1–21.

O'Grady, W., 2001a. An emergentist approach to syntax. Available at http://www.ling.hawaii.edu/faculty/ogrady/.

O'Grady, W., 2001b. Toward a new nativism. Studies in Second Language Acquisition 21, 621–633.

O'Grady, W., 2005. Syntactic Carpentry: an Emergentist Approach to Syntax. Lawrence Erlbaum Associates, Mahwah, NJ.

O'Grady, W., 2007. The syntax of quantification in SLA: An emergentist approach. In: O'Brien, M., Shea, C., Archibald, J. (Eds.), Proceedings of the 8th Generative Approaches to Second Language Acquisition Conference (GASLA 2006): The Banff Conference, Cascadilla Press, Somerville, MA, pp. 98–113.

Palmer-Brown, D., Tepper, J., Powell, H., 2002. Connectionist natural language parsing. Trends in Cognitive Science 6, 437–442.

Piaget, J., 1954. The Construction of Reality in the Child. Basic Books, New York.

Pinker, S., 1999. Words and Rules. Basic Books, New York.

Pinker, S., Jackendoff, R., 2005. The faculty of language: What's special about it? Cognition 97, 201–236.

Reinhart, T., 1983. Anaphora and Semantic Interpretation. University of Chicago Press, Chicago.

Ross, J.R., 1967. Constraints on Variables in Syntax. Unpublished Ph.D. dissertation. Department of Linguistics and Philosophy, MIT.

Smolensky, P., 1999. Grammar-based connectionist approaches to language. Cognitive Science 23, 589–613.

Sobin, N., 1997. Agreement, default rules, and grammatical viruses. Linguistic Inquiry 28, 318–343.

Steedman, M., 1999. Connectionist sentence processing in perspective. Cognitive Science 23, 615–634.

Stephan, A., 1997. Armchair arguments against emergentism. Erkenntnis 46 (3), 305–314.

Tomasello, M., 2003. Constructing a Language: a Usage-based Theory of Language Acquisition. Harvard University Press, Cambridge, MA.

Winograd, T., 1972. Understanding Natural Language. Academic Press, San Diego.

Wolf, F., Gibson, E., Desmet, T., 2004. Coherence and pronoun resolution. Language and Cognitive Processes 19, 665–675.

Further reading

Chomsky, N., 1977. Essays on Form and Interpretation. North Holland, New York.

Chomsky, N., 2001. Beyond explanatory adequacy. MIT Occasional Papers in Linguistics 20, 1–28.

Chomsky, N., 2002. On Nature and Language. Cambridge University Press, Cambridge.

Ellis, N., 2002. Frequency effects in language processing. Studies in Second Language Acquisition 24, 143–188.

Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., Plunkett, K., 1996. Rethinking Innateness. A Connectionist Perspective on Development. MIT Press, Cambridge, MA.

Gleick, J., 1987. Chaos: Making a New Science. Penguin Books, New York.

Hopper, P., 1998. Emergent grammar. In: Tomasello, M. (Ed.), The New Psychology of Language. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 155–175.

Jackendoff, R., Pinker, S., 2005. The nature of the language faculty and its implications for evolution of language (reply to Fitch et al.). Cognition 97, 211–225.

Larsen-Freeman, D., 1997. Chaos/complexity theory and second language acquisition. Applied Linguistics 18, 141–165.

Larsen-Freeman, D., 2002. Language acquisition and language use from a chaos/complexity theory perspective. In: Kramsch, C. (Ed.), Language Acquisition and Language Socialization: an Ecological Perspective. Continuum Press, New York, pp. 33–46.

MacWhinney, B., 1998. Models of the emergence of language. Annual Review of Psychology 49, 199–227.

MacWhinney, B., 2004. A multiple process solution to the logical problem of language acquisition. Journal of Child Language 31, 883–914.

Waldrop, M., 1992. Complexity: the Emerging Science at the Edge of Order and Chaos. Simon & Schuster, New York.