Peter Ludlow, "The Myth of Human Language" (2005)

God, having designed man for a sociable creature, made him not only with an inclination and under a necessity to have fellowship with those of his own kind, but furnished him also with language, which was to be the great instrument and common tie of society.

John Locke
Essay Concerning Human Understanding

The standard view about language is roughly the one advanced by Locke in the epigraph above (only sans God): Languages like Urdu, German, Polish and Portuguese are robust and fairly stable abstract systems of communication that are learned (with varying degrees of success) by human agents. Those agents in turn use the languages that they have learned to communicate ideas, perform certain tasks (by giving orders, instructions, etc), and in some cases as media for artistic expression. It is often argued that the better one learns a language, the better equipped one is to successfully communicate, accomplish complex tasks, etc. Sometimes the standard view uses the metaphor of language as a widely shared common currency that agents use to communicate, with individuals words being the common coins of the realm. Crucially, these common coins are fixed; as Locke argued, even Augustus, though he ruled the world, was unable to coin new Latin words.

This standard view, although not universally held, is at least widely held by academics and lay persons alike, ranging from philosophers and language instructors, to anthropologists and computational linguists, to politicians and pundits. Unfortunately, the standard view is badly mistaken, and its uncritical acceptance has had devastating consequences in all of these domains. It has led anthropologists to think that languages constrain the conceptual space of language users, and has led language departments into disastrous alliances with French philosophers of language and their word salad warfare on the hegemony of "language" and its alleged tendency to encode the interests of patriarchy and other evil powers. It has led to wooden approaches to language instruction on the one hand and to failed attempts at human/machine communication on the other. On the political end, it has led to silliness on both the left and the right by way of attempts to clean up or reform or otherwise render standard languages politically correct - a general sentiment that has led to downright discriminatory social policies like English Only laws, and in its extreme form, to attempts at language purification by Fascists like Mussolini.

It would make for an interesting book to sort out all the ways in which the standard view of language has crept into these domains of human activity and the corrosiveness of its effect, but my goal here is substantially less ambitious. I'll begin in part 1 with the negative case against the "common coin" view of language and offer an alternative in part 2. The alternative picture will be one in which there is a core part of our linguistic competence that is fixed by biology (perhaps by low level biophysical principles) but that this provides just a basic skeleton which is fleshed out in different ways on a conversation-by-conversation basis. To shift back to the monetary metaphor, there are some common coins, but we also have the ability to mint new coins on the fly in collaboration with our discourse partners, to control which of those common coins are in circulation at any given time, and to coordinate and precisify the shared meanings of those

common coins that are in use. As we will see, for most linguistic common coins the meaning is vastly underdetermined. I will suggest possible ways in which coins are minted and their values determined as discourse participants form dynamic communicative partnerships, resulting (if we really must deploy the term 'language') in what we might call micro-languages.   In part 3 I will provide an example of how the received view has led us astray.

1.0 The Myth of Human Language

Let's begin by following Chomsky (1986) in making a distinction between I-language and E-language. Thus far, when talking about "language," I have been talking about language in the sense of E-language - that is, the conception of language as an "external" stable abstract object that is that is deployed by a large population of language users for purposes of communication. I-language, on the other hand, is Chomsky's term for the internal cognitive mechanisms that underwrite our linguistic competence.

For most linguists, the traditional notion of an E-language like Portuguese or German or English is suspect at best (see Chomsky 1980; ch. 6). Consider the case of German, for example. In what sense is a "speaker of German" from the Dutch border of Germany and a "speaker of German" from Bavaria speaking the same language? (Especially given that their speech is not easily mutually intelligible?) The fact that we say these individuals speak the same language is more of a political decision than anything else, and indeed an individual raised in northern Germany and an individual raised in The Netherlands may find that their languages are more mutually intelligible than the two aforementioned German citizens do. As Max Weinreich is supposed to have said, "a language is a dialect with an army and a navy."

Saying that two individuals "speak the same language" is at best a loose way of talking about some superficial similarities in style of communication that have been raised to salience for political and social reasons rather than linguistic reasons. Chomsky (1994) compares it to saying that two cities are "near" each other; whether two cities are near depends on our interests and our mode of transportation and virtually not at all on brute facts of geography. The notion of "same language" is no more respectable a notion in linguistics than "nearness" is in geography. Informally we might group together ways of speaking that seem to be similar (relative to our interests), but such groupings have no real explanatory power.

One might try to retreat slightly by giving up on the idea of an E-language and endorsing a notion of E-dialect or E-idiolect, but even this retreat will not save the language-as-external-object position. Considerations that make it arbitrary when to say that two individuals speak the "same language" also apply to saying when they speak the "same dialect." That is, do the people in this village speak the same dialect as the people in the next village? Do we speak the same dialect as our next door neighbors? For that matter do we speak the same dialect as the people living under the same roof? My ex-wife pronounced the words 'Mary', 'marry', and 'merry' differently, while I pronounced them the same. Did we speak the same dialect? Do we? It depends on whether we want to identify with each other, and this is a kind of very local political decision.

At this point you might think we could retreat to a notion of E-idiolect - in other words, that there is an abstract thing that people speak, but it varies from person to person; each person has their own personal E-language. Even this is wrong (and it is certainly a mistake to conflate the notion of I-language with that of an idiolect - the notions cross-cut each other).

The problem with the notion of an E-idiolect is that we have no way of identifying the linguistic forms that would be part of a given individual's E-idiolect. Consider a hypothetical agent, Chesner. Chesner speaks in different ways with different groups of individuals (say Chesner uses a different vocabulary among philosophers than among family members) and indeed at different stages of life (contrast Chesner's use of language at age 3 and age 30). Do all of these ways of speaking count as being part of the same idiolect? What unifies them other than that they are ways in which Chesner happens to have spoken?

Of course one might argue that in this discussion I have more or less introduced a perfectly acceptable notion of language - one that is defined in terms of political identity etc. But the problem is that this notion of language, to the extent it is coherent, plays no explanatory role. Why do certain people communicate with each other so well? It is no answer to be told that they speak the same language. This is like saying that wood burns because it contains phlogiston or (to borrow a joke from Moliere) that opium makes one sleep because it has a "dormative virtue". In effect, we have expunged God from Locke's story about language, but the remaining picture still has all the explanatory power of intelligent design theory in biology: none.

What we are looking for here is some way of understanding how people are able to successfully communicate. Being told they speak "the same language" is no answer to this question, and in fact it leads us away from the answer; we want to probe the mechanisms that underwrite their ability to communicate. If we pursue these questions in a serious fashion, then we are driven far from the initial picture of a widely shared common coin system of communication.

2.0 The FLN and the LCS

Some writing on language tends to treat linguistic competence as a unified phenomenon, perhaps supervening on a single mechanism or module of human cognition. It seems more reasonable to suppose that the broad class of phenomena that we call "linguistic" or think of as having to do with "language" are supported by a combination of narrow mechanisms of the mind/brain (what Chomsky calls the FLN, for "faculty of language narrowly construed") and at the same time an entirely different set of abilities that are underwritten by world knowledge and various coordination strategies that we deploy. It seems doubtful that these coordination strategies are grounded in any single module of human cognition, but we can use the acronym LCS -linguistic coordination strategies - to speak of this broad class of abilities.

2.1 the FLN

Chomsky has argued that the FLN is a natural object that is part of our biological endowment. Recent work by Hauser, Chomsky, and Fitch (2002) has speculated that the mechanisms that underlie this core linguistic competence did not evolve gradually in response to selectional

pressures, but was sudden in evolutionary terms and involved what in effect amounts to a bio-physical wiring solution - a solution for hooking up the perceptual/articulatory system (the system that affords speech comprehension and production) with the conceptual/intentional system (the system that interprets and uses linguistic communication). The thesis is speculative, but not without supporting evidence. In the simplest form, support for the thesis involves the observation that low level physical and mathematical principles underlie many of the recursive patterns that we see in nature - ranging from the spiral patterns that we see in shells to the Fibonacci patterns we see in the distribution of seeds in a sunflower.

To illustrate the recursiveness of natural language, consider the following very simple case.

(1) This is the cat that ate the rat that ate the cheese that was made by the farmer that…

These sorts of patterns, in Chomsky's view, provide some of the evidence that the structure of the FLN is largely determined by basic biophysical properties.

It is reasonable to think that the FLN also underwrites significant aspects of the lexicon as well. This certainly seems to be the conclusion one would draw from work by Mark Baker (1988) which argues that that morphological and lexical properties are actually determined by the syntax. Following Higginbotham (1989) we can illustrate the basic idea by consider the following fragment from Lewis Carroll's poem "The Jabberwocky".

(2) Twas bryllyg, and the slythy toves did gyre and gymble in the wabe…

Just from the surrounding syntactic environment we can deduce quite a bit about the meaning of the term 'tove'. We know, for example, that toves are the sorts of things one can count (unlike substances like water), that they are spatially located and can move and undergo state changes (unlike numbers), they are quite possibly capable of acting under their own volition. All of this is defeasible, but these are reasonable suppositions to deduce from the surrounding linguistic structure.

Given the rapid acquisition of lexical items by children during their critical learning period (ages 1.5-6) and given their corresponding acquisition of and grasp of these basic thematic relations (provided only impoverished data, no reinforcement, etc) it seems reasonable to speculate that these thematic relations are part of the FLN. But as Bloom (2002) has argued, all of this just gives children a first pass at understanding word meanings. To flesh things out children also need to incorporate a broad range of contextual information real world knowledge. Of course, Bloom is assuming that there is "a meaning" to be learned. As I will argue below, it is more precise to say that children, like adults, must ultimately collaborate with their discourse partners to establish word meanings.

2.2 The LCS

While the FLN is often deployed in communicative endeavors, it is generally a very small part of such endeavors. As suggested earlier, it is, in effect, a kind of thin, rigid, skeleton onto which other types of abilities must be added for communication to be possible. What we need to

understand are the natures of those abilities - in particular the LCS -- the ways they interact with the FLN, and the consequences that can be extruded from the emerging picture.

The Dynamic Lexicon

minting new coins. One crucial element in this new point of departure has to do with the lexicon itself. As noted above, most philosophers of language suppose that the words we use are "common coin" in a broad social context -- that when you learn a language, among other things you learn a set stock of shared word meanings. Rejecting this picture, we can opt instead for the idea that many of the terms that we use are invented "on the fly" during individual conversations.

In effect this is a generalization of conclusions that have been reached by psycholinguists (e.g. Garrod and Anderson 1987, Brennan 1986, Brennan and Clark 1996, and Clark 1992) and their study of lexical "entrainment" -- a process whereby the use of certain words - sometimes novel words -- are introduced on the fly by discourse participants.

Studies on entrainment also undermine the myth of a common-coin lexicon by showing that even individuals who overhear or witness a conversation are in a much weaker position to understand what is being said than are the participants. Schober and Clark (1987), for example, show that discourse participants are in a much better position than eavesdroppers at understanding what is being said because participants are involved in the introduction of the lexical items that will be employed in the evocation of certain concepts.

Consider, for example, how much of a lecture you can comprehend by dropping in on a course in the middle of the term. If you are not familiar with the subject you may well be quite lost, and not just because you lack familiarity with the objects under discussion (if it is a philosophy class you might have dropped in on an unintelligible discussion of whether tables and chairs exist). One obstacle you may face is that you are unfamiliar with the terminology in play (of course, grasp of the terminology and knowledge of the subject matter are not so easily separated - more on this a bit later). You were not involved in the entrainment process whereby certain terms were introduced into the course. In such a situations you may dismiss the terms being used as "jargon", but this is just a way of saying that you don't understand the terms being deployed.

Common coins placed in and out of circulation. One important result of the entrainment experiments is that those common coins that do exist are not always in circulation, and indeed, are strategically retired and placed back into circulation depending upon the demands of the micro-language under construction. The situation is analogous to the position of the traveler who finds that various combinations of US Dollars, Euros, Yen, and Argentinean Pesos are accepted in different settings. Some are more widely accepted than others, and some can be introduced in the odd transaction with a bit of cajoling, but at the end of the day there are still establishments where only a Peso will do. Linguistic common coins are like this too, but their deployment is more strategic.

The experiments on entrainment are particularly illuminating here because they show that additional common coins are introduced into the micro-language in response to conversational demands on the discrimination of the concepts being deployed. If similar concepts are being

deployed (and the greater the need to discriminate concepts and kinds of objects), there is increased pressure to reissue certain coins.

common coins are thin. Linguistic common coins, whether in circulation frequently or rarely, are "thin." By that I mean that the shared part of the lexicon consists of just hints and clues (like one finds in dictionary entries) that may help us to deploy cognitive resources to flesh out the word meanings, and the way we flesh them out will vary according to contexts and social settings. A classic illustration would be the dummy terms like 'whatchamacallit' and 'thingamajigger', which are reissued often but typically with different denotations each time they are reissued.

Another example of this is the meaning of the term 'good'. This is a widely shared common linguistic coin, but there is much to its meaning that is underdetermined. For example, it is a typical phenomenon of sports talk radio to debate which of two sports stars is better. Was Mickey Mantle better than Barry Bonds at baseball? Well, one of them hit more home runs, but the other was on more championship teams. One of them may have cheated by using steroids. Should that be a factor? What is really up for grabs here is the question of what counts as a "good" baseball player - it is about the meaning of 'good'.

Jamie Tappenden (1999) offers a formal example of this phenomenon, introducing a language in which some meanings are open-ended and to be precisified at a later time. The language leaves "certain objects as 'unsettled' cases of a given predicate, in that it is open to the speakers of the language to make a further stipulation that the object is, or is not, to be counted as having the property in question."

As Tappenden notes, these cases happen frequently both unintentionally and intentionally outside of formal languages, with an example of intentional cases coming from the realm of law:

This happens with some frequency in law: it may be convenient to stipulate a condition for only a restricted range, leaving further stipulation for the future. There have been many different reasons for such reticence: courts have wanted to see how partial decisions fly before resolving further cases, higher courts may want to allow lower courts flexibility in addressing unexpected situations, legislatures may be unable to come to the needed political compromises without leaving 'blanks' for courts to fill in.

Tappenden is thinking of cases in which matters are intentionally left open, but we can imagine lots of reasons why aspects of word meaning might remain open as a kind of natural default state - it may simply be too complicated to specify everything (even for an expert) or it may be that crucial aspects of word meaning depend upon facts about the world that remain open. Or it may just be that the FLN is only accidentally suitable for communication and that for no reason in particular it just happens not to fix robust lexical meanings.

It would be a mistake, I think, to try an assimilate these cases of open meanings to that of vague predicates like 'bald'. Many of the disputes that arise have little to do with vagueness. Too see this, consider the following case from a dispute I heard on WFAN (a sports talk radio station in New York) when Sports Illustrated announced its "50 greatest athletes of the 20th Century". Some listeners called in complaining that a horse - Secretariat - had made the list, while host

Mike Francessa defended the choice. Clearly this is a dispute about what should be in the extension of 'athlete', and the callers wanted to argue that a horse had no place here. It is not as though the dispute would be resolved if Secretariat were a little bit faster or could throw a baseball, so it seems hard to imagine that these are vagueness cases.

This is also a good example of a case where fleshing out the meaning of the term is up to us and our communicative partners. So, even when we are deploying a common coin term (like 'athlete', for example) the extension of the term within a given context may be up for grabs and may require some form of coordination strategy - in the sports talk radio case the coordination took the form of a debate where discourse participants argued their respective cases.

At least in this narrow instance there is an obvious similarity to the legal realm, where competing parties may come together to resolve a dispute - in this case the way in which the term is to be understood with respect to the new cases in question (think of question of whether an existing patent "reads on" (applies to) some new technology). The key difference is that rather than taking place in a formal courtroom setting, these debates play out in less formal realms, ranging from sports talk radio to arguments with colleagues, friends, and partners.

Assigning meanings to common coins by jurisdiction. Tappenden's metaphor of court decisions can be extended in fruitful ways. Disputes over the best baseball player or whether a horse counts as an athlete are often just wheel spinning, but sometimes a consensus is achieved. This might be due to a series of rational arguments or it might be a simple case of someone asserting a claim and other participants deferring. In a bit we will look at how this kind of deference works, but first it is worth noting that when these disputes are resolved there are often jurisdictional limits.

When courts come to a decision on a particular dispute they set a precedent which may carry over into other jurisdictions. On the other hand it may not. Similarly, we may resolve a dispute or coordinate on the meaning of a term, and expect that to carry over into other micro-languages that we form. We may be disappointed to find we have to re-argue our point of view, or re-establish our credentials.

Alternatively, it may be that some of the disputes that we engage in (about sports, television, movies, and questions like "Is Chesner a smoker if he only smokes when drinking?") which appear trivial or silly are valuable precisely because they are litigating the content of certain key terms and this may be valuable in contexts where more is at stake and communication is critical. In other words, idle talk may well serve the function of helping us to calibrate our lexicons during periods of down time. These periods of calibration may be serve us well later when we later need to collaborate on some important project or problem.

Assigning meanings to common coins by deference. Sometimes we may not be involved in litigating the meaning of a term, but we may rather defer to the usage of someone (perhaps in the conversation, or perhaps in the greater community). To use a famous example from Hilary Putnam, we may defer to an expert on the proper individuating conditions of the terms 'beach' and 'elm'. There may be a social division of labor involved in fixing the semantic content of our utterances.

It is one thing to say that semantic deference takes place and quite another to explain how it works. Friend and Ludlow (2004) considered the thesis that deference-worthiness is earned discursively via a series of challenges. More precisely, we argued that expertise in a domain must be established via a series of interactive "partial knowledge proofs."

The phrase 'partial knowledge proof' is a riff on the notion of "zero knowledge proofs" in computer science (in particular in the field of public key cryptography). The basic idea of a partial knowledge proof is this: if I have a particular expertise, how can I prove to you that I have that expertise when it is something that you lack? To illustrate the idea, imagine a situation where we are hiring a philosopher in ancient philosophy but no one in the department is an expert in the area. We all have some knowledge of ancient philosophy, of course, but we are hiring in the area because we recognize we are not experts. We resolve this dilemma by issuing a series of challenges to the job candidate. With each question/answer exchange we learn more, allowing our colleagues to press on with deeper and more informed questions. In the end, via this interactive inductive proof procedure, we satisfy ourselves that the candidate is worthy. Or not.

Stacie Friend and I argued that this kind of procedure is more typical than one might think, applying even in cases like the meaning of the word 'cool' (in the social not the thermodynamic sense). We might think that Richie and Pottsie always blindly defer to Fonzie on the meaning of 'cool', but in fact there are times when challenges are issued, and there are at least person-internal debates about whether Fonzie is really the appropriate arbiter of the extension of the term. Fonzie's deference-worthiness is constantly subject to challenge, and may well be undermined as we encounter other arbiters of cool (as when Richie goes to college) or about Fonzie (as when he goes water skiing and jumps a penned up shark - definitely not cool).

It is an interesting question as to what counts in a decision to defer to Fonzie on the meaning of 'cool'. Presumably Richie and Pottsie had partial knowledge of the concept, and their deference is not tied to credentials possessed by Fonzie; Fonzie did not have a diploma from the College of Cool. In other cases, however, semantic deference does appear to be tied to credentials.

For example, one day a "tree guy" came to my house and while pruning some trees, identified the trees in my yard. Along the way he assured me he had gone to horticulture school. Did that provide him with the expertise to say what is a beech and what is an elm? Should I defer to him? Well, not much hung on the question, so I was perfectly happy to adopt his usage. For similar reasons I'm happy to defer to the doctor when she says I can't have arthritis in my thigh. But why do I defer?

Well, presumably it is not because these experts have pointy heads or impressive accents - it is because the credentials they hold (diplomas for example) that show they have been vetted by a kind of process not so different from the one we used to hire our Ancient philosopher - as students they were subject to an interactive inductive proof procedure which convinced their institutions that they had the relevant domain knowledge. It would be interesting to explore this process in more detail, when we turn to the semantics of word meaning a more pressing question arises: why does your domain expertise matter here?

The point of my question is that once domain expertise is established, the "semantic reach" of the domain expertise must be also established (e.g. should I defer to the materials scientist when she says that the glass in that window falls under the extension of 'liquid' in our conversation? Or is the materials scientist overreaching her jurisdiction when she asks us to adopt her linguistic usage?). In Ludlow and Friend (2004), we consider the thesis that this semantic reach must also be established discursively, via a series of challenges. But clearly this doesn't happen all the time. Nor should it.

Copycats. There is obviously a lot going on here - a combination of coordination strategies, litigation of deference and semantic reach, and the creative coining and fleshing out of lexical items. But in addition to all this there are presumably also cases of just flat out copying as well - cases where we blindly or at least indifferently adopt the linguistic practices of those around us, apparently for no reason at all. Well, maybe it is for no reason.

It is certainly the case the human agents are quite adept at simply doing as their neighbors do. In the United States someone started driving on the right hand side of the road, and everyone else followed suit. Similarly, we eat our pie with the wedge pointing towards us. It seems doubtful that anyone debated the issue. Joshua Epstein (2001) has shown that one can successfully model group political behavior with a population of cellular automata that basically just do what their neighbors do as long as no new agent comes along and violates conventions.

It is interesting to reflect on whether this behavior, hardwired or not, could count as being rational or normative in some sense. Surely some unreflective imitation must be warranted. It would certainly make for an interesting time if all conformity required pause for reflection. Quite apart from making driving an adventure, many of us would simply be paralyzed with indecision. One can easily imagine a semantic theory that established our warrant for reflexively following our neighbors when they introduce novel lexical items or when they offer precisifications and adaptations of those already in use. (The case for warrant would parallel recent work on perceptual warrant by philosophers like Burge (2003)).

There are moments however, when it is time to reflect on the imitative behavior we are engaged in, and this is true when we engaged in the assignment of word meanings no less than when we are engaged in the overtly political. Application conditions for terms and phrases like 'murder', 'life', 'family values', and 'good character' must be fleshed out and precisified, and it would seem to be a mistake to just blindly follow our neighbors or the powerful on their precisifications. Here is a place where we want to insist on deliberation and good reasons for a choice.

Power relations in lexical choice

Here is also the place where we see the role of power relations. It is not rare at all to find power relations in linguistic interactions, and they are reflected in everything from the pronunciation we adopt to the coining and precisification of lexical items. Sometimes discourse participants will attempt to "impose their will" on their discourse partners, even in the face of resistance, and sometimes we comply even when we don't like it. I suppose this happens in the classroom all

the time, but my favorite example comes from one of the entrainment experiments conducted by Susan Brennan (p.c).

In one of the entrainment experiments, subjects were asked to coordinate on the selection of terms for certain pictures that they would subsequently use when showing each other the pictures again later in the experiment. Transcripts of these experiments showed there to be little discussion of the choices - someone usually just called the shots. But who? And why did they get to call the shots?

In some cases the shot-caller appeared to have real world knowledge which they took to establish a kind of semantic authority (e.g. knowledge of cars). In other cases the transcripts are unilluminating, but one can imagine all sorts of factors ranging from relative age to dress to more discrete social factors. And, as I suggested, sometimes the deference relation is recognized but not appreciated. The case I have in mind involves an experimental subject that chose the word 'rice burner' to refer to a picture of a Japanese car. When the picture came up again the other subject complied with the coinage, but registered disapproval by referring to it as "your 'rice burner'."

Notice though, that we are never compelled to defer – we are never prisoners to our own or someone else's "language" as suggested by the following passage from Deleuze and Guattari, cited in Venuti (1995; 273).

How many people today live in a language this is not their own? Or no longer, or not yet, even know their own and know poorly the major language that they are forced to serve? This is the problem of immigrants, and especially of their children, the problem of minorities, the problem of a minor literature, but also a problem for all of us: how to tear a minor literature away from its own language, allowing it to challenge the language and making it follow a sober revolutionary path? How to become a nomad and an immigrant and a gypsy in relation to one's own language.

We cannot be imprisoned by something that does not exist. On the other hand we do need to be cautious in when and how we defer to the linguistic practices of our discourse partners, and we need to insist that semantic deference be paid only when it is warranted. More to the point, the proper response to attempts at regimenting our linguistic practice is not to retreat into infantile word games, but rather to exercise care and creativity in clarifying the terms we wish to deploy, explaining their usages, and making sure that our discourse partners understand and respect these usages. Care and clarity are the answer, not word salad.

3.0 An Illustration

My first job after I got my PhD in 1985, was not in academia, but working for the Intelligent Interface Systems Group of the Technology Strategies Center, run by the Honeywell Corporation. My first assignment was to study the then existent machine translation projects - an assignment that sent me traveling to research centers around the world. In those days, machine translation was crude, but in certain circumscribed contexts, it was economically viable

to have machines do rough drafts of certain documents. If necessary, the documents could be cleaned up by human translators.

Back then, my computer was an Apple II with 48K of ram, and the computers we used at the center (Symbolics Lisp Machines) had substantially less power than the low end laptops available for a few hundred dollars today. One might have thought that after 20 years significant advances in computing power we would also seen advances in machine translation and natural language "front ends" for data bases. But we haven't. And this is not the least bit surprising.

Most of the work on machine translation and natural language processing has been based on a mistake - the idea that one has to find an algorithm that can take some text in a "source language" and in one stroke translate the text into the "target language." But given that there are no languages, we can see that this is a confusion from the start.

The next time you go to a bank or a store with a particular request, think about the way your conversation plays out. Do you just make a request and receive an answer? How many times have you had to ask the teller or the clerk to clarify something? (The first time a bank clerk asked "Do you want that large?," I had no idea what she wanted to know.) How many times has the teller or clerk asked you to clarify what you need? How many times did you go back and forth with phrases like "sorry, did you mean…" or "I'm sorry, I didn't catch that" or "I'm not sure what it's called but I need something that…".

Interesting work has been done on the nature of conversational repair, even including the important role that is played by expressions like 'um', 'er', and 'hmmm' in communication – so called "disfluencies."   But even before we get to the role of these items we need to understand that the lexicon is not stable but must be recalibrated on a conversation-by-conversation basis. For example, in studying the way agents attempt to communicate with computers with natural language interfaces, Furnas et al. (1987) found that the likelihood that any two people would produce the same term for the same function ranged from only 7 to 18%. For example, when wishing to remove a file, persons used a broad range of terms including remove, delete, erase, expunge, kill, omit, destroy, lose, change, rid, and trash. You might think you could get around this problem by treating these terms as synonyms and having the system regard any of them as an instruction to delete a file, but as Furnas et al. discovered, even with as many as 20 synonyms for a single function, the likelihood of people generating terms from the synonym set for a given function was only about 80%. And then a new problem is generated. When people do use the same term, more likely than not they don't mean the same thing by the term. As Furnas et al. showed, even in a text editor with only 25 commands, if two people use the same verbal command, the chances that they intend same function by it was only 15%.

In the light of these considerations think about how silly it is to try and build a machine that "just understands you" when you walk up and begin talking to it. No human can "just understand you" and no machine will ever be able to do it - such a machine is a fantasy machine designed around the myth of language. We don't speak languages, so if machines did they would be no use in communicating with us anyway. If someone created a "perfect language" we would have no use for it.

This point can be extended to human translation as well. What is it that translators do? They surely don't translate from one fixed "source language" to another fixed "target language." To see this consider the situation faced by two Serbian friends of mine who are translators from "English" into "Serbian". One was translating Tolkein, the other was translating The Color of Purple. Exactly how does one translate Elvish expressions or Rural Black English Vernacular into "Serbian" (one common and very unhappy strategy in Serbia is to translate Black English Vernacular into "Bosnian"). In point of fact translators are not in the business of translating so much extending and morphing the "target language" so as to communicate the ideas in the book. Pannwitz (1917) had an interesting insight on this score:

The fundamental error of the translator is that he stabilizes the state in which his own language happens to find itself instead of allowing his language to be powerfully jolted by the foreign language. (from Venuti 1995; 148)

Of course on my view, it is not that the translator's language is changing so much as the translator is establishing a micro-language with the readers of the so-called translation. Direct coordination is out of the question, but assumptions about the knowledge and background of the audience will direct the way in which the microlanguage is constructed.

4. Conclusion

The real mystery in all of this is the question of why persons find the myth of language so persuasive when they have so much evidence against it. To sustain the myth they must ignore the dynamical nature of human communication, the widespread coining and reissue of novel terminology, the pervasive divergence in pronunciation and meaning, the repair strategies, the debates over and refinements of meaning and the alternating deference and conflicts over meaning, and still cling to the idea that there is a thing there – a language – that is helping us to communicate. One wonders why everyone doesn't come away with the impression expressed by James Joyce's in Finnegan's Wake:

Because, Soferim Bebel, if it goes to that… every person, place and thing in the chaosmos of Alle anyway connected with the gobblydumbed turkey was moving and changing every part of the time: the traveling inkhorn (possibly pot), the hare and the turtle pen and paper, the continually more or less intermisunderstanding minds of the anticollaborators, the as time went on as it will variously inflected, differently pronounced, otherwise spelled, changeably meaning vocable scriptsigns. No, so help me Petault, it is not a miseffectual whyacinthinous riot of blots and blurs and bars and balls and hoops and wriggles and juxtaposed jottings linked by spurts of speed: it only looks as like it as damn it.

Perhaps the problem is that we just cannot see how communication could emerge from this riot of blots and blurs and balls and hoops and wriggles; but somehow it does. If we are interested in how it does, then we need to retire the standard picture and begin investigating the nature of our linguistic coordination strategies, their origin, and use, and then we need to get about the business of rethinking those enterprises of ours that continue to rest upon the myth of human language.

Bibliography

Baker, M., 1988. Incorporation: a theory of grammatical function changing. Chicago: University of Chicago Press.

Bloom, P., 2000. How Children Learn the Meanings of Words. Cambridge: MIT Press.

Brennan, S. E., fothcoming. "The vocabulary problem in spoken language systems." In S. Luperfoy (Ed.), Automated spoken dialog systems. Cambridge: MIT Press.

Brennan, S. E. & Clark, H. H., 1996. "Conceptual pacts and lexical choice in conversation." Journal of Experimental Psychology: Learning, Memory, and Cognition 22:1482-93.

Brennan, S., and M. Schober, 1999. "Uhs and interrupted words: The information available to listeners". In Proceedings of the 14th International Congress of Phonetic Sciences, Satellite Meeting on Disfluency in Spontaneous Speech (pp. 19-22), Berkeley, CA.

Burge, T., 2003. "Perceptual Entitlement." Philosophy and Phenomenological Research 67, pp. 503-548.

Chomsky, N., 1980. Rules and Representations. New York: Columbia University Press.

Chomsky, N., 1986. Knowledge of Language. New York: Praeger.

Chomsky, N., 1994b. "Naturalism and Dualism in the Study of Language and Mind." International Journal of Philosophical Studies 2, 181-209.

Clark, H., 1992. Arenas of Language Use. Chicago: CSLI Publications and The University of Chicago Press.

Davidson, D., 1986. "A Nice Derangement of Epitaphs." In E. Lepore (ed.), Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson. Oxford: Blackwell, 433-446.

Deleuze, G., and F. Guattari, 1986. Kafka: Toward a Minor Literature, trans. By D. Polan. Minneapolis: University of Minnesota Press.

Epstein, J., 2001. "Learning to be Thoughtless: Social Norms and Individual Computation," Computational Economics 18, 9-24.

Furnas, G., T. Landauer, L. Gomez, L. M., & S. Dumais, 1987. "The vocabulary problem in human-system communications." Communications of the ACM, 30, 964-971.

Garrod, S., and A. Anderson, 1987. "Saying What Your Mean in Dialogue: A Study in Conceptual and Semantic Co-ordination." Cognition 27, 181-218.

Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). "The Faculty of Language: What Is It, Who Has It, and How Did It Evolve?" Science, 298:1569--1579.

Higginbotham, J., 1989. "Knowledge of Reference." In Alexander George, ed., Reflections on Chomsky, pp. 153-174. Oxford: Blackwell.

Lewis, D., 1969. Convention. Cambridge: Harvard University Press.

Ludlow, Peter, and Stacie Friend. "Disagreement and Deference: is Diversity of Opinion a Precondition for Thought?" Philosophical Perspectives 17: Language and Philosophical Linguistics, 2003, 115-139.

Schober, M., & H. Clark, 1989. "Understanding by addressees and overhearers." Cognitive Psychology 21, 211-232. Reprinted as chapter 6 of Clark (1992).

Stewart, I., 1998. Life's Other Secret. New York: John Wiley and Sons.

Tappenden, J., 1999. "Negation, Denial and Language Change in Philosophical Logic" In D. Gabbay and H. Wansing, (eds.) What is Negation? Dordrecht: Kluwer.

van Deemter, K. and S. Peters (eds.), 1997. Ambiguity and Underspecification. Stanford: CSLI Publications."
Venuti, L., 1995. The Translator's Invisibility: a History of Translation. London: Routledge.
Watts, D., 1998. Small Worlds: The Dynamics of Networks between Order and Randomness. Princeton: Princeton University Press.