An embodied grammar of words[1]

This paper presents the merits of integrating a science of language and a science of cognition and behavior. The science of language is empirical linguistics as developed by Zellig Harris and his students and colleagues. The science of cognition and behavior is Perceptual Control Theory (PCT), as developed by William T. Powers and his colleagues and students. Each establishes an empirically and mathematically sound foundation. Harris's theory of language and information is emergentist and amenable to dynamical and connectionist modeling. PCT is dynamical and connectionist. In neuroscience modeling, dynamical and connectionist approaches have supplanted the symbolic computationalism of classical Cognitive Psychology and its appendant field, Generative Linguistics.

Key words: Empirical linguistics, formal linguistics, Zellig Harris, control theory, negative feedback control, perceptual hierarchy, embodied cognition, dynamical systems, connectionism, neuroscience, objective information, subjective meaning, semantics, statistical learning, learning, language acquisition

I assume that the intended readers of this volume are familiar with Harris's empirical linguistics[2,3] and with Operator Grammar. For those who are not, there follows a very brief summary of features most relevant to the integration with PCT. However, because reviewers in the past have usually interpreted Operator Grammar in terms of the formalisms of Generative Linguistics, I have thought it best to begin with an account of how radically the latter has diverged.

Empirical Linguistics

Historically, the divergence began with a reinterpretation of what a transformation is. Harris was applying a term from mathematics which means a mapping from one set to another. (More on this below.) Disregarding Harris's definition, Chomsky took the term *transformation* from its usage in the formal syntax of logic (e.g. Carnap 1937, Rosenbloom 1950).[4] In this very different conception,[5] *rules of formation* generate simple propositions, and *rules of transformation* are used to infer more complex propositions. Chomsky's rules of formation generate phrase classes in a pseudo-hierarchy that terminates with word classes. This grammar of phrase classes and word classes, called Phrase-Structure Grammar (PSG), is a formalization of the immediate-constituent analysis pioneered by Leonard Bloomfield, implemented with a symbol-rewriting system of the type that Emil Post had invented (e.g. Post 1943).[6] Chomsky's rules of transformation were defined as deformation operations on the phrase structures generated by these rules of formation or by rules of transformation.

At first, the structures of PSG corresponded to actual phrases and sentences, and Generative transformational rules were merely restatements of the transformational relationships disclosed by

---

[2] This is one of two papers concurrently drafted. The other (forthcoming) was for the converse audience: readers unfamiliar with Harris's empirical linguistics but knowing or learning about PCT. Except where otherwise noted, figures are adapted from those in Nevin (forthcoming). For very helpful comments on both, I am indebted to Stephen B. Johnson, John Goldsmith, and Michael Gottfried from the point of view of linguistics, and to Kent McClelland, Martin Taylor, and Warren Mansell from the PCT perspective. Errors and oversights of course are my responsibility alone.

[3] Harris used the term formal linguistics at least as early as 1960. However, the association of 'formal' with 'formalism' invites confusion with a prevalent definition of linguistics as the study of systems of symbolic instructions intended to generate language. (Harrisian) empirical linguistics has always assumed corpus linguistics, as in Sampson (1980).

[4] He took every course offered by logician Richard M. Martin. Ironically, Martin esteemed the work of Harris but was later dismissive of Chomsky's Generative grammar (Meguire 2004).

[5] As Chomsky (1975:43) acknowledged.

[6] Chomsky studied Rosenbloom (1950), chapter 4 of which contains a nice summary of Post's work.

Harris's methodology,[7] but over time in order to make the transformational rules work the base structures and intermediate derived structures have become more and more complex and abstracted from the observed data of language. There is a recurrent problem of overgeneration, in which the systems of rules predict structures that do not actually occur in languages, and these exceptions require development of meta-rules, called constraints. A large family of these called island constraints apply to movement rules. In the most recent versions of the theory, a basic rule is move-α, with the effect that anything can be moved anywhere. (In Operator Grammar movement operations are few and very limited in scope, and these island phenomena are not problematic.) Because of this abstractness and complexity, and the scarcity of examples in adult speech from which a child might learn the exceptions to the rules, it is assumed that a child is able to 'acquire' language only with the assistance of a biologically inherited language module in the brain. The primary aim of Generative Linguistics is to characterize the properties of this 'language faculty'. These properties of the 'language faculty' are collectively referred to as Universal Grammar (UG). Empirical linguistics has no need of this postulate.[8]

Generative linguists assume that their symbolic formalisms characterize the essential nature of language. This conforms to, and is taken to confirm, the core metaphor of Cognitive psychology, which likens the functioning of the brain to information processing in a digital computer. "The standard view of classical cognitive science stated that cognition consists in the manipulation of language-like structures according to formal rules. Since cognition is 'linguistic' in itself, according to this view language is just a complex communication system and does not influence cognitive processes in any substantial way."[9] It is important to recognize the differences between language and the kind of 'language' that is generated by a formalism, as for example a computer programming language. Language varies in social and geographic space and changes through time, and the changes are contingent upon these variants, emerging from them by neo-Darwinist processes of blind variation and selective retention (Cziko 1995, 2000); a formal language is invariant in accord with the formal statements that define it, and changes only by redefinition. Language contains its own metalanguages; the metalanguage of a formal language is the external and antecedent means of defining it. Operations in language are contiguous; operations in a formal language make use of metalinguistic indices such as subscripts to coordinate operations over arbitrary spans. The structure of language (as disclosed by empirical linguistics) constitutes linguistic information; the semantic interpretation of a formal language must be provided by a separate symbol system which is ultimately dependent upon language (and which requires its own interpretation in turn).

For computability by algorithms in a digital computer, a grammar must be formalized. However, these essential differences, together with the proposed integration with PCT, call in question the value of measures of computability and conclusions about learnability that are derived from properties of symbolic formalisms. The divergence outlined above affects the definition of a science of language. Generative linguistics is the investigation of formalisms in the faith that language may be projected from the correct formalism, where linguistic data has the form of anecdotal examples and counter-examples. Empirical linguistics is the investigation of relationships in the data of language across a corpus with broad coverage.

We will now present the essential features of empirical linguistics in the form of a progression from a grammar of word-classes[10] to a grammar of words.

---

[7] Harris was producing a series of "structural restatements" of grammars written by other linguists when Chomsky became his student. He may have concluded that this is how one does linguistics. His Master's thesis (Chomsky 1951) is a formal restatement of materials on the morphophonemics of Modern Hebrew that Harris gave to him (Barsky 2007:148, Nevin 2010:117-118), and refers to morphophonemic rules as 'transformations'.

[8] It is a postulate and not an hypothesis because without it one is not doing Generative Linguistics. With Laplace, Harris might say "je n'ai pas eu besoin de cette hypothèse."

[9] Mirolli & Parisi (2009:517).

[10] I omit the grammar of sentence-expansions (Harris 1946), later reformulated as adjunction grammar (Harris 1963) and tree-adjoining grammar (Joshi 1985, 1999). These all avoid the ungovernable abstraction of PSG, and handles exocentric constructions properly. Adjunction misrepresents the order of adjectives and other modifiers, because adjuncts enter adjacent to their head. Reduction of conjuncts to adjuncts resolves this.

In linear algebra, a transformation is a mapping from one set to another, conserving some property. A linguistic transformation in this (original) sense is a mapping from subset to subset in the set of sentences. Each subset of sentences is indicated by a sequence of word-class names, called a sentence-form, such as $N_1 \, t \, V \, N_2$. The property that is conserved is word selection. Methodologically, this becomes the criterion for testing transformations experimentally by substituting different words as satisfiers of a pair of sentence-forms. (The core methodology of empirical linguistics at all levels of analysis is substitution tests. These are a linguistic specialization of the core methodology of PCT, the Test for the Controlled Variable.) A difference of acceptability between any two satisfiers of a given sentence-form also obtains between the corresponding two satisfiers of its transform.[11] Thus, the difference between *Alice ate chocolate* and *Vacuum purchased kelp* (satisfiers of $N_1 \, t \, V \, N_2$) is also seen between *Chocolate is what Alice ate* and *Kelp is what vacuum purchased* (satisfiers of $N_2 \, t \, be \, wh\text{-}Pro_2 \, N_1 \, t \, V$). This is not a test of the absolute acceptabilities of the sentences, which would be difficult to assess, but only a test of a parallel in how they differ in acceptability.

The next step toward a grammar of words was to factor these transformations (mappings) into simpler transformations, elementary sentence-differences. The incremental transformations add words, and the paraphrastic transformations change the phonemic shapes of words, and in certain limited cases their linear order.

Then in a parallel to the development in mathematics of operator theory from linear algebra, the sentence-differences were organized as Operator Grammar. Words are classified according to their argument requirement. The class N has a null requirement, and the various classes of operators O require some combination of N and O words in their argument, for example, O_n (*fall*, *be green*, *be a bird*), O_{nn} (*see*, *be sister of*), O_{no} (*expect*, *be aware*), and so on. As each operator enters into the construction of a sentence, the operator and its argument words are placed in a linear order (carrying with them the arguments of arguments, etc.). All such combinations are possible, but some are more restricted in acceptability than others.

As an entering operator introduces new word dependencies and new string relationships of words, the presence of a given word may become more predictable, with the effect of contributing less information to the sentence. As soon as this happens, the phonemic shape of such a word may be reduced, even to zero. These changes are called reductions. Harris referred to reductions as extended morphophonemics. In the relatively few cases in which a reduction is not optional, it is because it has become institutionalized and is very strongly expected as a convention of language use.

The main conditions for reduction are of two kinds of relationships between like words, dependency relationships and string relationships. A typical dependency relationship that enables reduction is parallel position under conjunction: *He came and went ↔ He came and he went*. Two kinds of string relationships between words are contiguity and proximity. Contiguity enables exact addressing of an antecedent. A repeated word may be reduced if it is contiguous to its antecedent and identified with it by a metalanguage assertion, as for example in the derivation of a relative clause: *John—John (same as prior) you met yesterday—borrowed a book ↔ John, who you met yesterday, borrowed a book*. Because the second argument of *meet* can be linearized at the front of the interrupting sentence (*John you met yesterday*), the two occurrences of *John* are contiguous; and indeed only those things which can occur at the front of a sentence can be relativized.[12] Subscript indices to show co-referentiality, as in a logical formalism, are not required. On the other hand, the unaddressed pronouns such as *he* and *she* require only that their antecedent be mentioned nearby (Harris 1982:127-135). Subscript indices are too precise in these cases where language requires only proximity.

The objective information in language is in the dependency relationships of words, primarily the operator-argument dependencies. Another form of dependency is repetition, and this includes repetition of words between sentences as well as within them. Harris recognized that a connective operator O_{oo} or O_{noo} requires that at least one word be repeated in its two arguments. This word-sharing

---

[11] More exactly, the difference-ranking is not reordered. Under some operators, such as *dream*, *imagine*, differences may be collapsed.

[12] This resolves a great many of the so-called island phenomena. There is no movement rule to constrain.

requirement enables the grammar to be regularized, and explains certain anomalies. Where a given conjoined sentence does not show this repetition overtly, there is justification for reconstructing an intervening conjunct under *and* which has been reduced to zero because it is obvious common knowledge. For example, we don't say *It might rain, and an umbrella shelters one from rain, so I'll take an umbrella*, we just say simply *It might rain, so I'll take an umbrella*. The very obvious intervening sentence might occur in a dictionary or an encyclopedia, but is not ordinarily spoken.

There are further regularities of word repetition across the sentences in a coherent discourse. Discourse analysis, in the form developed by Harris, uses distributional methods to establish equivalence-classes in a text (Harris 1952a,b, 1963). Texts in a restricted subject-matter domain have the same equivalence-classes in common (Kittredge & Lehrberger 1982, Grishman & Kittredge 1986, Harris et al. 1989)[13]. Some of the classifier words that are used to state definitions within the domain may also be used as names for the equivalence-classes, and therefore may also serve as the grammatical categories of the sublanguage, although algebraic symbols have more commonly been used. For example, in the pharmacology sentence *digitalis increases the beating of the heart* the first word (*digitalis*) is in the *DRUG* class, *increase* is a verb in the *AFFECT* class (which includes verbs like *affect* and *suppress*) and the phrase *the beating of the heart* is in the *SYMPTOM* class; the sublanguage sentence-form may be written as *DRUG AFFECTS SYMPTOM*. A member of an equivalence-class may be a phrase as well as a single word. The phrase *the beating of the heart* functions as a single 'word' in this sublanguage, whereas in physiology (an epistemologically prior science) *beat* and *heart* are separate words, as indeed they are in general usage.[14] In a sublanguage of healthcare, *cancerous growth* is a lexical item in the *INDICATION* class, which may be either a *SIGN* (observed by a medical professional) or a *SYMPTOM* (reported by a patient).[15] In the immunology sentence *nodes on the side injected with paratyphoid bacerin became slightly larger*, the phrase *became slightly larger* is in the *ENLARGE* family of histological operators.[16] Only certain sequences of sublanguage equivalence-class categories occur in texts that are recognized as belonging to the domain, and a sublanguage grammar for the domain specifies those class sequences.

In the development of a discourse, after a given word dependency occurs, there is increased likelihood of recurrence of a dependency between those words, or between classifiers of those words, or between words that can be arguments of those classifiers. Sentences with classifier words often occur in dictionary and encyclopedic sentences for the given domain.

Meaning is obviously not inherent in words. There is nothing in the sound of *cheval*, *Pferd*, *jahhóm*, *mǎ*, *horse* that associates these words with nonlanguage perceptions of large, four-hooved animals in the brains of French, German, Achumawi, Chinese, and English speakers, respectively. Maurice Gross (1994:213) proposed that "the linguistic unit of meaning is the elementary sentence." This can be understood as saying that meaning is in the word dependencies, and can be attributed to a word only when the word is part of an assertion. Gross confirmed for French verbs the conjecture by Henry Hiż (1961) that every word participates in a unique set of dependencies and reductions (Gross 1968, 1975, 1994). He also confirmed that a very high proportion of sentences have meaning that is not a compositional product of the meanings of the words in them. Compound words, idioms, and metaphors are the most obvious examples of non-compositional semantics. Such a sequence of words satisfies the operator-argument requirements of Operator Grammar, and can be graphed in a dependency tree like those in Figure 8, but nevertheless semantically (that is, distributionally) is an undecomposable lexical unit.[17]

---

[13] Les etiquettes *I* et *O* ici divent des mots Anglais *Input* et *Output*.

[14] Harris (1991.282).

[15] Sager et al. (1987) labels these members of the H-INDIC class in their computer implementation of a healthcare sublanguage grammar.

[16] In the information formats for the immunology sublanguage this is labeled the W class of histological operators, specifically those designated W_g. The example is adapted from sentence 792.1.2 of (Harris et al. 1989:134, 232).

[17] This is very difficult or impossible for symbolic rule systems to capture (Gross 1979). And as we have seen, this status (unitary vs. compositional) may vary, depending upon the subject-matter of the context.

The technical term for a dependency subtree regarded as a unit is *catena* (O'Grady 1998, Osborne, Putnam, & Groß 2012). A recurrent catena may become lexicalized, as though it were a single compound argument in Operator Grammar, and may be particularly subject to reduction. The operators and arguments within the unreduced catena retain a broader distribution, while the distribution of the lexicalized reduced form may become relatively restricted and so drift semantically from the etymological meaning of the unreduced catena. In English, *pro-life* is clearly two morphemes, the first of which can be seen as an alternate form of the preposition *for*; but in *proceed* the relation of the prefix is to *fore*, and in *profession* it is doubtful that the etymology has any current relevance. This is therefore an active locus of variation and change. Within a language, lexicalization of compound expressions can be seen to vary on a continuum of language use, and languages differ greatly as to which catenae are lexicalized and in the extent of reduction and lexicalization in compound words. The equivalence-classes that we see in the double array of discourse analysis identify phrases which function as lexical items in the grammar of each specific discourse. Their specification is sharpened in a survey of many discourses that are devoted to the same restricted subject-matter domain, and they become lexical items in the sublanguage grammar for that domain.

Linguistic information as disclosed by empirical linguistics is not the whole of semantics. One motivation for the integration of empirical linguistics with PCT is to explain how language users associate subjective meanings with the objective information that language transmits. A broader motivation is to provide a sound psychological account of the origin, development, and learning of language. A clear and explicit, testable model of language is essential for a PCT model of human behavior and interaction in social psychology, sociology, anthropology, politics, economics, history, and many other fields.

Perceptual Control Theory (PCT)

What promises to make this an embodied grammar is the theory of embodied cognition[18] known as Perceptual Control Theory.

Living things have a unique characteristic that differentiates them from non-living things. Consequently, the scientific study of living things presents demands, methods, and opportunities that differ from those of the physical sciences such as physics, chemistry, geology, and astronomy.

This unique characteristic of living things is called control. If you place a ball on an inclined plane, as Galileo did when he refuted Aristotle's theory of falling bodies, the ball invariably rolls downhill. The ball can do nothing about it. If you place a mouse, a pigeon, or a human on an inclined plane, what the living organism does there depends not just upon gravity, inertia, and friction but also upon what it wants. What happens depends upon the effects of its own actions combined with the effects of those same physical forces that determine the path of the ball. We living things are different from non-living things because we act so as to resist or evade influences that affect perceived aspects of bodies or our environment in ways that we do not wish. More generally, we living things have preferences, and we use energy metabolized from nutrients to resist unpredictable influences in our environment that would otherwise make our perceived experience deviate from those preferences. For this reason, the sciences of life require methods that do not apply to the physical sciences. This section will describe the methodology of Perceptual Control Theory.

James, Dewey, and others were on the right track around the turn of the 20th century when they spoke of behavior as the attainment of a fixed goal by variable means, but their ideas were dismissed in the rise of behaviorism. How could present actions be determined by a goal, a future state that does not yet exist? In the effort to be scientific in linear cause-and-effect terms that are appropriate for the physical sciences, behaviorists asserted the dogma that stimuli in the environment cause an organism to 'emit' or 'produce' responses. In the 1950s and 1960s, the Cognitive Revolution (Gardner 1985) made much of superseding the stimulus-response theory of Behaviorism. However, Cognitive Psychology differs from Behaviorism only by the vague proposal that between the stimulus and the response the brain performs 'information processing' computations. Events in the environment are still regarded as

---

[18] Wilson & Foglia (2011) provides an overview of proposals about embodied cognition. They do not mention PCT.

stimuli, with a linear, sequential chain of cause and effect leading from stimuli into the brain and out again to behavior. In this view, sensory organs send signals to the brain, the brain does some 'information processing' to integrate these signals with cognitive maps of the environment, and then issues commands to the muscles.

The methodology of Cognitive Psychology has also changed very little in its fundamentals (Marken 2009, 2014). Conventional methodology of psychology is still modeled after the IV-DV methodology of the physical sciences. All the many interdependent variables in an experimental situation are supposed to be held constant except for two. The independent variable (IV), called the stimulus, is experimentally manipulated in a measured way, and changes in the remaining dependent variable (DV), called the response, are also measured. Because the responses are variable ("the attainment of a fixed goal by variable means"), a statistical correlation is sought between the IV (stimuli) and the DV (responses) over many experimental runs. But the statistical correlations are at a level that would never be accepted in the physical sciences. Anything better than a coin-toss is generally considered significant in the psychological and social sciences. Ah well, they say, that's just how it is with the behavior of living things. But this is not a science of behavior. Rather than a science, it resembles a cargo cult, emulating the methods of the physical sciences with the hope that the successes of the physical sciences will be delivered. (There is some merit as an investigation of how to influence the goals of behavior, which we will consider below, but not as an explanation of behavior.)

An abundant literature demonstrates that behavior does not correlate with stimuli. Some of this literature is reprinted in Powers (1989, 1992) and Marken (1992, 2002, 2014). It is not difficult to demonstrate this for yourself, for example using the interactive demonstrations that are provided with Powers (2008).[19] In each of those demonstrations you will be asked to control a perceptual variable, and you will experience for yourself how the only perceptual input that you could conceive of as a 'stimulus'—the difference between the state of that variable and the state that you want it to be in—is as close to zero as you can make it, even as your behavioral efforts are varying considerably in order to make it so. You will see that the variations in your behavior correlate inversely with a variable that you do not perceive at all. Your actions counteract an unseen and unpredictable disturbance to the state of the perceptual variable that you are controlling. This disturbance cannot be a 'stimulus' that causes your actions, because your very actions conceal it from you.

The gist of the matter is that living things do not control their behavior. Rather, behavior is an organism's means of controlling perceptions that matter to it. The behavior varies in whatever way is needed to make a perceived state of affairs be the way the organism wants it to be. We *use* our actions as the means to add our influences to those of the environment in order to perceive a desired result. The desired perceptual result is called the Controlled Variable (CV).[20] To control a perception is to make it conform to an internally maintained preference or standard for it (which itself may vary). Any living thing controls perceptual variables that matter to it by varying its actions—by varying the behavioral outputs that influence the state of the CV in accord with the properties of the environment. The CV is also exposed to other unpredictable and largely unperceived influences in the environment. These, we call disturbances to the CV. As unpredictable disturbances in the environment make the input perception deviate from the internal standard, the actions of the organism continuously act precisely so as to counter that influence.

Figure 1 shows these relationships in a simple control loop diagram. The dotted horizontal line is the demarcation between the nervous system and the environment. Perceptual organs receive feedback from the environment. From this feedback the perceptual input function *I* (comprising sensors and peripheral nervous system) constructs the perceptual signal *p*. The reference signal *r* enters the comparator *C* from above as an excitatory signal, and the perceptual feedback *p* enters from below as an inhibitory signal (hence, *negative* feedback). Each of these signals is a rate of firing that can be

---

[19] The interactive programs are freely available at http://www.billpct.org/. The book explains the phenomenon of control and provides tutorials for each program. Other interactive demonstrations may be accessed at http://www.mindreadings.com/demos.htm and http://www.mindreadings.com/javademos.htm.

[20] Do not confuse this with the IV-DV notion of a "control variable", one which does not change during the course of an experiment. In yet another usage, a control variable in a computer program is analogous to the reference value (setpoint) in a control system.

modeled quantitatively. The effect of combining the two signals is to subtract the perceptual feedback quantity from the reference quantity. The difference *r-p* is output as an error signal *e*. The output function *O* transforms any change in *e* to a change in actions influencing the state of the controlled variable *CV*. One or more disturbances *d* out in the environment are also influencing the state of the *CV*. The control system maintains an equivalence $p \cong r$ by varying its behavioral output so as to exert an influence on CV that opposes the influence of *d*. The maximum amount of deviation of *p* from *r* depends upon loop gain. Loop gain is the sum of amplification factors in *O* and elsewhere in the loop.
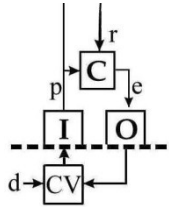

Figure 1. A simple control loop diagram.

This is not a stepwise sequence, as in the TOTE scheme proposed for Cognitive Psychology (Miller et al. 1960). The phases of the control loop—perceptual input, comparison with reference, error signal, action output, effect on the CV, perceptual input—are all concurrently continuous. The state of each variable in the loop is concurrently both the effect of the prior and the cause of the next. Transmission lags through the nervous system and through the environment affect only the time it takes to regain control in the face of a swift-acting disturbance, they have no material effect on the real-time phenomenon of ongoing control. The feedback loop is a whirling wheel of circular causation (Figure 2) which the disturbance *d* attempts to deflect, but it cannot because the continuous transformation of each variable into the next maintains the equivalence $p \cong r$.
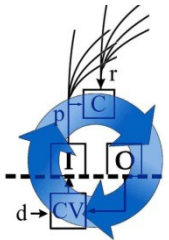

Figure 2. The loop is a whirling wheel of continuous circular causation with two inputs, the reference r *and the disturbance* d*. Viewed this way, its only output is the perceptual signal* p *passed up to higher levels, because behavioral outputs from* O *are part of the 'wheel'.*

This simple diagram maps a single negative-feedback control loop. Familiar examples include a thermostat maintaining your room temperature and a cruise-control system maintaining the speed of your car. Living things are more complex. To understand how you accomplish something as seemingly trivial as maintaining an upright posture while reading this, instead of collapsing into the collection of sticks and elastic bands that physics says you are, requires combining simple control loops into the perceptual hierarchy. The perceptual signal *p* continues upward and branches to the perceptual input functions of higher-level systems, and the strength of the reference signal *r* is proportional to the error output *e* from one or more[21] systems that control perceptions at higher levels of the perceptual hierarchy.[22] The strength of *r* specifies how much of the referenced perception is requested, that is, how strong the perceptual signal *p* should be that is constructed by the perceptual input function.

In a PCT model of imagination, a copy of the reference signal branches back so as to contribute to the perceptual input that is sent up to higher systems. For clarity, Figure 3 represents the input function I

---

[21] Internal conflict occurs when two higher-level systems control disparate values of the same perceptual signal *p* at the same time. The application of PCT to psychotherapy, called the Method of Levels, enables the subject to redirect attention away from the conflicted variable so that resolution of the conflict becomes possible (Carey2006, 2008). There may also be more than one source of disturbances *d*.

[22] There is no infinite regress of levels; see discussion of the reorganization system in Powers (1989:287-293), Runkel (2003:227-242), Powers (1995).

and output function O as small unlabeled circles. A third circle above the comparator C represents the reference input function that creates the reference signal *r* from a weighted sum of higher-level output signals. Two branches of the reference signal *r* descend from the reference function. The main branch sends an excitatory signal to the comparator as in Figures 1 and 2. The other branch adds an unamplified copy of *r* to the perceptual input signal above where it enters the comparator C.
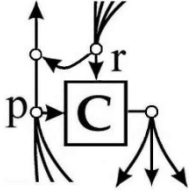

Figure 3. Detail showing branch of reference signal merging with perceptual input.

When there is no input from a lower level, then this relatively weak copy of the reference signal *r* is the only perceptual input that is sent upward. This is how we control aspects of a perception that are not directly accessible to the senses (the other side of the box that we reach to grasp, the dog attached to the tail that we see through the doorway). When there is perceptual input from lower levels, then the contribution of this signal is overwhelmed by the input from lower levels.[23]

Figure 4 shows these connections in context of several levels of the control hierarchy.[24] The connections on the right-hand side of Figure 4 do not reach all the way to the environment, but instead control in imagination at different levels of the hierarchy. If the absence of sensory input is at a relatively high level, the perceptions that are controlled are relatively abstract and devoid of detail, the stuff of thinking and planning. Hallucinating and dreaming are more vivid because signals from relatively lower levels provide more detail in the experience.

---

[23] Powers (2005:223) proposes an active switching process. This is schematically clear, but the dynamism of it seems to me implausible. Research will eventually decide between an analog model based on relative signal strengths or a discrete-state switch-flipping model that depends upon the ability of neuron motility to make and break these connections with precise appropriateness and timing.

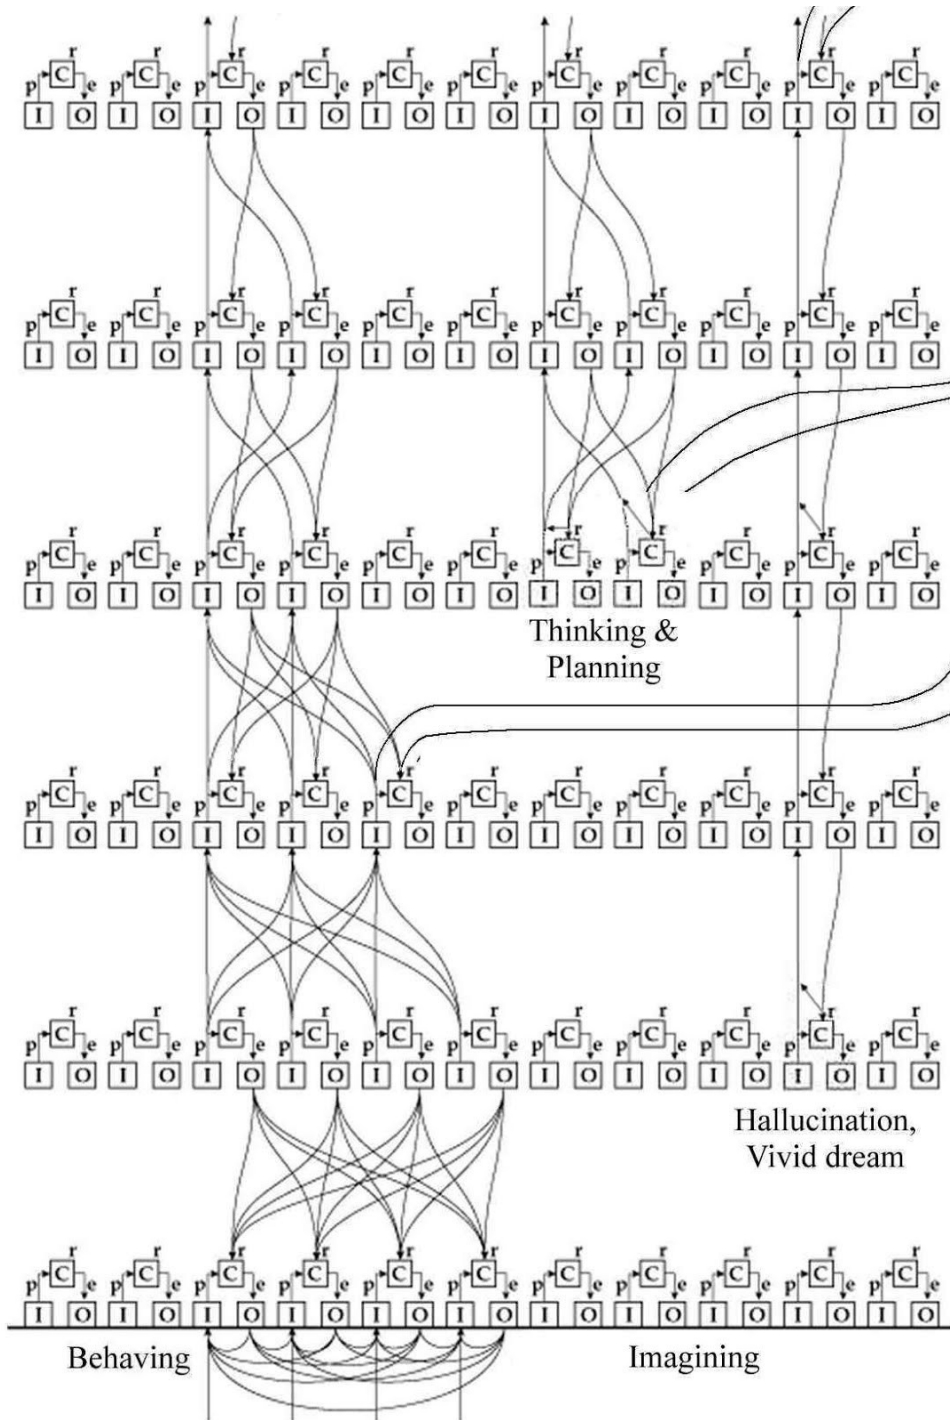[24] Adapted from the figure at (Powers 1988:278, Runkel 2003:196).

Figure 4. Connections and types of control in a hierarchical control system. The top three levels (Sequences, Programs, and Systems Concepts) are not shown

This model predicts that every perception includes some elements that are imagined. 'Live' input from lower levels, when it is available, makes an overwhelmingly stronger contribution to perceptions. However, when input is incomplete, or when input from one sensory modality is inconsistent with input from another, imagination plays a stronger role. This is the basis of perceptual illusions. Confirmation bias (or "myside bias"), the tendency to favor information or interpretations of experience that confirm one's existing beliefs and disbeliefs,[25] is not often thought of as a perceptual illusion, but it is indeed.

---

[25] Nickerson (1998).

A perceptual illusion that is especially pertinent here is the McGurk effect. A perception of what syllable or word a person is speaking is constructed by combining perceptual input from several sensory modalities. So-called lip-reading is more important than many of us realize. If a video is made of a speaker producing an utterance and it is is dubbed with a sound-recording in which one phoneme is changed, the brain overrides the actual acoustic input with a different auditory perception. This is why infants require visual perception of people's faces to learn language. This is discussed in more detail in the section on language perceptions, below.

Obviously, the strength of the reference signal is determined at higher levels, but other associative links can provide input to an imagined perception. Memory is a recreated neural signal. Just as perceptions are constructed, memories are also constructed. It appears that memory is local to each synapse.[26] Neurochemicals at the synapses sustain the neural signals that constitute short-term memories. To establish a long-term memory, firing of the synapse alters gene expression in the cell nucleus of the neuron, affecting protein synthesis mediated by RNA. (Neurotransmitters are membrane transport proteins.) This, too, is local to each synapse. A perceptual input signal from below or an output signal from above can evoke a memory by retracing the neural pathways in which the remembered perceptual signal originated. The growth and maintenance of new branches and synaptic terminals, which is the general means for amplifying a neural signal, strengthens memory and makes it more persistant (Kandel 2006). This process is called memory consolidation, and if it is interrupted long-term memory may not form. Conversely, an established memory can change in a process called reconsolidation (Loftus 1998, 1999).

Thus, when it is said that memory of a certain kind is located in a certain part of the brain, it apparently means that systems for recognizing and controlling that kind of perception, or important components of it, are located there. In addition, the process of establishing and strengthening memories appears to involve associative links between different parts of the brain. For example, the amygdala mediates those bodily conditions and states which we subjectively experience as feelings, and a memory that includes strong emotion involves signals to and from the amygdala.[27] These links are the basis of Pavlovian conditioning (Mirolli et al. 2009). Associative links from an imagined/remembered signal to the amygdala and back can result in considerable strengthening of the imagined perception. This is presumably the basis of 'wishful thinking.'

It appears likely that associative memory is a function of the architecture that we have just outlined. A higher-level system for recognizing and controlling a certain perception sends reference signals to lower-level systems, which return perceptual signals whose strength above that of the (branched) reference signal is determined according to their perceptual input. Conversely, a perceptual input arriving from a lower level results in an error signal (unless it just happens to match the reference input). The output function sends signals which contribute to their reference inputs of lower systems. A 'reminding' process spreads across the hierarchy.

Often, we control so as to recreate a remembered perception.

> A choir director blows a note on her pitch pipe and one second later, fifty voices recreate the same pitch in 102 ears. A man shows a boy how to make a multiple jump in checkers and half an hour later the boy does it to the man. A person memorizes a street address and six months later drives slowly along a street until he perceives a matching house number. (Powers 2005:219)

The choir members begin controlling the pitch in imagination immediately while the pitch pipe is sounding. The street address is rehearsed in imagination, and perhaps by also speaking aloud, to

---

[26] (Powers 2005.207–230, Kandel 2006, Rudy 2008).

[27] There is a reciprocal relationship between the amygdala and the right ventrolateral prefrontal cortex (RVLPFC), such that talking about emotions increases activity in the RVLPFC and reduces activity in the amygdala (Lieberman et al. 2007). This is the basis of certain kinds of interventions in emotionally challenging situations, and is an inherent aspect of the Method of Levels, an application of PCT to psychotherapy.

commit it to memory.[28] By iterating a request for that desired state of the perception, a higher-level system strengthens the synaptic memory at the reference input.

Figure 5 shows eleven proposed levels or orders of perceptions in the perceptual hierarchy. In Powers et al. (1960) and Powers (1973, 2005) these levels were provisionally identified on largely phenomenological grounds, with experimental support most strong for the first three levels. Of the varieties of experimental corroboration since then, among the most persuasive is a series of observational studies conducted by Frans Plooij, Hetty van de Rijt, and their colleagues. They demonstrate that the emergence of each successive level of the perceptual hierarchy corresponds behaviorally to predictable regression periods in the maturation of primate and human infants, each regression followed by growing competence controlling the new order of perception.[29]
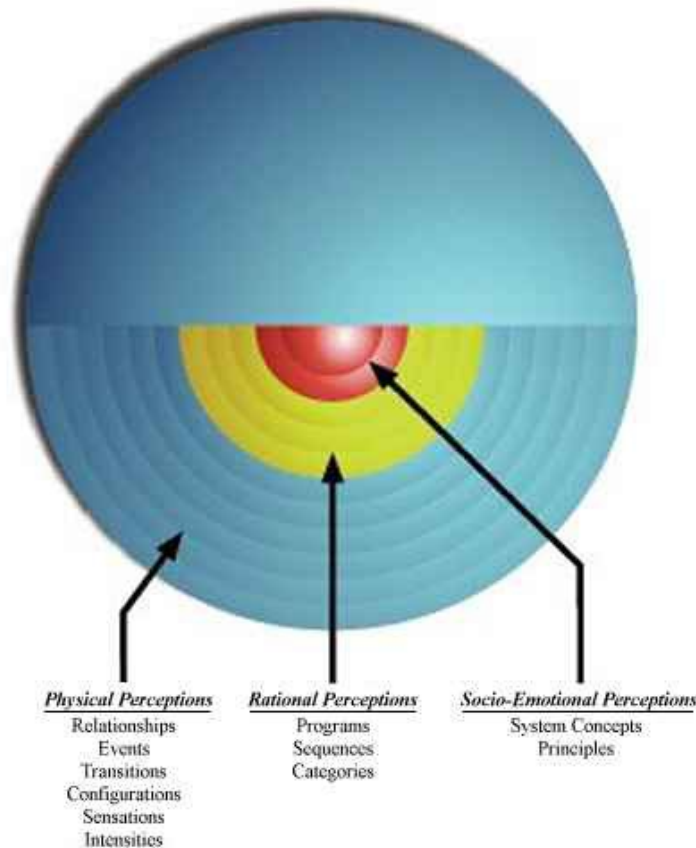

Physical Perceptions
Relationships
Events
Transitions
Configurations
Sensations
Intensities

Rational Perceptions
Programs
Sequences
Categories

Socio-Emotional Perceptions
System Concepts
Principles

Figure 5. Eleven orders of perception represented concentrically in three groups.[30]

McClelland (2011) has organized the eleven orders of perception into three groups. The lowest six levels are *physical* perceptions:

1. **Intensities**: a signal from a sensory organ, e.g. how bright, or with how much pressure.
2. **Sensations**: a vector of intensities, e.g. green, or the taste of an apple.
3. **Configurations**, e.g. discrete objects, or the extent to which a limb is bent.
4. **Transitions**, e.g. rotating counterclockwise, or receding.[31]
5. **Events**: sequences that are short and well practiced or well skilled, e.g. a syllable or word.

---

[28] The checkers example probably involves the 'memory' in a structural change rather than a reference value stored in memory. The learning just makes the system controlling the procedure "jump the opposing piece" iterable with the same jumping piece.

[29] Plooij (1984), Plooij & van de Rijt-Plooij (1990), Heimann (2003). Van de Rijt & Plooij (2010) is a very accessible account directed at parents of young children

[30] Adapted from a figure in McClelland (2011).

[31] Yin (2014) locates Transition control in the basal ganglia.

6. **relationships**, e.g. under, inside, or adjacent.

The next three levels McClelland's termed *rational* perceptions:

7. **Categories**, the results of processes of generalization or analogy.
8. **Sequences**, e.g. the steps in a recipe or in directions for finding a hotel.
9. **Programs**, e.g. deciding which item on a menu to order.

The last two, McClelland calls *socio-emotional* perceptions:

10. **Principles**: heuristics, and goals which typically might be expressed as maxims (e.g. "honesty is the best policy").
11. **System Concepts**: collections of principles which are perceived as a unity, e.g. The Red Sox, the Marines, physics, 'the government'.

Just as the phenomenon of control differs fundamentally from the material actions and reactions studied in the physical sciences, so the interaction of autonomous control systems—controlling one's perceptions of other living things in the environment—differs from controlling one's perceptions of the inanimate world. Most obvious are the possibilities of conflict, accommodation, and cooperation. Of these three, conflict is most extensively discussed in the literature of PCT. Kent McClelland (2004, 2011, forthcoming) has shown that the effects of collective control are perhaps even more ubiquitous in the built environment, but conflict is much more obvious because it causes trouble, and collective control avoids it.

McClelland has demonstrated how an environmental variable is stabilized when two or more individuals concurrently control their respective perceptions of it (and that this is so even if they are in conflict). When a variable so stabilized then becomes relied upon as public means for controlling other perceptions, its very stability as a feedback path through the environment can itself become a controlled perception. The maintenance of its stability cannot be accomplished by any individual alone. This appears to be the basis of what we call a social norm or expectation. Because the several members of a public are controlling the stability of common environmental feedback paths, no one member has to expend much energy or devote a lot of attention to controlling it.

Attention[32] is a limited resource. The eyes can only be directed to a limited portion of the environment, one's hearing attends to some sounds while ignoring others,[33] and so on. When a social norm is not met, any control loop that passes through that feedback path in the environment is disturbed, and attention must be diverted from control of other perceptions. Regaining control places demands on our available resources for perception and control, perhaps even including higher-level strategies for solving problems. Even though no one devotes much conscious attention to controlling these norms, because their maintenance is shared and their use is established habit, deviations are quickly resisted, indicating that they are in fact controlled with high gain.[34] There is a further ramification. What is at stake when a convention is violated is a perception of the reliability of others to cooperate in goals that require mutual aid, and their reliability to avoid interfering (coming in conflict) with personal goals. When trust is weakened or lost, one must be more attentive to perceiving and averting potential sources of conflict.

---

[32] Attention, awareness, and consciousness are rather mysterious, because the vantage point of one's awareness is always different from that of which one is aware. In PCT, this dualism, by shifting active control from a conflict to a 'neutral corner' of observation or commentary, is the basis of the Method of Levels (MoL) (Carey 2006, 2008; see also Powers et al. 2011:13 & fn. 8). The mystery of who or what the experiencer of perceptions may be is the basis of koans, those Buddhist puzzles that are said to lead one to a realization of the nature of mind: Can that within us which is aware of a perception and attends to it be aware of and turn attention to itself? What are we aware of in the gap between one focus of attention and the next? Who do you think you are?

[33] The mechanism for ignoring perceptual input may find an account in the parallel processing architecture introduced in the next section.

[34] Loop gain is a function of amplification at several points in the control loop. It affects the rate of increase in the strength of behavioral outputs in proportion to the rate of increase of the effect of disturbances on the controlled variable. The effect of high gain is not high output (unless the disturbance is strong), but low tolerance for error (see Powers 1973:290). This distinguishes circular causation in a negative-feedback loop from sequential linear causation of responses by stimuli, as e.g. T.O.T.E. (Miller & Galanter 1960), where amplification in successive stages can only increase output.

PCT provides a quantitative, precise, testable model for the actions and interactions of all living things, from molecular biology to sociology. That said, a disclaimer is in order before the discussion of language as the control of perceptions. PCT claims only the empirical fact of control and the strong hypothesis that when an implementation of the PCT model in a computer simulation closely replicates the behavior of the individual whose behavior was measured, the control structures in the simulation inform us about unseen control structures within the individual. Just how those structures are implemented in neurological and biochemical 'wetware' is largely unknown, but the PCT model proposes specific bold hypotheses that can be used by neuroscientists to guide research. Every PCT simulation and every interactive PCT demonstration is an *experimentum crucis* testing the PCT model. However, the specific models of language proposed here have not yet been implemented as simulations, so they are no more than proposals within the general PCT model. And, I should emphasize, neuroscience is not my field, so I am relying largely on secondary sources.

PCT explains how language is at the same time an individual endowment and a social product. The elements of language discussed in the next section result from and are maintained by collective control.

Language perceptions

A language is used by members of a speech community as means of achieving aims that require coordination with others. To an infant, the collectively controlled regularities of the language are as much a part of the environment as food, toys, and smiles.

When a baby is born, its brain does not yet have all of the neurological systems that are needed recognize the full hierarchy of perceptions that an adult does. Intensity signals from the sensory organs develop in utero, and on that basis successive layers of recognition and control of perceptions are added one at a time. With each addition of a level, the child quite literally lives in a new perceptual universe, and has to learn how to control a newly recognized type of perceptions. Van de Rijt-Plooij & Plooij (1992, 1993) and Plooij (2003) found in their research that as a new level of the perceptual hierarchy emerges and begins learning how to use the prior levels as its means of control, it disrupts the child's newly-gained competence (and confidence) controlling at the prior levels. The child who was confident and happy suddenly becomes insecure. This is the basis of predictable 'cranky' regression periods that are well documented in child development.[35]

The child's brain gains control and develops skill by trying out different connections and different signal weights, both in the error signals branching downward to evoke reference signals at lower levels, and in the perceptual signals branching upward from lower levels to the input functions that construct perceptual signals at the new top level. PCT models this as a process called reorganization. As neural systems at the new level gradually become more skillful at varying the reference signals at lower levels, the child's confidence and good humor return.

Linguists specializing in child language development have identified predictable stages in Children's learning of the language (or languages) that they come to speak as adults. Newborns go through an inarticulate "phonation" stage for about 8 weeks while the vocal tract lengthens enough for mature articulation to be possible. This is followed by what is called a "gooing" stage for another 11 weeks or so, then for about 7 weeks during the so-called "expansion" stage there is heard a growing ability to produce clear vocalizations with the diaphragm and larynx and to restrict the vocal tract, producing interruptions in the flow of vocalization that sound more and more like consonants.

In the gooing stage, and even more in the expansion stage, infants play with all the ways that they can produce sounds, manipulating all the organs that affect the flow and vibration of air between the larynx and the apertures of the nose and mouth. For any given articulation and sound, however strange it seems, it is almost always possible to find at least one language in the world that makes phonemic distinctions by that means, but most of the sound-gesture events that the child produces at this stage are not used to distinguish words in the language that they will eventually learn to speak. For example,

---

[35] Van de Rijt-Plooij & Plooij (1992, 1993), Plooij (2003).

even a child of English-speaking parents occasionally produces sounds by extending the epiglottis toward the aryepiglottic folds at the top of the larynx. Very few languages use this speech gesture— Austronesian languages, Semitic languages, and a few languages of the Pacific Northwest region of North America. During this same period, the infant also waves about its arms, legs, fingers, and toes, exploring what they do, reorganizing internally, and gradually developing control systems with which to move them purposefully.

It is not until 6 or 7 months, as Relationship perceptions consolidate their control of Event-level perceptions, that the first overt signs of language appear, the babbling of simple syllables. A syllable is an acoustic and gestural perceptual unit at the Event level.

Child language researchers have long observed how parents and other adult caregivers employ peculiar modifications of language called 'parentese' when talking to small children, and that infants clearly prefer it. The contrasts of words and the rhythms and intonation patterns of the adult language are exaggerated. It is not difficult to see that this assists the development of neural systems in the child's brain that learn to recognize and then control these perceptions. As the Canonical Babbling stage begins, the child increasingly uses the phonemic contrasts of the adult language to form syllables, and uses sequences of syllables to control the language-specific intonation contours of assertion, exclamation, question, etc. By this stage, the baby has developed the first six levels of the perceptual hierarchy. Table 1 shows how these stages of language development line up with the emergence of successive levels of perception and control.

Table 1. Emergence of perceptual levels aligned with stages of language development

wks	mos	yrs	Perceptual Order	Duration	Language Stage
0			1. Intensities	5 wks	Phonation
5	1		2. Sensations	3 wks	Phonation
8	2		3. Configurations	4 wks	Gooing
12	3		4. Transitions	7 wks	Gooing
19	4		5. Events	7 wks	Expansion
26	6		6. Relationships	11 wks	Canonical babbling
37	9		7. Categories	9 wks	
46	11		8. Sequences	9 wks	
55	13	1	9. Programs	9 wks	~ 12 months: words
64	15		10. Principles	11 wks	
75	17		11. Systems		
	24	2			~ 2 years: short sentences
	48	4			~ 4 years: Theory of Mind
	60	5			~ 5 years: pragmatics

Level 1: The perception of **intensities** had already developed *in utero*. Examples of intensity perceptions include brightness of light, loudness of sound, the intensity of a taste, the degree of pressure of a touch. From a newborn, we hear cries and unorganized phonation.

Level 2: At around 5 weeks, the baby's brain begins to construct changing **sensations** from combinations of intensities. In visual perception, the brain now constructs a range of colors from the intensities of signals from red-, green-, and blue-sensitive cone cells in the retina. Inarticulate phonation continues.

Level 3: Configurations and patterns are constructed from sensations at about 8 weeks. With control of the configuration of the vocal folds there is more play with the sounds of the voice in what is called the 'gooing' stage.

Level 4: Transitions (smooth changes in configurations) emerge at about 8 weeks. The 'gooing' stage of vocalizing continues, but vowel sounds begin to approximate those of adult speech, and there is more play with squeals and other pitch variation.

Level 5: An **Event** is a short, well-practiced combination of configurations and transitions. Syllables are Event-level perceptions. In the 'expansion' stage of language development, at about 4 months, the child begins to recognize and control simple consonant-vowel (CV) syllables.

Level 6: At around 6 months, control at lower levels is destabilized by a growing ability to recognize **relationships** such as spatial distance. It is not until this new level begins to settle down that the infant begins canonical babbling, playing with the relationship of contrast between the audible event

perceptions that we recognize as syllables. Control of phonemic distinctions emerges as a function of perceiving herself making syllables of the sort that she hears adults around her making, and having some of them recognized by adults as words in correct context (such as *oo* for "oops!" and "achoo!" when someone sneezes). As syllables come under control, she practices the pitch and amplitude configurations of intonation contours and the temporal configurations of speech rhythms.

Each of these first three stages of language development lasts about a month (shown in the **duration** column in Table 1). In the phonation stage, two levels of perception emerge during 8 weeks; in the gooing stage, two more levels emerge during 11 weeks; in the expansion stage, during a period almost as long (7 weeks), the child comes to recognize Event perceptions, including syllables. After that, "Canonical babbling" of increasingly well-formed syllables and intonation contours extends over a period three or four times as long, 30 weeks or more.

During this period, as Relationships, Categories, and Sequences are recognized and come under control, a lot of passive learning of vocabulary is going on. For language development, this is an actively receptive learning period. Words learned as passive vocabulary provide a scaffold for categorizing (Mirolli & Parisi 2005) as the child very studiously sorts experiences into different kinds.

Level 7: Most PCT researchers refer to this as a **Category** level of perception, emerging at about 9 months. Evidence adduced for Category perceptions often is a projection of the observer's use of language. I will argue, as Taylor (forthcoming) does independently, that categorization is a process at every level, as though on a 'vertical' plane orthogonal to the levels, much as language perceptions interface with non-language perceptions. For the present discussion, it is immaterial whether we call these Category perceptions or merely increasingly complex kinds of Relationship perceptions

A more subtle and crucially important capability emerges at 9 months: the capacity to imagine perceptions from another's point of view. It is this which enables humans to cooperate, and to learn from one another not just the goal and means of an activity but the manner of executing it. The ability to replicate successful technique has the effect of a 'rachet' of incremental improvements which, over generations, has enabled the complexity of human cultures (Tomasello 1999, 2014; Herrmann 2007; summarized in Stix 2014). Separation anxiety begins as the child becomes aware that caregivers have other things on their minds and are capable of abandoning them.

Level 8: At the **sequence** level, beginning at about 11 months, the child becomes fascinated how one event follows another, and how objects fit together in a certain way. Neural systems at the sequence level construct perceptions that can look at words and control them.

Level 9: After another 9 weeks (at about 13 months), **program**-level perceptions logically organize and combine sequences so as to reach goals in a variety of ways, with some flexibility of choice along the way. Building a tower of blocks, getting dressed, eating breakfast are examples of goal-perceptions that are controlled at the program level. It is not until the Program level begins to connect words to their instrumental and pragmatic uses ("If I say this, then she does that") that adults recognize words in a child's utterances, typically somewhere between the first birthday and 13 months. After that comes an 'explosion' of active vocabulary as higher-level systems learn what they can do by means of that backlog of recognized passive vocabulary.

The mastery of language continues as the child gains skill at all levels of perceptual control. Above the Program level are heuristics and other **Principles** by which one selects an appropriate program (at about 15 months), and **System concepts**, including a developing self image and a concept of family.

What is called Theory of Mind (ToM) develops gradually, culminating at about age 4 (Astington et al. 2010). At about 2 years of age children are speaking short sentences, they use tense appropriately, they ask questions, and they talk about such attributes as colors and sizes. At that age, they also understand that when people get what they want they are happy and when they don't they are sad, and they are beginning to see that there may be a difference between what they want and what another person wants. By age 3, they realize that others may have different wants and feelings, likes and dislikes, and also talk about what people think and know, but it is not until about age 4 that children recognize that

others may harbor beliefs that are not true and may be ignorant of what they themselves know. This makes possible pretense, trickery, and prevarication, all of which involve controlling (in imagination) a perception of another person's understanding, different from one's own.[36] At about 5 years, they begin to develop more control of what is called pragmatics, the contextual appropriateness of their language use. No higher levels of the perceptual control hierarchy are required for these later developments. They seem to depend on growing competence with controlling perceptions within the hierarchy which is in place by 18 months, especially control involving the interrelation of language perceptions with non-language perceptions. Throughout this development, control of language perceptions becomes a more and more important means for organizing other perceptions.

Empirical linguistics has identified the kinds of perceptions that a baby must learn to recognize and control so that the toddler can understand sentences and produce them:[37]

1. Phonemic distinctions.
2. Words and morphemes, with their main meanings.
3. Word dependencies (the argument requirement of each word).
4. The selection of each word (the dependencies that have greater than average likelihood).
5. The canonical or preferred linearization (word order) and its alternatives.
6. The main reductions (variant shapes of words), their domain (a particular word or all words in a position) and the conditions for applying them.

Phonemic distinctions. The perceptions that are relevant for a science of language, and for a baby learning to speak, are not speech sounds, but speakers' judgments of contrast. Phonetic features of speech are associated with the contrasts by the pair test, an especially rigorous form of the substitution tests which are used for all distributional analysis.[38] The contrasts can be represented in various ways. For the further work of grammar, it really does not matter much which representation is chosen. Dialectologists are concerned with the detail of a 'narrow' transcription, and phonologists are concerned with the simplicity and generality of their description of the sound system, morphology prefers a representation that simplifies the description of morphophonemic alternations, which often also best suits the purposes of comparative reconstruction of language history. For language pedagogy and for the further work on syntax and semantics a vernacular orthography is almost always used, however 'imperfect' it may be for other purposes. Any of these representations can be reached by a distributional analysis, and they are all 'phonemic' to the extent that the original judgments of contrast are recoverable.[39]

When we represent the locations of contrast alphabetically, they are called phonemes. Our familiarity with writing with alphabetic symbols prejudices our view of these matters. Speech gestures and their acoustic effects in fact overlap and don't fit into tidy alphabetic segments.[40] Most illiterate speakers or speakers literate only in a non-alphabetic writing system are unable to segment utterances into individual speech sounds (Hyman 2010:281n41). The important point here is that it is not the specific sounds, phonetically, but the phonemic contrasts that are the fundamental or primitive elements of a language.[41] We are so accustomed to represent the points of contrast with alphabetic letters that we think that the letters, or the sounds they appear to represent, are the basic elements, and indeed it is

---

[36] ToM develops earlier when the child is often engaged in pretense play (Garfield et al. 2001), and when parents give reasons for correcting the child's misbehavior and talk to the child about thoughts, wants, and feelings (*op. cit.*). Needless to say all of these things require considerable competence with language.

[37] Harris (1991.404–405). Also, with slight differences, in Harris (1988.111–113).

[38] Non-linguists should understand that the word *distribution* here does not have the meaning employed in statistics and probability. The linguistic usage is analogous to that in ecology, referring to environments of occurrence.

[39] The judgments of contrast are primitive, not 'phones', *pace* e.g. Bloch (1948), Chomsky (1964).

[40] Browman & Goldstein (1989, 1992).

[41] "The fundamental data of descriptive linguistics are ... the distinctions and equivalences among utterances and parts of utterances. The elements of [phonological analysis], including the setting up of phonemes … are manipulations of these distinctions on the basis of distribution" Harris (1951a.33).

convenient to talk of them in such terms, but in fact each letter, or phoneme as they are called, is merely a marker for a place where the given syllable or word contrasts with possible other syllables or words. A given contrast can be maintained by diverse and variable means. The perceptual requirement to distinguish words is primary, and the variable means of doing so secondary and derivative.[42]

The investigation of language development (Table 1) shows how as children we learn to control syllables first, and then subsequently refine our control of the contrasts between them. In the transition from canonical babbling to words at about 12 months, it appears that words are constructed of syllables (event-level perceptions), and that the syllables are differentiated by phonemic distinctions but not necessarily constructed as sequences of segmental phonemes.[43] Even after syllables are well formed with orthodox phonemes, the choice and sequence of phonemes is idiosyncratic for some time. It is common for immediate caregivers to tune into the child's way of controlling the matrix of phonemic contrasts before anyone else (even a trained linguist) can recognize words in their utterances.

A simple example will illustrate. My youngest daughter had reached a stage at which we began recognizing many of her words. But she said [tʌks] as she struggled to get out of her highchair. Where are the ducks, we wondered. Getting to her feet on the seat of the chair, she triumphantly stamped and crowed [tæns]! (with an unaspirated [t] that sounds close to our *d*). ("Yes! You're dancing!" affirmed her mother, herself a dancer.) Looking out the window at the wintry landscape, she said [nows]. Is her nose cold? It was not until she pointed out the window and said [nowsmæn] that light dawned. She was pointing at a snowman. Suddenly, like a key turning in a lock, it became clear that where we benighted adults began a syllable with initial *s* plus another consonant, she put the *s* at the end of the syllable: [tʌks] = *stuck*, [tæns] = *stand*, [nows] = *snow*, [nowsmæn] = *snowman*. All the phonemic distinctions were present, though not in the conventional sequence.

A number of other linguist parents have reported the same phenomenon, and I have no doubt that this or similar transpositions often go unrecognized. Needless to say, as we began repeating the words that our daughter intended to say we provided acoustic feedback that resulted in her changing her pronunciation until what she said sounded more like what we said. The process by which she accomplished this is illustrated in Figure 6, and in the discussion there of how we control our pronunciation indirectly, adjusting how the gestures of talking feel until it sounds right.

Pronunciation is controlled at lower levels of the hierarchy, up to event and sequence. Reference values for pronunciation are set by two kinds of processes. We are most aware of what our speech sounds like. We are generally not aware of controlling what it feels like, but we do the former by means of the latter.

Figure 6 illustrates the relationship between these two control processes. We can't control our speech sounds in real time,[44] because by the time we hear them it is too late to change them. If we make a mistake in pronunciation, we can only repeat the utterance with the mistake corrected. The perceptions that we control in real time are our tactile and kinesthetic perceptions of the organs of speech articulation. However, the reference values for articulation are adjusted over time as means of making our pronunciation sound right to us. We may draw an analogy to learning the skill of skeet shooting or trap shooting by repeated practice, as opposed to real-time control in aiming a rifle by aligning the

---

[42] For discussion of how relationship is primary and relata derived, see Bateson (1979:170) and Brown (1972): "the universe comes into being when a space is severed or taken apart [...] We take as given the idea of distinction and the idea of indication, and that we cannot make an indication without drawing a distinction. We take, therefore the form of distinction for the form." (*ibid*. v & 1)

[43] Phonemes are often thought of as categories of phonetic segments or phones. That this is a mistake can be seen in the developmental process outlined here. Phonemic contrasts make the difference between imitation (phonetically similar) and repetition (phonemically same). Alphabetic phonemic segments are no more than a convenient representation of the contrasts, and they are just one way to represent them. See Harris (1951.33–35, 75–76, 361), Nevin (1993a), Browman & Goldstein (1989, 1992).

[44] Artificially prolonged vowels and sonorants (*m*, *n*, *l*, and the like) can be corrected in real time, but that is not normal speech.

silhouette of the sight with the image of the target.[45] The relationship between auditory and articulatory control is a bit easier to notice when one practices repeating a tongue-twister (*she sells sea shells*, or the like).
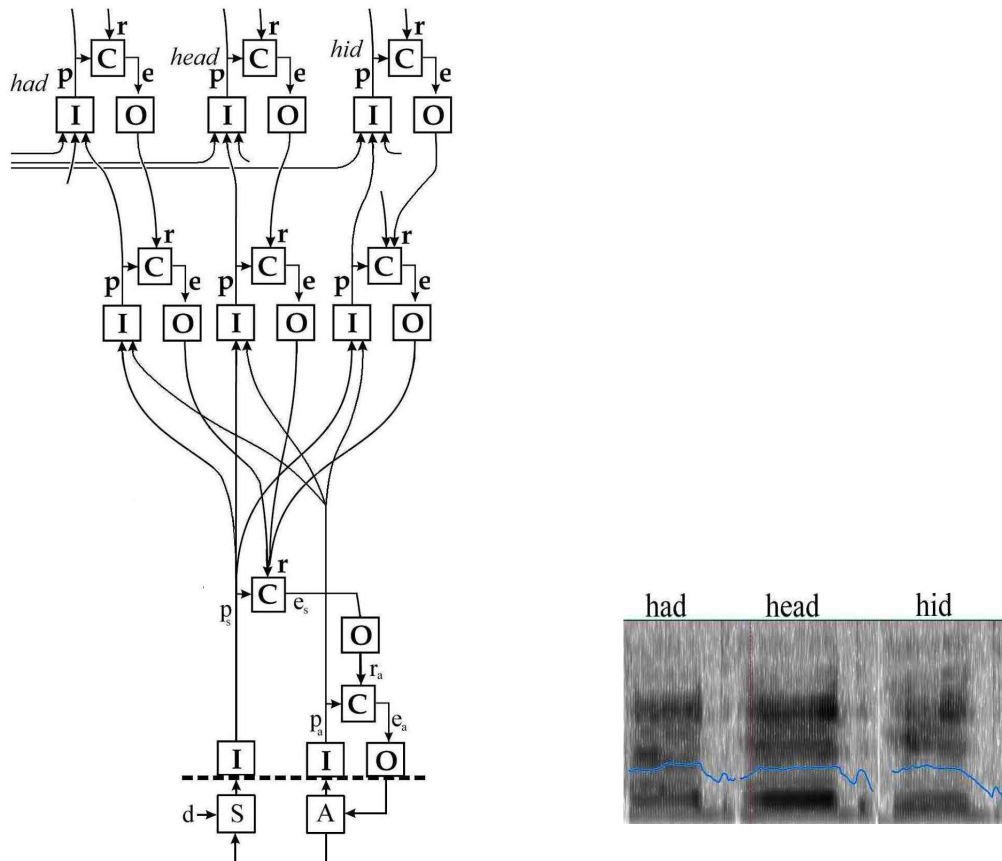

Figure 6. Two kinds of reference-setting in speech. The sound of one's own speech (labeled S) is controlled by adjusting the kinesthetic and tactile reference perceptions for articulation (labeled A) over repetitions of a given phoneme, syllable, word, or phrase. Recognizers (controllers) are shown here only for the medial vowels of the words had, head, *and* hid.

The model diagrammed in Figure 6 predicts that if we were able to disturb the sounds of their own speech that speakers heard in headphones, so that for example when they said *head* repeatedly it gradually came to sound to them more like *hid*, they would resist that disturbance by shifting their pronunciation toward *had*. There are two reasons for disturbing only the vowel. The first is that, unlike consonants, vowels involve few and slight tactile perceptions limited to the edges of the tongue. This is why the most common dialect differences in pronunciation (rather than word choice) are differences in the vowels. The second reason for disturbing the vowels is because their production is continuous, whereas the transient changes in formant frequencies that distinguish one consonant from another are very brief and precisely timed, so they are more difficult to disturb.

In fact, this experiment has been done by researchers who have no awareness of Control Theory[46] (Katseff 2010; Katseff & Houde 2008; Katseff, Houde, & Johnson 2008, 2010), and that prediction has been shown to be correct. The PCT model also correctly predicts that subjects do not completely resist the disturbance to the vowel that they hear. The diagram shows that articulatory perceptions continue upward and contribute to the perception of the word *head* and its neighbors, *hid* and *had*. The disturbed acoustic input gets corrected at the cost of moving the articulatory perceptions away from

---

[45] The example is due to Bateson (1979:216-219).

[46] They have not responded to correspondence.

their references. This puts the two in conflict, and the result is an optimization in which each controls as well as it can but neither is satisfied.[47]

Acoustic and articulatory perceptions are combined with perceptual input from several sensory modalities to construct a perception of what syllable or word a person is speaking. For example, so-called lip-reading is more important than many of us realize. In the McGurk effect,[48] if a video is made of a speaker producing an utterance and it is is dubbed with a sound-recording in which one phoneme is substituted by another, the brain overrides the actual acoustic input with a different auditory perception. When e.g. the nonsense syllables /ba-ba/ are dubbed into a video showing the lip movements of /ga-ga/, what you inexorably hear is /da-da/. Close your eyes, and you hear /ba-ba/; open them, and you hear /da-da/. Given visual perception that the speaker's lips are not moving as needed to produce /ba-ba/, the brain creates a mistaken audio perception that supplants actual audio input. This is why infants require visual perception of people's faces to learn language. Also important among the language perceptions that influence the control of word perceptions such as those shown at the top of Figure 6 are word dependencies, each word's selection, linear word order, and the reductions that change the overt shape of words. These are discussed in the subsections below.

The control of phonemic contrasts establishes discrete elements in the continua of speech, satisfying the first requisite for a language: the elements must be discrete, socially predetermined, and arbitrary.[49]

Words with their dependencies & selections. In each language, only some of the possible sequences of phonemes constitute words. Words and morphemes are discernible as discrete elements because they recur in different contexts and because the set of possible next phonemes is greatest at morpheme boundaries. Stochastic processes can identify where their boundaries lie (Harris 1955, 1967, Saffran et al. 1996a, Goldsmith 1998, 2001, 2009).

Meanings are obviously not inherent in words. There is nothing in the sound of *cheval*, *Pferd*, *jahhôŏ m*, *mǎ*, *horse* that associates these words with nonlanguage perceptions of large, four-hooved animals in the brains of French, German, Achumawi, Chinese, and English speakers, respectively. There is abundant evidence that we learn meanings of words from their placement in context of other words. This is why dictionaries include example sentences. Gleitman (2002) summarizes research demonstrating "a perceptual learning process sensitive to interacting distributional regularities at many levels of analysis" (Gleitman 2002:213) and reports how surprised she and her colleague felt when they "discovered that the first verb in a blind child's vocabulary was liable to be *look* or *see*, and that these words—learned from the normal usage of sighted parents—were used sensibly to describe haptic exploration and apprehension" (Gleitman 2002:215).

Words are grouped into classes according to their substitutability in contexts of repetition. Each word-class is an approximate statement of word dependencies, because it is defined in terms of the classes of words that can co-occur with it. The work of analysis refines these word classes and subclasses, the co-occurrence restrictions between them, and the characterization of word dependencies.

Various grammars of word-classes have been devised. As a formalization of Bloomfield's Immediate Constituent analysis, PSG defines a hierarchy of phrase-class names as abstract symbols manipulated by rewrite rules. Methodologically, the phrases are determined according to the gestalt psychology notion that people intuitively know how to break a complex item into constituents. Harris (1946) treats constituent structure instead as word-expansions determined by substitution tests. Harris immediately recognized the relation of these expansions to conjoined sentences and began teaching transformational analysis in the late 1930s.[50] Work with Aravind Joshi and others on computability of

---

[47] The PCT model also explains other details of the experimental results, but that is out of scope for this paper.

[48] McGurk & MacDonald (1976). A good demonstrations of the *b-g* swap is at http://auditoryneuroscience.com/McGurkEffect. A BBC clip demonstrating the McGurk effect with a *b-f* swap is at http://www.faculty.ucr.edu/~rosenblu/lab-index.html.

[49] Harris (1968:6-11).

[50] Lisker, in email of 1 March 2000 to Bruce Nevin.

language reformulated these expansions as strings in an adjunction grammar (Harris 1962, Joshi, Kosaraju & Yamada 1972a,b). Joshi and his students combined rules of different types in tree-adjoining grammars (TAGs), in which rewrite rules (Joshi 1969, 1985) or dependency rules (Joshi & Rambow 2003) generate the center strings and adjunct strings for adjunction rules. TAGs avoid problems of PSG due to uncontrolled proliferation of abstract preterminal symbols (Joshi 2002) and problems that generalized dependency grammars encounter treating adjunct dependencies in the same way as operator-argument dependencies. They combine the strength of rewrite or dependency rules in specifying endocentric constructions (operator-argument dependencies) with the strength of adjunction or expansion grammars in specifying exocentric constructions (dependencies reshaped and partially obscured by reductions). It is even possible to formulate a dependency grammar encompassing both types of constructions (Joshi & Rambow 2003), which may be useful for formalizing Operator Grammar for computer processing purposes.

However, while symbol-manipulating rules are necessary for natural language processing with a digital computer, despite the centrality of the 'information processing' metaphor to Cognitive Psychology there is no direct evidence that the human brain requires anything more than word dependencies and equivalences of phrases with their reduced forms, as specified in Operator Grammar.

The partition of the base vocabulary into N and the several classes of operators O gives an approximation to word selection that is more fine-tuned than that seen in the various grammars of word classes, but to reach a grammar of words, this gross classification must be further refined by specifying the selection of each word.

Some sets of words have a particular selection in common. To take an obvious case, many operators of the O_{no} class, such as *know*, *think*, *say* require as their first argument a subset of N which refer to human or human-like subjects. Some words have exceptionally broad selection, e.g. indefinites such as *someone*, *something* and the demonstrative pronouns derived from them (Harris 1982:201-205). That they contribute correspondingly little information points up the correlation of selection with meaning. A reflex of word selection is seen in the way that acceptability varies depending upon context, which provide a criterion for transformation. A plausible context might even be found for *Vacuum purchased kelp*, such as a dream, or *Vacuum* as a name.[51] Selection restrictions are notoriously resistant to formalization, for reasons that will become more apparent when this discussion returns to a PCT framework for Operator Grammar.

Linearization. Prior to speaking (or writing) the dependency relations among word-perceptions in the brain are presumably not temporally ordered. As an operator enters into the construction of a sentence, the operator and its arguments necessarily are linearized. In the course of speaking, we frequently begin with one linearization and, interrupting, shift to one that is more opportune. This is hidden in written texts but something of it may be seen in successive revisions of a finished text. One demonstrates understanding by recapitulating the word dependencies in what was heard or read, but this is seldom an exact reiteration of word sequences. PCT can help us to understand how this is not mere error-prone 'performance' obscuring an underlying linguistic 'competence', but rather is inherent to language as the control of perceptions, and can explain why paraphrase is a better measure of comprehension than rote repetition is.

Every language has a preferred word order for an operator entering upon its arguments, SVO for English, SOV for Latin, and so on, together with secondary alternatives as in *This I refuse* and *John you met yesterday*. As mentioned, alternative word order is important for juxtaposing repeated words as a condition for reduction. Argument-indicators help the recipient to sort out the intended dependency relationships. For example, in English a preposition marks the third argument of an O_{nnn} word, e.g. *to* after *give*, *donate*, etc.: *John donated the rare book to the library*; also *The rare book John donated to the library, To the library John donated the rare book*, and in an older poetic style even *John the rare book donated to the library*. After *give*, the most frequently used of these O_{nnn}

---

[51] A number of writers have invented plausible contexts for Chomsky's celebrated example *Colorless green ideas sleep furiously*, including a poem by Dell Hymes.

operators, the *to* is so predictable that it can be zeroed in the canonical word order: *John gave the library the rare book* (but not after the other operators of this O_{nnn} set, e.g. **John donated the library the rare book*; and not in other linearizations, e.g. **The library John gave the rare book*).[52] Conversely, when words are linearized in other than the canonical word order, the argument-indicators are especially important, and indeed languages in which argument-indicators are organized in elaborate case systems are more free to vary the word order. The operator is marked by tense and aspect suffixes, with the default or neutral form in English being the suffix *–s*, the so-called present tense[53] (which is actually atemporal). The operator-indicator is placed after those operators that we think of as verbs, as in *He seeks fame*, and it is placed before operators that have more durative meaning, which we think of as adjectives, using *be* as a carrier: *The book is rare.*[54] These morphological details that help with keeping track of word dependencies are not learned immediately by children, and may be learned first as frozen morphophonemic alternants of particular words before it is grasped as an active affixation.

Reductions. The reductions are the part of Operator Grammar that is difficult to formalize for computer processing. Fortunately, that is not our present project. There is no *a priori* reason to suppose that the learning of the reductions is fundamentally different from the learning of morphophonemic alternation and suppletion. The differences are in scope, scale, and complexity, but not in kind. The 'contraction' *can not ↔ cannot ↔ can't* is commonly described as morphophonemic; the reduction *I expect John to be here. ↔ I expect John.* is not, nor are the reductions

John is here. I saw John (same as nearby). ↔ John is here. I saw him.
John washed John (same as nearby). ↔ John washed himself.
I gave a male person (same as nearby) that; that (same as prior) is a book. ↔ I gave him the book.

Improbable word sequences like that last example 'disappear' immediately, as it were, being reduced to conventional form as soon as the words enter into the relationship that enables the reduction, but a proposal that children actually learn word sequences like that, even tacitly, would be difficult to support. PCT can shed a great deal of light on the proposition that not only the word dependencies but also all the possible linearizations and reductions are all available concurrently when an operator is asserted. This is a practical consequence of Gross's demonstration that a unique set of dependencies and reductions is a lexical property of each word. But first, let us briefly consider an essential characteristic of language that presents a very similar challenge, that is to say, the ways in which variant forms—variant in geographic and social space and changing through time—are equivalent.

Language variation and change.

We can say *John gave the library the rare book*, but not **Jean a donné la bibliothèque du livre rare* ; we can say *He comes and goes*, but not **Il vient et va*. Such differences between specific reductions illustrate how languages differ in their reductions (Salkoff 2002), as well as in the forms of words.

Language changes through time, and the collectively controlled perceptions of different communities diverge in the course of time if their people are not in frequent communication as peers. (1) Phonetic differences in pronouncing the same phonemic shape were discussed at the beginning of this section. (2) Vocabulary may be coined or borrowed as cultural needs change. (3) Different communities may use different phonemic shapes with the same non-language meaning and the same distribution relative to other words, as for example regional differences between those who say *soda* or *pop* or *cola* or *soft drink*, or those who say *sofa* vs. *divan* vs. *chesterfield*. This may be seen as a limiting case of (4) differences in word selection or the acceptability ranges of different operator-argument combinations: The acceptability ranges are essentially the same, but two forms may alternate or compete and there there may follow a complete substitution of the phonemic shape of an item. Switching from one such

---

[52] Other prepositions are operators, e.g. *on* in *the book is on the table*. Some linguists have argued that the argument indicators are an outward manifestation of a psychological reality, notably Fillmore (1968) as developed especially by Cook (1989).

[53] Explanation of tenses is straightforward, but would take us too far afield here.

[54] There is also a verb *be*, as in *God is*. The indefinite word *do* is also a carrier for tense and aspect in certain well defined reductions.

dialect to another amounts to adding a reduction operation to one's mental grammar. And of course (5) there are dialect differences in details of the reductions, an obvious example being those who naturally and freely use *ain't* and those for whom it is an affectation. Reductions may change or vary as to their domain, that is, which words are subject to it, as to the string or dependency conditions for carrying them out, and the frequency of carrying out a particular reduction may vary or change. In each of these five domains, analogy may play a role, but (6) analogy is particularly important when the selection of two lexical items intersects, and the distribution of one is extended to include all or part of that of the other. This very widespread historical process almost always is by way of a grammatical construction that is well established for one item and extended to the other, and thus can be accounted for within the system of reductions.

Operator Grammar identifies three distinct areas in which aspects of language structure may be in flux at the growing edge of language. The acceptability of such forms is fuzzy, uncertain, or disputed. To see these clearly, it is helpful to consider a partition of language into a Report structure and a Paraphrase structure (Harris 1969). Without any reductions the operator-argument dependencies generate an informationally complete sublanguage of sentential constructions that are maximally explicit but in many cases unspeakably awkward, and some are not said at all without obligatory reductions. The (mostly optional) reductions generate streamlined paraphrases which convey the same information more efficiently and in conformity with conventions of the language community.

Within the Report structure, some word combinations are grammatical but their acceptability is marginal or is contextually restricted, as for example *purchase* entering on the pair *vacuum*, *kelp*. Reductions may be so strongly expected or customary for a given form that all sentences of that form are of marginal acceptability. An example was seen earlier in the rather stilted paraphrases of relative clauses as secondary sentences, as in *John—said John you met yesterday—gave a book ↔ John, who you met yesterday, gave a book*. In Operator Grammar the definite article *the* is reduced from *that* plus a relative clause: *John gave the book ↔ John gave that which is a book ↔ John gave that; that (same as prior) is a book*. This derivation accounts nicely for the many idiosyncrasies and special semantic characteristics of the definite article (Harris 1982:237-244). The definite article is a relatively recent innovation in only some of the Indo-European languages, and with many difference of detail across those languages which do have it. In fact, our *the* derives historically from the Old English demonstrative *se*, *sǽ o, þæt* from which we also have our present-day *that*.

The reductions generate some constructions which are undeniably grammatical but beyond what is comfortably in the language. For example, walk is in the O_n class in *The dog walked*. The apparently transitive verb in *He walked the dog* can be derived by way of a non-extant intermediate resultant of reductions (Harris 1982:317-318, 1991:110): *He enabled the dog to walk ↔ †He walk-enabled the dog ↔ He walked the dog*.[55] The dagger character † marks an "infra-sentence" (Harris 1982:18). Figure 7 is an impressionistic representation of the variable, fuzzy boundary of fully accepted reductions within the set of all grammatical resultants of reductions, and may serve to suggest by analogy the indeterminacy of the above types of language variation and change within a determinate structure.

---

[55] Another word than *enable* might be found. The related causative *He caused the dog to walk* has obviously a different meaning.

Figure 7. A suggestive representation of the partition of language into two systems, Report and Paraphrase, showing within the set of fully grammatical paraphrases an irregular and variable boundary separating marginal and unattested products of reductions from the majority which are fully accepted.

Control of language perceptions

Earlier, we reviewed some of the influences in the perception of a word as contrasted with other similar words. We glanced at experimental work which placed control of auditory feedback at odds with control of tactile and kinesthetic (muscle tension) perceptions by disturbing the sound of a subject's speech in headphones in real time, and at the McGurk effect in which visual 'lip-reading' can cause an imagined perception to supplant auditory perception entirely.

A more important and pervasive influence on word perception comes from perceptions of operator-argument dependencies. As a simple illustration, in the context *the baseball hit his _____*, a mispronunciation as *hid* (or *had*) will not prevent us from perceiving the word *head*. There are at least two reasons for this.

First, from the bottom up: the context *h___d* being constant, perceptual input of the sound [hɛd] best matches the input functions for the phonemes in *head*, but the filtering of harmonics in the vocal tract is not absolute, there is also some acoustic energy at frequencies that activate the input functions of the adjacent vowel phonemes that would be heard in *hid* and *had* (phonetically [hɪd] and [hæd]).

Secondly, from above downward: *hit* is an operator of the O_{nn} class, so the input function of the recognizer for *hit* includes a requirement for two words of the *N* class. In the context *the baseball hit his _____*, where the blank has not yet been filled, the perception of *baseball* has satisfied the first argument, but the *hit* recognizer is still looking for its second argument by controlling it in imagination. If there is a reference signal entering a word recognizer and no perceptual input, an error signal goes to the reference inputs of lower-level systems, initiating reference signals there. Recall that a branch of the reference signal loops into the perceptual input function (Figures 3 & 4). If there is no perceptual input from below, it is this local echo of the reference signal that is controlled. Any perceptual input from below, however weak, is added to and strengthens that minimal (imagined) perceptual input. In this sense, the control system that recognizes the word *hit* is "looking for" its second argument word. Possibly it is controlling a word with broad selection, a classifier word (as e.g. *mammal* is a classifier for *dog*, *cat*, etc.) or an indefinite word of broadest selection such as *something*. Then the mechanism of "looking for" the proper argument word in input would be by associative memory linking a set of words whose narrower selection fits within the broad selection of the indefinite or classifier word.

To flesh out this proposal, we will now turn to how signals of different strengths compete in a massively parallel PCT model of cognitive functioning.

Parallel processing in the perceptual hierarchy

Discussions of hierarchical control in the PCT literature generally concern how a perception is controlled by dynamically varying the reference values for perceptions at lower levels.[56] That's on the output side, the generation of behavioral outputs as means of controlling perceptual inputs. The proposal here concerns a corresponding dynamism in the recognition of perceptions on the input side, where ambiguity must be taken into account.[57]

This may be applied to Operator Grammar is as follows: An operator word O_i is input and recognized. For each argument requirement of O_i, the recognizer of O_i sends a signal to the reference input of a recognizer for each of the words $A_1, A_2, .. A_n$ that could be in that argument position. Each recognizer is a 'demon' which (because of that reference input) is now 'looking for' its input and 'yelling' to the

---

[56] For example, "Each new level of perception creates a new class of entities that can be controlled by varying reference signals at the next lower level" (Powers 1988:279).

[57] Selfridge (1959) is a somewhat similar bottom-up model of how we perceive the same visual configuration even as the perceived object is rotated, translated (shifted around), and resized.

extent that it finds it; that is, each is now sending up a perceptual signal with strength proportional to the extent that its input function is able to construct a signal from the perceptual input that comes to it. Initially, this comes from the weak 'imagination' copy of the reference input that branches back to the perceptual input. If word A_i is recognized (if input matches reference), one branch of the perceptual signal for A_i is input to the recognizer of O_i, and consequently a stronger O_i signal is passed along to potential operators over it.

There is a dependency relationship between an operator and its arguments. Since an operator is asserted of its argument, it cannot be asserted unless its argument requirement is met. This may be represented graphically as a dependency tree, as in Figure 8.
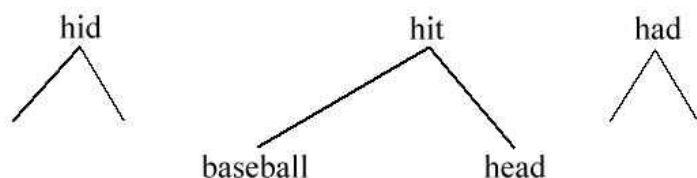

Figure 8. Dependency tree diagrams showing the relationship of the operator hit *to words that satisfy its argument requirement, and the operators* hid *and* had *whose argument requirement has not been met.*

Hit is in the O_{nn} class, that is, it requires two primitive arguments *N*, two words that do not require any other words to be present. The words *baseball* and *head* satisfy the argument requirement of *hit*, and that strengthens the perceptual signal for the word *hit*, which might in turn meet the argument requirement of some second-order operator (such as *remember* in *I remember that the baseball hit his head*, or *but* in *the baseball hit his head but he wasn't hurt*). As described above, the acoustic and articulatory perceptions activate recognizers for *hid* and *had* as well as *head*, but they are operator words, so they cannot satisfy the argument requirement of *hit*. Furthermore, the signals from their recognizers are comparatively weak because their argument requirements have not been met.

The more complete a partial dependency structure is, the stronger the signal from the recognizer of its top operator. However, weaker signals are still present. This can account for the sudden shift of construal that is most obviously seen in garden-path sentences such as *The horse raced past the barn fell down*, or Lashley's example,[58] *Rapid righting with his uninjured hand saved from loss the contents of the capsized canoe*, where until the last words are heard the first word is understood to be *writing*.

The population of words that are candidate arguments of an operator is determined not just by what is immediately spoken or written, because words and recurrent phrases are associatively linked as they are stored in memory, and this helps us to associate current input with relevant background information. The associations in memory derive from all the word-dependencies of previously-assimilated discourses, including the larger-scale structures of language due to word-sharing across conjunctions and repetition in successive periods of a discourse. The hypothesis here is that each such association is the effect of physical neural connections sending nerve-signals to recognizers of input words and to the operators and arguments with which those words are associated in memory. In the PCT architecture, when the perceptual input is *it's going to rain*, common-knowledge word dependencies associated with the word *rain* are activated.

Of course non-language perceptions are also associated with the words and controlled in imagination, and these provide additional pathways for associative memory. Non-language perceptions have a much broader role which we will take up in due course, but for the present discussion we will pretend that only word associations are involved. (This simplifying fiction is necessary for any grammar that is not embodied with sensory input, or a simulation thereof.)

Classifier words and indefinites (words of inspecific reference, such as *a thing*, *something*) support associations from specific statements to generalizations of them. Conversely, generalizations using classifier words provide a basis for assessing the acceptability or expectability of novel assertions. Harris (1982.6, 1991.78) suggested how to extrapolate the acceptability of a novel utterance from the

---

[58] Lashley (1951.120). Cited by Harris (1991:98).

acceptability of a familiar one with classifier words in order to explain how "sentences which we are certain were never said … are immediately understood as grammatical." Thus, for the sentence (1) *The very young thrips maneuvered itself into a narrow crack in the gray bark*,

> *… is an insect* is in the selection of the infrequent word *thrips*, as in the classifying sentence *A thrips is an insect*; and the common *move* is in the classificatory selection of the somewhat less frequent *maneuver oneself*. These selectional relations classify the word combinations of (1) under the word combinations of a sentence which is far less infrequent or unusual than (1), namely *The insect moved into a crack*.[59]

In a perceptual control hierarchy, all that is required is that the recognizer for each of the potential argument words send a branch of its perceptual signal to the input function of the classifier word, and that the classifier word send a signal to each potential argument word. When this weak signal to the recognizer for some word is strengthened by other perceptual signals, the effect is of that recognizer 'yelling' that much louder.

What is proposed here amounts to a knowledge base of word dependencies. When classifier words are used to support analogic extension and generalization from familiar to unfamiliar word dependencies, it constitutes an acceptability model for determining how sentences differ in acceptability depending upon context. In the PCT model of brain function, these differences of acceptability correspond to differences of signal strength. Where the incoming words fit into more than one possible arrangement of operator-argument dependencies, the catena with the strongest signal is integrated into the knowledge base in preference to other possible readings. (The topmost operator in a catena sums the strengths of its arguments and of their arguments in turn, etc.) Ambiguity is far more pervasive in language than we realize (unless we are trying to parse it with a computer), because this automatic process in our brains eliminates all but the strongest resolutions without our ever becoming aware of the weaker alternatives. This is also the basis of the judgment of relative likelihood which is the criterion for transformation.

The recurrence of domain-specific classifier words in a discourse, and the efficacy of using domain-specific definitional and encyclopedic 'common knowledge' sentences to fill in gaps so as to satisfy the word-sharing requirement for conjoined sentences and to regularize the repetition requirement for discourse coherence, provides a basis for identifying the subject-matter domain to which a discourse belongs.

Dependency trees and subtrees may be linearized in various ways into the sentences of discourses. I refer to this as discourse-level periphrasis, as distinct from sentence-level paraphrase.[60]

---

[59] Harris (1991.78). This is related to the 'Pronominal Approach' for descriptive linguistics developed by Karel van den Eynde and his colleagues. See van den Eynde et al. 1988, van den Eynde et al. 1998, van den Eynde & Mertens 2001, van den Eynde et al. 2002, van den Eynde & Mertens 2003, Kirchmeier-Andersen 1995, Melis & Eggermont 1994. Independently, Mattick suggests a kindred method for characterizing a 'laboratory language' "for the description of techniques applicable to a number of related fields" (Harris et al. 1989.156).

[60] I introduced this distinction in my A.M. thesis in 1970.

Nonverbal perceptions and subjective meanings

In addition to the objective information that is disclosed by the distributional methodology of empirical linguistics, there are other, subjective aspects of meaning. Language perceptions are associated in memory with the universe of non-language perceptions. Many of these associations are 'subjective', that is, they are idiosyncratic to the individual language user's experience, and not conventionalized in the mutual calibration and intersubjective agreement by which we continuously recreate and slowly change our language in the course of using it.[61]

The perceptions that constitute language are necessarily of the same types as those that constitute the rest of our perceived universe, constructed in the same hierarchical levels of perceptions. Figure 9 illustrates perceptual inputs and error outputs for language in parallel with inputs and outputs for non-language perceptions in a diagram similar to that presented in Figure 4. Because of the complexity of the drawing, only one interconnection is shown between language and non-language. Powers (1973) posits that the perceptual signal for a word, at the Event level, contributes to the perceptual input for a category perception.[62] Most words of the primitive argument *N* class refer to concrete objects, that is, perceptions on the Configuration level. In the model generally accepted among PCT researchers, a Configuration perception (with all its lower-level detail of color and texture sensations and intensities) is input to a Category-level input function that also receives a signal from the Event perception of the word *brick*. Perhaps a Sequence-level perception calls for a brick as part of its next step, that sequence being controlled in turn by a Program-level perception "building a wall".

---

[61] Subjective meanings are also associated with aspects of language other than the informational structures constituted by operator-argument dependencies and the lexicalization of phrases, and some of these are standardized in other ways, notably perceptions of identity and identification that are studied in sociolinguistics.

[62] The alert reader may be disturbed by the recollection that a syllable is also at the Event level. A perception at a given level may be constructed of other perceptions at that same level. Consider the perception of a letter A (a configuration) the pixels of which are also letter-configurations.

The two highest levels (Systems and Principles) are not shown here.

Programs

Sequences

Categories

Relationships

Events

Transitions

Configurations

Sensations
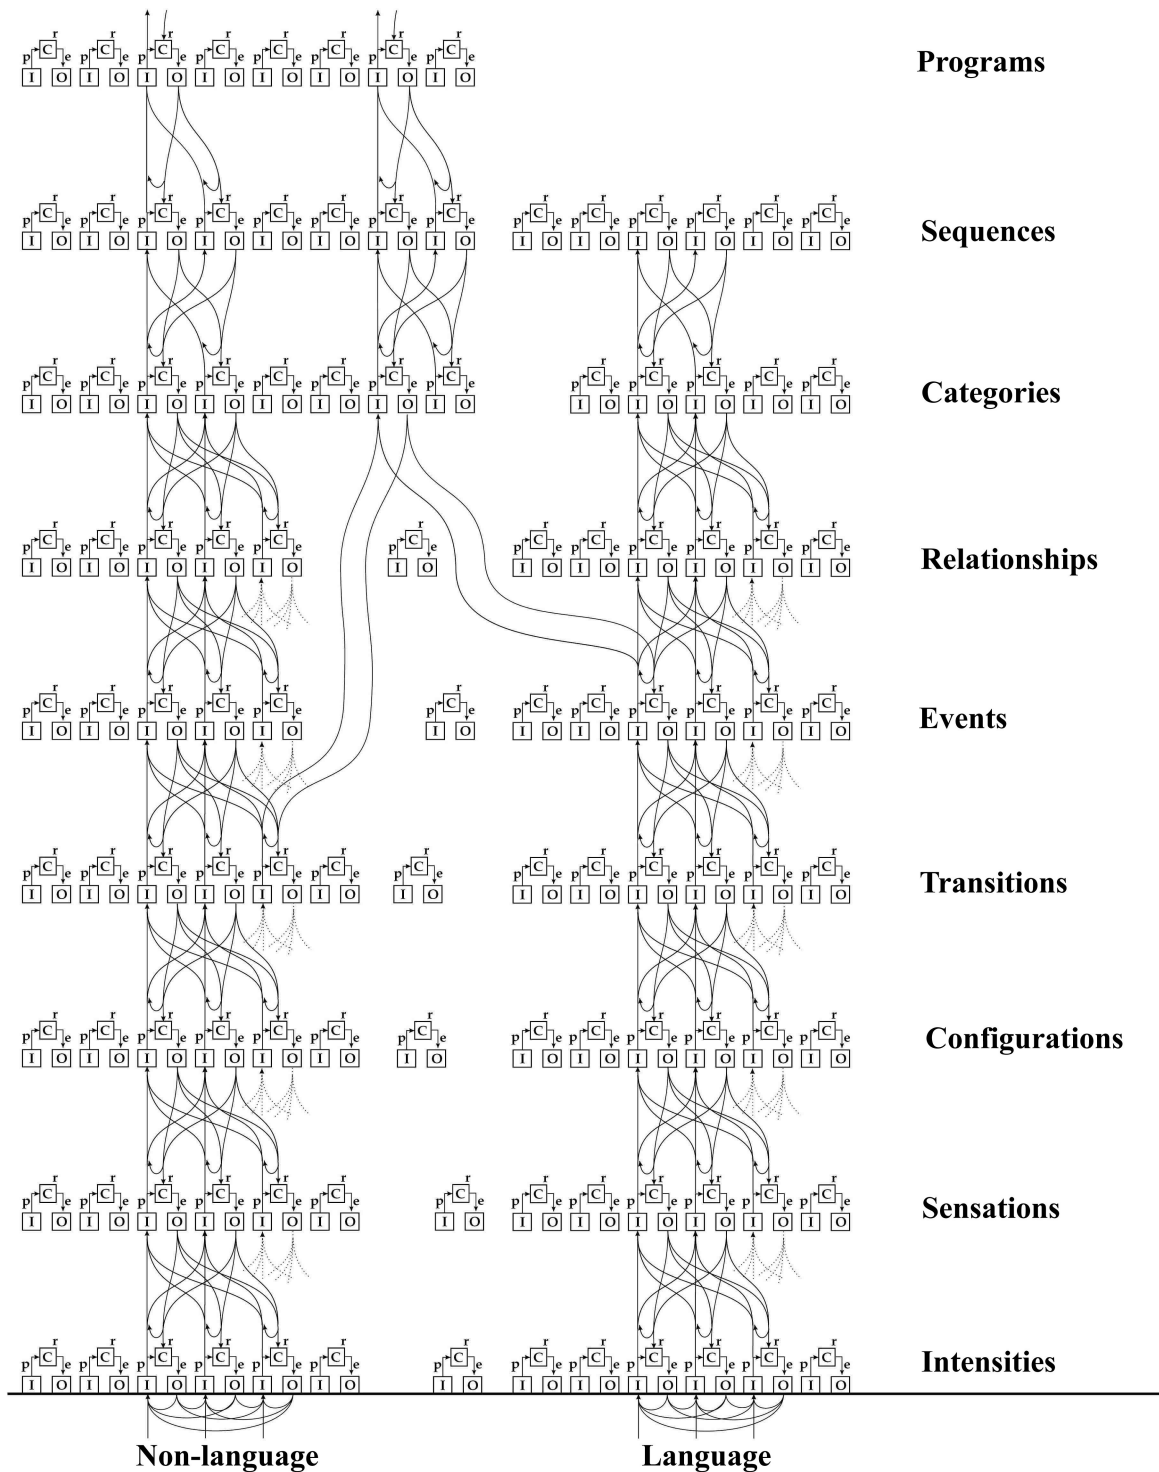
Intensities

Non-language

Language

Figure 9. Connections between the language perceptions that constitute objective linguistic information and the non-language perceptions that are experienced as subjective meanings. At the top, a program perception is controlling argumentation while concurrently controlling referenced perceptions. (Only a few input or output lines are shown connecting levels.)

This diagram is of course vastly simplified. The average human brain has about 86 billion neurons. Each receives neural impulses through many dendrites, with a total of some ten thousand branches. Each sends neural impulses through an axon, with as many as a thousand branches. Each branch terminates in a synapse and may have *en passant* synapses partway along their length. The number of dendrites connected from a single synapse to a single cell, and especially the number of dendritic spines,[63] weights the received signal. Neurons are capable of constantly strengthening and weakening these connections, making and breaking them. But while these facts of neurophysiology may emphasize the ludicrous simplification of the diagrams in this paper, they should also enhance the credibility of the PCT model which they so inadequately limn. The point of this diagram in particular is to communicate the relationship between the orders of perception that are constructed in the brain, and the crucial role of language as a socially constructed system of collectively controlled perceptions.

Anything can be organized in categories, just as we can speak of anything for which we have vocabulary. Whatever systems there may be that categorize perceptions or recognize categories, their inputs can include signals from the systems that recognize words, linguistic constructions, and catenae, and from and the systems that recognize dependencies among these as linguistic information.

The several kinds of linguistic perceptions are hierarchically related in our conception of grammar. In particular, the dependencies of operators and arguments are hierarchical in each sentence. However, it may be important to emphasize the obvious point that they do not reside at successively higher levels of the perceptual hierarchy. For example, both phonemes and words occur in strings that are controlled at the sequence level, and both phonotactics and operator-argument dependencies are controlled at the relationship level. It is true that most primitive arguments correspond to configuration perceptions (discrete concrete objects), but e.g. *vacuum* does not; first-order O_n, O_{nn}, and O_{nnn} therefore generally also correlate with low-level perceptions, unless metaphoric, but the nonverbal perceptual correlates (subjective meanings) of second-order operators like *surprise*, *cause*, *begin*, *(be a) fact*, vary widely, depending upon their arguments.

The relationship of language perceptions to non-language perceptions is not as simple as that of symbol to referent. The denotation of even a concrete noun in the *N* class is in the non-language perceptions with which it is associated. (We have no non-perceptual knowledge of it.) The non-language perceptions that are associated with a lexical item (a word or a recurrent catena) may be products of collective control, that is, manufactured and maintained according to certain public specifications: a AA battery, a yardstick. It is implausible to say the same of natural phenomena (a pulsar, a chickadee). In both cases, regularities of language are projected onto the non-language perceptions, categorizing them, but in the latter (and far vaster) case, the attributes of the category are not enforced by collective control. Denotations are not so easy to distinguish from connotations, then, since both reside in the non-language perceptions, and both appear to be organized as such by language. A diachronic perspective may be helpful here. Language and the associated categorization of non-language perceptions presumably have evolved together (Mirolli & Parisi 2006) over the span of an estimated 3000 generations that humans have been speaking language.[64] Collective control of language perceptions has enabled more effective collective control (and social standardization) of non-language perceptions, and vice versa.

Some consequences of PCT for linguistics

PCT shows that any proposal in linguistic theory must demonstrate that the objects and relations that it proposes can be perceived, are in fact perceived by language users, can be controlled, and are in fact controlled by language users. Standards of parsimony (Ockham's Razor) forbid the multiplying of explanatory principles if you can do without them. There are many theoretical constructs in the field, such as binary merge and feature checking, that do not meet this standard (Clark & Lappin 2011.214). empirical linguistics does, because everything in the resulting theory of language and information derives from tests for controlled perceptual variables.

---

[63] Sala & Segal (2014).

[64] Based on the current estimate of 45,000 BCE for the exodus from Africa, presumed to have been enabled by and preceded by the development of language, and 15-20 years per generation.

Methodological consequences

The basic methodology of PCT is to construct and test a simulation of the continuous behavior of an individual organism. Runkel (2007) calls this the method of 'testing specimens'. Prior to PCT, the only way that researchers in the life sciences knew how to be scientific was to emulate the IV-DV methodology of the physical sciences. The standard methodology of psychology and the social sciences, which Runkel calls 'casting nets', seeks statistical correlations between 'stimuli' (independent variables) and a range of 'responses' (dependent variables), averaged across many experimental runs with many individuals.

Population statistics are useful, so long as they are not misapplied to a different population. They do not explain the behavior of individuals.[65] However, in statistical learning[66] there is a curious intersection of the properties of populations and the properties of individuals. Our means of recognizing and controlling perceptions are established by evolution and refined by learning, and learning and evolution both result from Neo-Darwinist processes of blind variation and selective retention (Cziko 1995, 2000). Random variations across a population of individuals are the basis of evolution. Learning processes appear to depend upon the frequency and accumulation of perceptual signals and changes in the size of populations of neural connections.

Research enterprises like operant conditioning impose 'selective pressure' on these internal 'populations'. In a typical experiment of this type, an animal is starved to 85% of its normal body weight. The experimenter is coercively disturbing the level of blood sugar, an 'intrinsic variable' that is essential to life, using coercive means that overwhelm the subject's capacity to control. In this 'establishing condition' for the experiment, the animal explores the artificially constrained experimental environment and by trial and error it hits upon some activity that produces a bit of food. By repetition, it then establishes a reference for controlling the relevant perception (e.g. pressure on a bar) as means of controlling the perception of eating and restoring its blood sugar to its biologically established setpoint. No explanation of what the organism does is given or even attempted. This is no more than an instrumental means of influencing the organism's setting of its own internal reference values. It is not a theory of behavior.

A PCT simulation is not very complicated mathematically, the simultaneous integration of three equations, but its working is subtle and accounts for all the complexity of behavior. To be accepted, such a simulation must replicate the measured behavior with greater than 95% accuracy (and preferably greater than 99% accuracy).

The methodology of PCT requires

    A. Identifying a controlled variable (CV). The Pair Test and other substitution tests of empirical linguistics are tests for CVs.
    B. Measuring the state of that variable, disturbances affecting that state, and actions of the subject which counter those disturbances.
    C. Inferring the subject's internally-mainained reference value from the value at which the subject controls the CV.
    D. Creating a working generative computer model that controls the state of the CV according to that reference, resisting random disturbances with only as much precision as the subject did.

One cannot overstate the importance of the Test for the Controlled Variable. Here is a detailed description:[67]

    1. Select a variable in the environment that you think subject *S* might be controlling.
    2. Identify by what sensory paths *S* perceives the environment variable *EV*. If *S* cannot perceive *EV*, return to (1) for another guess. Just because you can perceive *EV* does not mean that *S* can perceive it, or that *S* perceives it in exactly the same way as you do.

---

[65] Runkel (2007), Gould (1996), Heimann (2003), Marken (1992, 2002), Plooij (1984, 2003), Powers (1973, 1989, 1992, 2005, 2008), Powers, Clark, & McFarland (1960), van de Rijt-Plooij & Plooij (1992, 1993).

[66] On statistical learning theory, see e.g. Saffran et al. 1996b, Misyak et al. 2012, Karuza et al. in press

[67] Based on Runkel (2003:77-79), which in turn was based on Powers (1973:232-246; 2005:233-248; 1979:110, 112), and (I am sure) correspondence and conversations with Powers as e.g. in Powers & Runkel (2011).

3. Predict the effect of applying a disturbance *d* to *EV* if *S* was not controlling *EV*. If possible, measure the effects of disturbances on *EV* when *S* is not present and you know that *EV* is not controlled.

4. Experimentally apply various amounts and directions of disturbance *d* directly to *EV*. These should be gentle disturbances, not to overwhelm *S*'s ability to control. Avoid unintended side effects that might disturb other variables that *S* might be controlling.

5. Measure the actual effects of *d*.

6. If the effects are as predicted in (2), stop, and return to (1) for another guess.

7. If the effects are markedly smaller than predicted, identify the cause of opposition to the disturbances. That cause may be caused by *S*'s behavior. You may have identified the environmental feedback function.

8. With reference to what you learned in (2), block the sensory pathway between *S* and *EV*, so that *S* cannot perceive it. If *d* continues to be opposed, return to (2) to identify the sensory input by which *S* is perceiving *EV*. Consider that *S* may be using more than one sensory modality concurrently (as in the McGurk effect).

9. If after blocking sensory input in (8) the values of *EV* as disturbed by *d* are as predicted in (3), or if *S* acts to remove the block to sensory input (treating it as a disturbance to control), then *EV* is a controlled variable *CV*.

Having identified a *CV*, the further methodology of PCT is to collect data to be modeled and then to construct a computer simulation that uses negative feedback control to achieve the same results. For models of control by means of the lowest four levels of the hierarchy (Intensities, Sensations, Configurations, Transitions), this means that one must carefully measure three quantities as they co-vary:

qd   The disturbance quantity
qi   A quantity representing *S*'s sensory input
qo   A quantity representing *S*'s behavioral output

Computer simulations of control of (or by means of) the lower four levels have been very successful—extraordinarily so, by the standards of the social sciences. In psychology and sociology generally, any statistical results that are even a little better than a coin-toss are considered significant. A PCT simulation that is less than 95% accurate is not yet ready for publication. However, almost all modeling thus far has been at fairly low levels of the hierarchy, for example Marken's control model of how a baseball player catches fly balls (Marken 2005). The organization of a model that demonstrates the effect of disturbing the sounds of the subject's own speech in real time was sketched above, but lacking data from the experimenters, or access to the specialized computer hardware that they used, no model has been constructed.

An immediate question for linguists is: How do we measure the perception of a word, phrase, sentence, or text as a variable quantity? Lower-level perceptions are analogue variables, but from the Event level up perceptions are binary, either *yes* the perception is there or *no* it is not. Such a perception can be felt to be more strongly or less strongly represented by a particular exemplar, there can be uncertainty and disagreement about marginal cases, and priming is prejudicial (by activating references that are controlled in imagination), but any perception at a Relationship, Sequence, or higher level disregards such distinctions.

For modeling of control of words and word dependencies, it is useful to remember that each level of the hierarchy exists in a perceptual environment constructed by the levels below it. This suggests that we do not have to recapitulate all of phonetics and phonology in a model of Operator Grammar. Unless we are actually concerned with modeling variations of pronunciation, intonation, etc., we can adopt the convenient fiction that discrete words and morphemes are the lowest-level perceptions to be modeled.

For word dependencies, one type of disturbance might be a willful misinterpretation exploiting the ambiguity that is always potential among the paraphrase set of sentences produced by the reductions.

It can be difficult to disturb language variables experimentally, and could encounter ethical constraints, but often naturalistic observation suffices. (McClelland (forthcoming) has pertinent

remarks on uses of naturalistic observation in sociology.) For example, people often correct perceived errors in another's speech, either overtly with prescriptive speech, or in the course of repeating or paraphrasing what was said. There is a problem somewhat analogous to indeterminacy in physics, called the 'observer's paradox': people change how they talk when they know that their way of talking is being observed as such. We need "to find out how people talk when they are not being systematically observed; yet we can only obtain these data by systematic observation" (Labov 1972:209). Eliciting techniques have been developed by linguists (e.g. Harris & Voegelin 1953) and sociolinguists (e.g. Labov 1970, 1972, Labov & Waletzky 1967). This is complicated by ethical requirements for informed consent of subjects. Avoidance of these difficulties is one of the attractions of corpus linguistics working with large collections of naturalistically collected texts.

PCT research is a largely untouched field, where anyone who ventures in has an excellent chance of making valuable and lasting contributions.

Connectionism vs. computationalism

Connectionism emerged in the late 1980s (Rumelhart et al. 1986a,b) as a way of simulating Hebbian learning in a digital computer, that is, the strengthening and weakening of synaptic connections and neural signals as a consequence of experience. Cognitive structure is seen as emergent from these processes. In many respects, connectionism strongly resembles the reorganization system in PCT. The crucial difference is that connectionist systems are opaque, and can only be described in very general and low-level terms, e.g. by specifying the learning algorithm and the number of units involved.

Theoretically, the perceptual control hierarchy could emerge from a connectionist model, and could be demonstrated to control its input perceptions, using standard PCT methods. This is uncomfortably close to the infinite monkey theorem (that a random process will 'almost surely' produce any given body of text, such as the writings of Shakespeare, eventually). The perceptual control hierarchy is the product of two reorganization systems operating on the principle of blind variation and selective retention. The primary one has produced variants over a very large number of generations of autonomous control systems, of which only some succeeded in bringing progeny to reproductive maturity; the secondary process tunes perceptual input functions and reference input functions. The first is evolution, the second learning.

Computationalism is the view that cognition is accomplished by the syntactic manipulation of symbols. The philosophy of mind associated with Generative linguistics and with classical Cognitive Psychology is closely identified with computationalism, and strongly resists Connectionism. Computationalists propose that specialized modules of the brain are responsible for domain-specific modules of cognition; connectionists assume only relatively low-level general-purpose neurological structures. Computationalists propose syntactic rules for internal manipulation of explicit symbolic structures (mental models); connectionists propose computational models of the weighting of connections between neurons to simulate Hebbian learning from environmental stimuli. Computationalists often posit domain specific symbolic sub-systems designed to support learning in specific areas of cognition (e.g., language, intentionality, number), whereas connectionists posit a small set of very general learning mechanisms. Demonstrations that connectionist systems can accomplish high-level processes such as logical reasoning have settled the matter in favor of connectionism among the majority of neuroscientists, but these recent developments have not yet affected psychology, philosophy of mind, and the affiliated branches of linguistics.

More recently, dynamical systems have been proposed that recognize the continuous mutual influence of cognition, behavior, and the environment. They are formalized by several differential equations describing the system's state over time in a trajectory space (van Gelder & Port 1995, van Gelder 1998). While some dynamicists would frame the above debate as computation vs. dynamic systems, and even urge that this approach supplants connectionism, connectionist and dynamic systems seem rather to be complementary. PCT can be understood as a dynamical system that is computationally much, much simpler, because it works within the hierarchy of negative-feedback control systems that has been established by evolution. Connectionist and dynamical schemes may have a place describing

reorganization, both in evolution and in learning. However, so far random reorganization in the manner of e. coli chemotaxis has been sufficient (Marken 1989, Pavloski et al. 1990).

Consequences for the metalanguage

PCT suggests a reconsideration of metalanguage assertions that two word-occurrences have the same referent; such assertions enable reduction to the *wh-* pronoun of a relative clause. For example, *I passed a pen, which I had borrowed, to my neighbor* is reduced from *I passed a pen—a pen [same as prior] I had borrowed—to my neighbor*. In a PCT account, we may suppose that the work of the metalanguage assertion *same as prior* is done by the association of both occurrences of *pen* with one non-language perception of a pen. Although I did not have PCT at the time as a theoretical framework to justify it, this was the basis of the proposal in (Nevin 1970) that, just as word-recurrences are grouped in the same column for discourse analysis, so the dependency hierarchies for successive sentences in a discourse (or set of discourses in a sublanguage) could be merged into a heterarchy or network in which the repeated words occur only once. Alternative linearizations into sentence-sequences I referred to as periphrasis, distinguishing that discourse-level phenomenon from sentence-level paraphrase.

According to the theory outlined in Section 3, a word is one of many perceptions entering the input function of a Category recognizer. When you ask me if I have *a pen*, the word identifies the perceptual category; any pen will do. If you ask me to *lend the pen on the left*, the phrase identifies an individual member of that category, distinguished by its spatial relationship to something else (perhaps another member of the same category). Operator Grammar expands that phrase to *lend that which is a pen; a pen (prior same as penult) is on the left of something*.[68] The word *that* in the source of the definite article *the* (Harris 1982:237) specifies the individual instance. One is left to wonder how that metalanguage sameness assertion *prior same as penult* becomes performative as well as descriptive. In a PCT model of an embodied grammar, the sameness is performative because the two word-occurrences in the linearized sentence are both linearized from the same node of the pre-linear network of word dependencies.

Harris noted that "the meaning of *same* (or *said*) as condition for *wh-* is not precise and may vary with the circumstances…."[69] To illustrate this, in the first sentence of the first paragraph of this section above, the clause that is conjoined by semicolon could be reduced instead to a relative clause, yielding

> *PCT suggests a reconsideration of metalanguage assertions that two word-occurrences have the same referent which enable reduction to* wh- *pronoun of a relative clause.*

The phrase "*that two word-occurrences have the same referent*" makes it difficult to recognize the antecedent of the relative pronoun *which*.[70] If we elide that phrase, the relationship is clear:

> *PCT suggests a reconsideration of metalanguage assertions [...] which enable reduction to* wh- *pronoun of a relative clause.*

In Operator Grammar, the source of both the former and the latter is

> *PCT suggests a reconsideration of metalanguage assertions that two word-occurrences have the same referent; metalanguage assertions that two word-occurrences have the same referent [same as prior] enable reduction to the wh- pronoun of a relative clause.*

It is not so easy to see how a phrase as complex as *metalanguage assertions that two word-occurrences have the same referent* is associated in any simple way with a single non-language perception. It seems much less problematic to accept that the sentence is after all affirming that both operators over that phrase, *enable* and *reconsider* (which becomes *reconsideration of* under *suggest*), have the same argument in the non-linear word-dependencies from which the sentence was linearized.

---

[68] Harris (1982.87–97, 227–229, 237–243).

[69] Harris (1982:227). At the ellipsis in this quotation there is a reference to pp. 89-92 of this grammar.

[70] By its length and because of ambiguity: the pronoun could be a reduction of *occurrences* instead of *assertions*. This is why the semicolon and *such assertions* appears instead in the original text above.

They are different word-occurrences in the linearized output, but in the word-dependencies prior to linearization they are the same. The metalanguage adjunct *same as prior* asserts this sameness in a linearized form, but need not be represented by control of metalanguage word-perceptions in either the speaker who is formulating the sentence or the recipient who understands it and integrates it into her own internal knowledge-base of word dependencies. The feasibility of this is an open question for research.

Formalisms, computers, and robots

To be implemented on a digital computer, Operator Grammar must be formalized. The operator-argument dependencies are easy to formalize. The reductions are very difficult to formalize. An attempt might be made on the lines of Lexicon-Grammar (Gross 1994). In Lexicon-Grammar, each word in the lexicon labels a row in a matrix, and each column of which represents a transformation. A cell is marked if the word on its row participates in the transformation on its column. Such a matrix for reductions might require more than two dimensions.

Computer implementations typically represent semantics of language with reference to a structured database, with an interface converting sentences to database queries. Generalization, metaphor, and other processes of analogy that are common in human cognition are not easily expressed in such an implementation, and of course it cannot claim to model how the human brain works. The semantics of Operator Grammar could be represented by a set of domain-specific acceptability models. An acceptability model is a knowledge base resulting from the analysis of a set of discourses. This naturally extends to a set of intersecting sublanguage grammars, with phrasal lexical items and stricter threshholds of acceptability, and classifier vocabulary would naturally support processes of analogy.

Computer languages are symbol systems generated by formal grammars of the PSG family. (Backus-Naur Form, or BNF, can be seen as a notational variant of PSG rewrite rules.) The semantics of a computer language is in its transformation (by compiler, interpreter, or the like) into machine-language instructions which set bit patterns in memory media. This is clearly not how the brain works.

A more human-like computer language, implementing Operator Grammar, would probably have to be embodied in a more human-like computer which interacts by controlling feedback from its environment concurrently with feedback from its internal environment of linguistic acceptability-models. Ideally, this would be an analog computer, rather than a simulation of analog processes on a digital computer.

Some epistemological considerations

Epistemology is the philosophical inquiry into what knowledge is and how we acquire or construct it. Whether knowledge comes to us by way of our senses or by some other means such as 'innate ideas', our awareness of it is a subjective awareness of perceptions. According to PCT—and I am aware of no evidence to the contrary—perceptions are neural signals in the brain, and our awareness of perceptions, of the qualia of experience, is something that still eludes scientific specification. No mechanism for "the rider in the chariot"[71] has been identified.

It follows that we can have no direct knowledge of the phenomena that are studied by the natural sciences. All that we know of the world is perceptions constructed in our brains. Even when we use more or less sophisticated tools and instruments of observation and measurement, we have no direct apprehension of what is Really going on, Kant's famous *Ding an sich* (the "thing in itself"). We live in a universe of perceptions which we take to be Real. We get away with this because our experience controlling our perceptions in our environment continually tests the veridicality of our perceptions, and we do this using perceptual organs that our species have been putting to the test of survival for a very long time. Other species, many of them with far longer experience, construct very different perceptual universes; Lettvin et al. (1959) gives the well-known illustration of what the frog's eye tells the frog's brain.

---

[71] *Katha Upanishad*, 3.3. In that metaphor, PCT is about how the chariot works.

The conclusion is actually rather obvious that the map of perceptions is not the territory which is perceived. There are no cabbages, kings, or quarks in the human brain, only perceptions represented by the staccato electrochemical discharge of neurons and experienced subjectively by as yet unspecified means.

> [E]ven our actions have consequences that are known to us only after they have been encoded by the nervous system. To ask how the objects and events of the real world are encoded into neural signals is to miss the point: objects and events are already in the form of neural "codes" by the time we become aware of them. This is how neural codes look, taste, feel, and so on. The situation was put best by von Glasersfeld: the brain is not the Black Box: the environment is. Our actions seem to affect a real objective world, whether those actions be the informal ones of everyday living or the formalized ones of scientific experimentation. To recognize that the knowable effects of actions are really activities in the brain, already encoded and transformed, is to be forced into an unorthodox view of scientific objectivity. What can objectivity mean when it is applied to a world that is subjective from the start? (Powers 1988.176–177)

To answer that question—What can objectivity mean?—it will be of no help to go down a rabbit hole with Alice regarding the undeniably creative role of perception, as reflected in the idiom "what do you make of that?" But nevertheless on the way to an answer we must acknowledge that all of culture, and all of the notorious divergence of standards and expectations of communities from primitive tribes to scientists working in paradigms, all the babel of human diversity, results from the activity of our brains constructing complex perceptions from simpler ones.

We must acknowledge, too, that an unruly proportion of what we perceive is supplied in imagination. How imagination works was described earlier. A principal purpose of scientific method is to stabilize against 'subjective' influences in the construal of data, such as confirmation bias, a well-known but too often overlooked perceptual illusion.

We can be aware of nothing except perceptions. A theory or an abstract model is itself a new perceptual construct. It can be tempting to take a successful model to be Reality, but this is the error called reification. Whether "the fool on the hill" sees the sun going down or the world spinning round, both are perceptions. A map is not the territory that it describes.[72] Science proves nothing, its conclusions are always provisional, and proof is possible only for mathematics and logic. The best answer that the physical sciences can offer to our question is that scientific objectivity is a function of "intersubjective agreement" among researchers who replicate experiments and report the same results.[73]

The natural sciences such as physics and chemistry construct abstract theoretical models because they must. However, empirical linguistics and PCT each resolve this dilemma in a way that is different from that of the physical sciences, because the phenomena that they investigate are fundamentally different from the phenomena studied in the physical sciences. The physical sciences seek to learn about the Reality beyond our perceptions, but the subject matter of empirical linguistics and PCT consists only of perceptions.

---

[72] Korzybski (1994.58). There are more difficult challenges. "The doctrine that the world is made up of objects whose existence is independent of human consciousness turns out to be in conflict with quantum mechanics and with facts established by experiment" (d'Espagnat 1979). See Gröblacher et al. (2007), summarized in Cartwright (2007). See also the comprehensive account by Ryckman (2005).

[73] Lack of replication is becoming recognized as a grave problem; for discussion, see the references cited at http://goo.gl/4CmLIi. Perhaps because of a social need for reliable authorities, it is not always remembered that scientific theories are never proven. Proof is possible only for mathematics and logic. At best, a theory may resist disproof better than any alternative theories long enough to become accepted as the consensus view within a field.

The core methodology of PCT, the test for identifying controlled variables,[74] has the effect of establishing intersubjective agreement between the investigator and the subject organism. In the measurement of input, output, and disturbance quantities, PCT is under the same constraints as the physical sciences.

Phonetics is among the physical sciences, but the elements of language, beginning with the phonemic contrasts, and continuing through the words and morphemes, the word dependencies, the selection of each word, the canonical and alternative linearizations, and the reductions, are all perceptions. More particularly, they are perceptions of the perceptions that others in the speech community are controlling, collectively controlled perceptions reflecting and embodying intersubjective agreements. Because language is a product of intersubjective agreement, and consists only of perceptions, empirical linguistics is objective in ways that other sciences cannot be. And for this reason, the information carried by language (the departures from equiprobability) is objective information.

Empirical linguistics characterizes the objective information in language, and PCT shows how an embodied cognitive system associates subjective meanings with that objective linguistic information.

Conclusion

Perceptual Control Theory (PCT) can integrate all the disciplines concerned with living things into a common methodological and conceptual framework with a unified scientific foundation. An understanding of language as the control of perceptions is crucial for applying PCT to the human sciences. Empirical linguistics as developed by Zellig Harris is a science of language as a perceptual construct. Its methodology essentially uses the PCT 'Test for controlled variables' to identify phonemic contrasts, to test the substitutability of words in context of other words and more generally in sentence-forms, and to differentiate satisfiers of sentence-forms as to their acceptability. Tests of acceptability-difference establish the transformations and guide the factoring of transformations into operators, arguments, and reductions and the development of grammars of discourses and of sublanguages. The rest is combinatorial ("distributional") analysis of the current state of the shifting patterns of stability that users of language have created and recreated in processes of collective control over countless generations. Other studies that are also possible on this basis, such as investigations of variation and change in language, the possible further evolution of language (which is probably only beginning), the formal correlates of argumentation and other metadiscourse, and the relation of the formal information structures of language to mathematical and logical formalisms, are among future prospects for the unification of empirical linguistics and PCT.

References

Astington, Janet Wilde, & Margaret J. Edward. 2010. "The development of Theory of Mind in early childhood". *Encyclopedia on early childhood development*. [http://www.child-encyclopedia.com/documents/astington-edwardangxp.pdf]

Barsky, Robert F. 2007. *The Chomsky effect: a radical works beyond the ivory tower*. Cambridge, MA: MIT Press.

Bloch, Bernard. 1948. A set of postulates for phonemic analysis. *Language* 24.1:3-46.

Carey, Timothy A. 2006. *The method of levels: how to do psychotherapy without getting in the way*. Hayward, CA: Living Control Systems Publishing.

Carey, Timothy A. 2008. Perceptual Control Theory and the Method of Levels: Further Contributions to a Transdiagnostic Perspective. International Journal of Cognitive Therapy: Vol. 1, The Transdiagnostic Approach to Cognitive Behavioral Therapy, pp. 237-255.

Carnap, Rudolf. 1937 [1934]. *The logical syntax of language*. London & New York: Kegan Paul. [English transl. of *Logische Syntax der Sprache* (1934).]

Chomsky, A. Noam. 1951. Morphophonemics of Modern Hebrew. Unpublished Master's thesis, The University of Pennsylvania. [Held in the Rare Book collection of the Van Pelt Library. Extensive rev. Garland Press, Outstanding Dissertations in Linguistics Series, New York, 1979.]

---

[74] A detailed description of the test for controlled variables is given by Runkel (2003:76-79)

Chomsky, A. Noam. 1964. Current issues in linguistic theory. The Hague: Mouton.

Clark, Alexander & Shalom Lappin. 2011. *Linguistic nativism and the poverty of the stimulus*. New York: Wiley-Blackwell.

Gardner, Howard. 1985. *The mind's new science: A history of the cognitive revolution*. New York: Basic Books.

Garfield, Jay, Candida Peterson, & Tricia Perry. 2001. Social Cognition, Language Acquisition and the Theory of Mind. *Mind and Language* 16.9:494-541. [http://philpapers.org/rec/GARLPH]

Gleitman, Lila. 2002. Verbs of a feather flock together. *The legacy of Zellig Harris: language and information into the 21st century, Vol. 1: Philosophy of science, syntax, and semantics*, ed. by Bruce E. Nevin. Amsterdam & Philadelphia: John Benjamins, pp. 209-232.

Goldsmith, John. 1998. "An algorithm for the unsupervised learning of morphology". *Natural language engineering*. 12:4.353–371.

Goldsmith, John. 2001. Unsupervised learning of the morphology of a natural language. *Computational Linguistics* 27:2.153-198.

Goldsmith, John. 2009. "Segmentation and morphology". Alex Clark, Chris Fox, & Shalom Lappin (eds.) *Computational linguistics and natural language processing handbook*, pp. 364-393.

Gould, Stephen Jay. 1996. *The mismeasure of man*. New York: W.W. Norton & Co.

Grishman, Ralph & Richard Kittredge (eds.). 1986. *Analyzing language in restricted domains: sublanguage description and processing*. Hillsdale, NJ: Lawrence Erlbaum.

Gross, Maurice. 1993. Les phrases figées en français. *L'information grammaticale* 59:36-41.

Harris, Zellig S. 1946. From morpheme to utterance. *Language* 22.3:161-183.

Harris, Zellig S. 1952a. Culture and style in extended discourse. In Sol Tax & Melville J[oyce] Herskovits (eds.) *Indian tribes of aboriginal America* (=Selected papers from the 29th International Congress of Americanists, New York 1949, vol. III), 210-215. New York: Cooper Square Publishers. (Entire vol. repr. New York: Cooper Press, 1967; paper repr. in Harris 1970:373-389.)

Harris, Zellig S. 1952b. Discourse analysis. *Language* 28.1:1-30. (Repr. in Harris 1970:313-348 and Harris 1981:107-142.)

Harris, Zellig S. 1955. From phoneme to morpheme. *Language* 31.2:190-222. (Repr. in Harris 1970:32-67.)

Harris, Zellig S. 1962. *String Analysis of Sentence Structure*. (= Papers on Formal Linguistics, 1.) The Hague: Mouton. (2nd ed., 1964; repr., 1965.)

Harris, Zellig S. 1967. "Morpheme Boundaries within Words: Report on a computer test". (= Transformations and Discourse Analysis Papers, No.73.) Philadelphia: University of Pennsylvania. (Repr. in 1970.68-77.)

Harris, Zellig S. 1963. *Discourse analysis reprints*. (=*Papers on Formal Linguistics* 2). The Hague: Mouton. (Repr. of TDAP 3b, 3c, & 4a, Transformations and Discourse Analysis Project, The University of Pennsylvania.)

Harris, Zellig S. 1988. *Language and Information*. (=Bampton Lectures in America, 28.) New York: Columbia University Press, ix, 120 pp. [From lectures given at Columbia University, New York City, in Oct. 1986.]

Harris, Zellig S., Michael Gottfried, Thomas Ryckman, Paul Mattick, Anne Daladier, Tzvee N. Harris, & Suzanna Harris. 1989. *The form of information in science: analysis of an immunology subanguage*. Preface by Hilary Putnam. (=*Boston studies in the philosophy of science* 104.) Dordrecht/Holland & Boston: Kluwer Academic Publishers.

Harris, Zellig S. 1991. *A Theory of Language and Information: A mathematical approach*. Oxford & New York: Clarendon Press.

Harris, Zellig S. 1962. *String Analysis of Sentence Structure*. (= Papers on Formal Linguistics, 1.) The Hague: Mouton. (2nd ed., 1964; repr., 1965.)

Heimann, M. (Ed.). 2003. *Regression periods in human infancy*. Mahwah, New Jersey: Erlbaum.

Herrmann, Esther, Josep Call, María Victoria Hernàndez-Lloreda, Brian Hare & Michael Tomasello. 2007. Humans Have Evolved Specialized Skills of Social Cognition: The Cultural Intelligence Hypothesis. *Science* 317:1360–1366 (September 7).

Joshi, Aravind K. 1969. Properties of formal grammars with mixed types of rules and their linguistic relevance. *COLING '69: Proceedings of the 1969 conference on Computational linguistics*, 1-18. Association for Computational Linguistics Stroudsburg, PA [DOI: 10.3115/990403.990450]

Joshi, Aravind K. 1985. Tree-adjoining grammars: how much context sensitivity is required to provide reasonable structural descriptions? *Natural langauge parsing*, ed. by D. Dowty, L. Karttunen, & A. Zwicky, 206-250, Cambridge: Cambridge University Press

Joshi, Aravind K. 1999. Tree-adjoining grammars. In G. Rosenberg & A. Salomaa (eds.) *Handbook of formal languages* 69-123, Berlin: Springer Verlag.

Joshi, Aravind K. 2002. Hierarchical structure and sentence description. *The legacy of Zellig Harris: language and information into the 21st century, Vol. 2: Computability of language and computer applications*, ed. by Bruce E. Nevin & Stephen B. Johnson. Amsterdam & Philadelphia: John Benjamins, pp. 121-142.

Joshi, Aravind K., S. Rao Kosaraju, H. M. Yamada. 1972a. String Adjunct Grammars: I. Local and Distributed Adjunction. *Information and Control* 21.2:93-116. [DOI: 10.1016/S0019-9958(72)90051-4]

Joshi, Aravind K., S. Rao Kosaraju, H. M. Yamada. 1972b. String Adjunct Grammars: II. Equational Representation, Null Symbols, and Linguistic Relevance. *Information and Control* 21.3:235–260. [DOI: 10.1016/S0019-9958(72)80005-6]

Joshi, Aravind & Owen Rambow. 2003. A Formalism for Dependency Grammar Based on Tree Adjoining Grammar. *Proceedings of the [first] Conference on Meaning-Text Theory*, Paris, France.

Kandel, Eric R. 2006. *In search of memory: the emergence of a new science of mind*. New York: Norton.

Karuza, E.A., Newport, E.L., Aslin, R.N., Starling, S.J., Tivarus, M.E., & Bavelier, D. 2013. The neural correlates of statistical learning in a word segmentation task: An fMRI study. *Brain & Language* 127.1:46-54. [doi: 10.1016/j.bandl.2012.11.007]

Katseff, Shira E. (2010). *Linguistic constraints on compensation for altered auditory feedback*. Ph.D. Dissertation, U.C. Berkeley.

Katseff, Shira & John F. Houde, (2008). "Compensation ?=? Mental Representation". *LabPhon11 abstracts*, edited by Paul Warren, Wellington, NZ.

Katseff, Shira, John F. Houde, & Keith Johnson, (2008). "Partial compensation in speech adaptation". *2008 Annual Report of the UC Berkeley Phonology Lab*, 444–461.

Katseff, Shira, John F. Houde, & Keith Johnson, (2010). "Auditory feedback shifts in one formant cause multi-formant responses". *Journal of the Acoustical Society of America*, 127(3), 1955.

Kittredge, Richard & John Lehrberger (eds.). 1982. *Sublanguage: Studies of language in estricted semantic domains*. Berlin & New York: Walter de Gruyter.

Lieberman, Matthew D., Naomi I. Eisenberger, Molly J. Crockett, Sarina M. Tom, Jennifer H. Pfeifer, and Baldwin M. Way. 2007. "Putting feelings into words: Affect labeling disrupts amygdala activity in response to affective stimuli." *Psychological Science* 18.5:421-428. [PDF: http://www.scn.ucla.edu/pdf/AL(2007).pdf]

Labov, William. 1970. The study of language in its social context. *Studium Generale* 23:30-87.

Labov, William. 1972. *Sociolinguistic Patterns* Oxford: Blackwell.

Labov, William & J. Waletzky. 1967. Narrative analysis: oral versions of personal experience. In J. Helm (ed.) *Essays on the verbal and visual arts*. Seattle: University of Washington Press 12-44. (Repr. in *Journal of narrative and life history* 7:3-38.)

Loftus, Elizabeth F. 1998. Illusions of Memory. *Proceedings of the American Philosophical Society* 142:60–73.

Loftus, Elizabeth F. 1999. Lost in the mall: Misrepresentations and misunderstandings. *Ethics & Behavior* 9.1:51–60. [doi:10.1207/s15327019eb0901_4. PMID 11657488.]

Marken, Richard S. 1989. "Random-walk chemotaxis: Trial and error as a control process". *Behavioral neuroscience* 103:1348-1355. [Repr. in Marken (1995:87-105).]

Marken, Richard S. 1995. *Mind readings: experimental studies of purpose*. Gravel Switch, KY: CSG Press.

Marken, Richard S. 2002. *More mind readings: methods and models in the study of purpose*. Chapel Hill: New View Publications.

Marken, Richard S. 2005. Optical Trajectories and the Informational Basis of Fly Ball Catching. *Journal of Experimental Psychology: Human Perception and Performance* 31.3:630–634. [DOI: 10.1037/0096-1523.31.3.630]

Marken, Richard S. 2009. You say you had a revolution: Methodological foundations of closed-loop psychology. *Review of General Psychology*, 13.2:137-145. [http://dx.doi.org/10.1037/a0015106]

Marken, Richard S. 2014. *Doing Research on Purpose: A Control Theory Approach to Experimental Psychology*. Chapel Hill: New View Publications.

McClelland, Kent. 2004. "Collective Control of Perception: Constructing Order from Conflict." *International Journal of Human-Computer Studies* 60:65–99.

McClelland, Kent. 2011. "Environmental stability and collective control". Paper presented at the 2011 Conference of the Control Systems Group in Boulder, CO. [Kent McClelland <mcclel@grinnell.edu>, Department of Sociology, Grinnell College, Grinnell, IA 50112 USA.]

McClelland, Kent. [forthcoming]. A PCT perspective on sociology. In Warren Mansell (ed.), *Living control systems IV: the world according to Perceptual Control Theory*. New Canaan, CT: Benchmark Publications.

McGurk, Harry & John MacDonald. (1976). "Hearing lips and seeing voices." *Nature* 264(5588):746–8. [doi:10.1038/264746a0. PMID 1012311.]

Meguire, Philip. 2004. "Richard Milton Martin: American Logician". *The Review of Modern Logic* 10 (31): 7–65.

Miller, George A., Eugene Galanter, & Karl H. Pribram. 1960. *Plans and the Structure of Behavior*. New York: Holt, Reinhart & Winston.

Mirolli, Marco & Domenico Parisi. 2005. Language as an aid to categorization: A neural network model of early language acquisition. In A. Cangelosi, G. Bugmann, & R. Borisyuk (Eds.) *Modeling Language, Cognition and Action: Proceedings of the 9th Neural Computation and Psychology Workshop*. Singapore: World Scientific, 97-106. [http://goo.gl/IFlmXz]

Mirolli, Marco & Domenico Parisi. 2006. Talking to oneself as a selective pressure for the emergence of language. In Cangelosi, Angelo, Andrew D. M. Smith, & Kenny Smith (eds.) *The evolution of language: proceedings of the 6th international conference on the evolution of language*, Singapore: World Scientific, 214-221.

Mirolli, Marco, Francesco Mannella, & Gianluca Baldassarre. 2009. The Roles of the Amygdala in the Affective Regulation of Body, Brain and Behaviour. Connection Science 22.3:215-245. [http://goo.gl/QmR0mI]

Mirolli, Marco & Domenico Parisi. 2009. Language as a Cognitive Tool. *Minds & Machines* 19:517–528 [DOI 10.1007/s11023-009-9174-2]

Misyak, Jennifer B., Michael H. Goldstein and Morten H. Christiansen. 2012. *Statistical-Sequential Learning in Development*. (P. Rebuschat & J.N. Williams, Editors) Boston & Berlin: Walter de Gruyter. [DOI: http://cnl.psych.cornell.edu/pubs/2012-mgc-SL-LA.pdf]

Nevin, Bruce E. 2010. Noam and Zellig. In Douglas A. Kibbee (ed.) *Chomskyan (R)evolutions*. Amsterdam/Philadelphia: John Benjamins, 103-168.

Nevin, Bruce E. [forthcoming]. Language and thought as control of perception. In *Living Control Systems IV: the world according to PCT*. New Canaan, CT: Benchmark Press.

Nickerson, Raymond S. 1998. "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises". *Review of General Psychology* **2** (2): 175–220. doi:10.1037/1089-2680.2.2.175

Pavloski, R. P., G. T. Barron, & M. A. Hogue. 1990. Reorganization: Learning and Attention in a Hierarchy of Control Systems. American Behavioral Scientist 34.1:32-54. [doi:10.1177/0002764290034001005]

Plooij, Frans X. (1984). *The behavioral development of free-living chimpanzee babies and infants*. Norwood, N.J.: Ablex.

Plooij, Frans X. (2003). "The trilogy of mind". In Heimann (2003:185-205).

Plooij, Frans X. & Hetty van de Rijt-Plooij. (1990). Developmental transitions as successive reorganizations of a control hierarchy. *American Behavioral Scientist*, 34:67-80.

Post, Emil L. 1943. "Formal reductions of the general combinatorial decision problem". *American Journal of Mathematics* 65.2:197-215.

Powers, William T. 1973. *Behavior: The control of perception*. Chicago: Aldine de Gruyter. [2nd exp. ed. = Powers (2005)].

Powers, William T. 1979. The nature of robots. Part I: Defining behavior, *Byte* 4.6:132-134, 136, 138, 140-141, 144; Part II: Simulated control system, *Byte* 4.7:134-136, 138, 140, 142,144, 146, 148-150; Part III: A closer look at human behavior, *Byte* 4.8:94-96, 98, 100, 102-104, 106-108, 110-112, 114, 116; Part IV: Looking for controlled variables, *Byte* 4.9:96, 98-102, 104, 106-110, 112. Peterborough, NH: Byte Publications.

Powers, William T. 1988 [1989]. "An outline of control theory". *Conference workbook for "Texts in Cybernetic Theory"*. American Society for Cybernetics, Felton, California, October 18-23, 1988, 1-32. Reprinted in Powers (1989:253-293).

Powers, William T. 1989. *Living control systems: Selected papers of William T. Powers*. Gravel Switch, Kentucky: The Control Systems Group, Inc.

Powers, William T. 1992. *Living control systems II: Selected papers of William T. Powers*. Gravel Switch, Kentucky: The Control Systems Group, Inc.

Powers, William T. 1995. The origins of purpose: the first metasystem transitions. In F. Heylighen, C. Joslyn, & V. Turchin (eds.), World Futures vol. 45 (special issue on The Quantum of Evolution) 125-138.

Powers, William T. 2005. *Behavior: The control of perception*. New Canaan, CT: Benchmark Publications. [2nd exp. ed. of Powers (1973)].

Powers, William T. 2008. *Living control systems III: the fact of control*. New Canaan, CT: Benchmark Publications.

Powers, William T., R. K. Clark, & R. L. McFarland, 1960. "A general feedback theory of human behavior". *Perceptual and Motor Skills* 11, 71-88 (Part 1) and 309-323 (Part 2). [Reprinted in *General Systems V*, 63-83, 1960. Partial reprint in Smith, A. G., *Communication and Culture*, New York: Holt, Rinehart, and Winston (1966).]

Powers, William T. & Philip J. Runkel. 2011. *Dialogue Concerning the Two Chief Approaches to a Science of Life: Word Pictures and Correlations versus Working Models*. Hayward, CA: Living Control Systems Publishing.

Rosenbloom, Paul C. 1950. *The elements of mathematical logic*. New York: Dover.

Rudy, Jerry W. 2008. *The neurobiology of learning and memory*. Sunderland, Massachusetts: Sinauer Associates.

Rumelhart, David E., James L. McClelland, & PDP Research Group. 1986a. Parallel Distributed Processing, Volume 1, Explorations in the Microstructure of Cognition: Foundations. Cambridge, MA: MIT Press.

Rumelhart, David E., James L. McClelland, & PDP Research Group. 1986b. Parallel Distributed Processing, Volume 2, Explorations in the Microstructure of Cognition: Psychological and Biological Models. Cambridge, MA: MIT Press.

Runkel, Philip J. 2003. *People as living things: The psychology of perceptual control*. Hayward, CA: Living Control Systems Publishing.

Runkel, Philip J. 2007[1990]. *Casting nets and testing specimens: Two grand methods of psychology*. Hayward, CA: Living Control Systems Publishing. [1990 ed. New York: Praeger.]

Sala, Carlo & Menaham Segal. 2014. Dendritic Spines: The Locus of Structural and Functional Plasticity. *Physiological Reviews* 94.1:141-188 [DOI: 10.1152/physrev.00012.2013]

Saffran, Jenny R., Elissa L. Newport, Richard N. Aslin. 1996a. Word Segmentation: The Role of Distributional Cues. Journal of Memory and Language 35.4:606–621.

Saffran, Jenny R., Richard N. Aslin, & Elissa L. Newport. 1996b. Statistical Learning by 8-Month-Old Infants. *Science* 274.5294:1926-1928. [DOI: 10.1126/science.274.5294.1926]

Sampson, Geoffrey. 1980. *Schools of linguistics: Competition and evolution*. Stanford: Stanford University Press.

Selfridge, Oliver G. 1959. Pandemonium: A paradigm for learning. In D. V. Blake and A. M. Uttley (eds.), *Proceedings of the Symposium on Mechanisation of Thought Processes* (held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958, Vol. 1, Symposium number 10, Session 3, Paper 6), London: H.M. Stationery Office, pages 511–529. [Repr 1962, URL: http://webdocs.cs.ualberta.ca/~sutton/pandemonium.pdf]

Stix, Gary. 2014. The "It" Factor. Scientific American (September 2014), 311:72-79 [doi:10.1038/scientificamerican0914-72]

Taylor, Martin M. [forthcoming]. <Title>. In Warren Mansell (ed.), *Living control systems IV: the world according to Perceptual Control Theory*. New Canaan, CT: Benchmark Publications.

Tomasello, Michael. 1999. *Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.

Tomasello, Michael. 2014. *A Natural History of Human Thinking*. Cambridge, MA: Harvard University Press.

van de Rijt-Plooij, Hetty & Frans X. Plooij. 1992. "Infantile regressions: Disorganization and the onset of transition periods". *Journal of Reproductive and Infant Psychology* 10:129-149.

van de Rijt-Plooij, Hetty & Frans X. Plooij. 1993. "Distinct periods of mother-infant conflict in normal development: Sources of progress and germs of pathology". *Journal of Child Psychology and Psychiatry* 34:229-245.

van de Rijt, Hetty & Frans X. Plooij. 2010. *The Wonder Weeks. How to stimulate the most important developmental weeks in your baby's first 20 months and turn these 10 predictable, great fussy phases into magical leaps forward*. Arnhem, The Netherlands: Kiddy World Publishing

van Gelder, Tim. 1998. The dynamical hypothesis in cognitive science. *Behavioral & Brain Sciences*, 21, 615-665.

van Gelder, Tim & Port, Robert F. (1995). It's about time: An overview of the dynamical Approach to cognition. In R.F. Port and T. van Gelder (Eds.), *Mind as motion: Explorations in the Dynamics of Cognition*, 1-43. Cambridge, MA: MIT Press.

Wilson, Robert A. & Lucia Foglia. 2011. "Embodied Cognition", *The Stanford Encyclopedia of Philosophy* (Fall 2011 Edition), Edward N. Zalta (ed.) [URL = http://plato.stanford.edu/archives/fall2011/entries/embodied-cognition/].

Yin, Henry H. 2014. How Basal Ganglia Outputs Generate Behavior. *Advances in Neuroscience* 2014.768313. [doi:10.1155/2014/768313]