# UC Merced

## Proceedings of the Annual Meeting of the Cognitive Science Society

**Title**
Aspects of the Logical Structure of Conceptual Analysis

**Permalink**
https://escholarship.org/uc/item/4dt5k8t6

**Journal**
Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

**ISSN**
1069-7977

**Author**
Niyogi, Sourabh

**Publication Date**
2005

Peer reviewed

# Aspects of the Logical Structure of Conceptual Analysis

**Sourabh Niyogi**
MIT
Cambridge, MA 02139
niyogi@mit.edu

## Abstract

Modern day lexical databases are not constructive, but differential (Miller et al 1990). We outline the logical structure of the task of a conceptual analyst who wishes to construct a constructive lexicon based on a Universal Theory Model of Concepts. Elementary notions of descriptive and explanatory adequacy are developed within this model. The diverse streams of evidence available for the conceptual analyst to engage in empirical inquiry are reviewed in the domains of light and perception.

## Theories for Constructive Lexicons

Major projects have been conducted for centuries to record, in one form or another, representations of word meanings. This has been in the form of the lexicographer's dictionary or the more modern computational linguist's electronic databases (WordNet, FrameNet, VerbNet) e.g. last year's symposium (Miller, Fillmore, Palmer, Lenat and Hayes 2004). The degree to which computational linguists and cognitive scientists draw upon these resources as accurate descriptions of lexical knowledge is astonishing. Researchers speak of "putting meaning in your trees", providing "deep semantics", automatically labeling "semantic roles", or using WordNet as the authority for word sense disambiguation, to name some key projects. (Palmer et al 2002, Fillmore et al 2001, Gildea and Jurafsky 2002) Given the increasing reliance to these databases, it may appear that the terms *meaning* and *semantic* are not being used glibly and that the theoretical foundations of these databases were sound.

Even the most cursory analysis shows that this is not so. Consider the distinction made by the originators of WordNet between a constructive vs. differential lexicon: In a *differential* theory of the lexicon, meanings can be represented by any symbols that enable a theorist to distinguish among them; In a *constructive theory* of the lexicon, the representation should "contain sufficient information to support an accurate construction of the concept (by either a person or a machine)" (Miller et al 1990). Today's dictionaries and all of today's lexical databases are *differential*: the intension of synsets of WordNet are just sufficient so that someone who already knows English can distinguish among synsets, while the thematic roles used in VerbNet and FrameNet are notoriously difficult to define: primitive terms such as Agent,

Patient, Perceiver_passive, Phenomenon, etc. are unexplained primitives with few constraints on their generation (Fodor 1998, Newmeyer 1998). If the semantics problem is kicked upstairs, it is unclear what, exactly, is being accomplished.

In this paper, we consider the task of the *conceptual analyst* who wishes to construct a *theory-based lexicon* that is constructive rather than differential. By a theory-based lexicon, we mean a lexicon based on a technical explication of the traditionally nebulous "theory theory" (Gopnik and Meltzoff 1997, Carey 1991). Recent developments suggest that a technical explication of what is meant by "theory theory" is tractable (c.f. Gopnik et al 2003, Niyogi 2005). The simplest model of this is diagrammed in Figure 1, which I call the *Universal Theory Model of Concepts*. Via a set of primitives and experience, a Theory Acquisition Device (TAD) acquires some theory $T^*$, an explicit representation that attempts to explain the law-like regularities of domains and entities hypothetical and of the world. The TAD yields person P's "naive theory"; If P's TAD outputs $T^*$, then $T^*$ is P's theory, and generates the possible internal states his mind may have of a set of observable and unobservable variables in various domains. A concept generator $G$, applied to $T^*$, generates a hypothesis space of word meanings $G(T^*)$. Then, a Vocabulary Acquisition Device uses experience and $G(T^*)$ to output lexicon $L$.

For example, recall the well-known observaton that children, when probed indirectly for the meaning of *uncle* at one stage (T1), thinks it means "friendly middle aged men", and when probed at a later stage (T2), thinks it refers to a "parent's sibling" (Keil 1989). The standard explanation is the child at T1 does not have the kinship framework that the child at T2 has. In this UT model, T1 and T2 are two possible states $T^*$ of the TAD (all else held constant). In $L$, the word *uncle* is mapped to an element of the set $G(T1)$ when $T^* = T1$, and then mapped to an element of the set $G(T2)$ when $T^* = T2$.

This *theory acquisition device* (TAD) has an initial state $T^*(t = 0)$, and a set of possible states it may assume. The theory of the initial state of TAD, we will call Universal Theory (UT), analogous to Universal Grammar (Chomsky 1981). UT constrains the set of possible states that the TAD can be in. The TAD state may evolve to improve on the generative model, given either observations in the domain or by mental fiat.
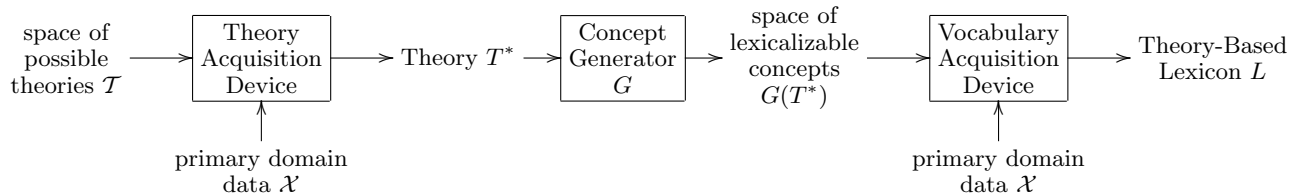
Figure 1: The Universal Theory Model of Concepts. A Theory Acquisition Device selects $T^*$ from a set of possible theories $\mathcal{T}$. A Concept Generator $G$ maps $T^*$ to a set of lexicalizable concepts $G(T^*)$. A Vocabulary Acquisition Device outputs theory-based lexicon $L$ using $G(T^*)$ as a hypothesis space; The task of conceptual analysis is to provide an account of $\mathcal{T}$ and $G$ that achieves explanatory adequacy and an account of $T^*$ and $L$ that achieves descriptive adequacy.

## The Task of Conceptual Analysis

In the Universal Theory model of concepts the analyst is faced not with just the task of describing lexicalized concepts in some form or another (as would be sufficient for a *differential* theory of the lexicon), but is faced with several *interrelated* tasks:

**Task (1)** describing a space of possible theories $\mathcal{T}$

**Task (2)** describing a *particular* TAD state $T^*$ for a particular person P at a particular moment in time $t$; this depends on Task 1.

**Task (3)** describing the concept generator $G$, that maps a *particular* theory state $T^*$ to a set of concepts $G(T^*)$; this depends on Task 2.

**Task (4)** describing a *particular* Lexicon $L$ as a map from a set of sounds to a set of elements of $G(T^*)$; this depends on Task 3.

Tasks (1) and (3) concern the study of universal properties of all persons P, independent of experience; Tasks (2) and (4) concern descriptions of particular states of P at a particular moment of time.

We may adapt, more or less directly, the notions of descriptive adequacy and explanatory adequacy developed in generative grammar by Chomsky (1965) to the above. The task of the conceptual analyst is to construct descriptively adequate accounts of $T^*$ and $L$ and explanatorily adequate accounts $\mathcal{T}$ and $G$, given evidence available to the conceptual analyst.

### Accounts of the TAD

An account of the TAD with *descriptive adequacy* must describe the precise internal structure of a particular person P's TAD state $T^*$ at a particular point in time. It may be reasonable for a conceptual analyst to consider $P$ and a somewhat different $Q$ as representative of two populations – say of 3-year olds vs 10 year olds, or of English-speaking vs Chinese-speaking adults – and give explicit descriptions of T1 and T2 for each population. It may also be reasonable for a conceptual analyst to consider relatively modular fractions of the TAD state to make the task of describing $T^*$ somewhat more managable – e.g just those concerned with "perception", "disease" or "kinship". If the internal state of a person's P TAD internalized some notions not in the analyst's description,

then we would say the analyst's account would not be descriptively adequate.

In contrast, an account of the TAD with *explanatory adequacy* must demonstrate how the set of Ts that can be hypothesized and evaluated by a TAD is fully determinable by primary domain data $\mathcal{X}$ or sheer mental fiat. By *primary domain data $\mathcal{X}$*, I mean observations available in a particular domain; it is this data that a TAD state $T^*$ may hope to explain, in that $T^*$ is an adequate generative model for the observations. By sheer mental fiat, I mean every cause of TAD state change *not* driven by primary domain data: abduction, analogical matching and so on. The former process may be given a Bayesian/MDL treatment (Tenenbaum, Griffiths and Niyogi, in press); the latter processes remains mysterious, although a very important topic in cognitive science.

A theory of the TAD that seeks explanatory adequacy must explain how $\mathcal{X}$ is mapped to a particular reachable T (or a set of reachable Ts) in $\mathcal{T}$. Theory acquisition would be impossible unless there were a circumscribed set of Ts – circumscription is necessary fact of *any* induction problem, as is well-known (Goodman 1951). Implicitly, the analyst must answer: *How are all possible theories alike?* The central questions are: what is the initial state of the TAD, and given the primary domain data, how does it evolve to future states?

Adapting Chomsky (1965, pg. 30), a TAD must have the following:

**(i)** a technique for representing observations in the domain $(w)$

**(ii)** a technique for representing the unobservable states in the domain input structurally $(SD)$

**(iii)** some delimitations on the class of possible hypotheses $\mathcal{T}$ concerning the structure of the input (Universal Theory, or UT)

**(iv)** a method for determining what each hypothesis $T_i \in \mathcal{T}$ implies with respect to the primary domain data $\mathcal{X}$

**(v)** a method for evaluating the hypotheses in $\mathcal{T}$ with respect to the primary domain data $\mathcal{X}$

This description allows for theory change based solely on the primary domain data $\mathcal{X}$; The simplest model for theory change by sheer mental fiat is one where the TAD is always constrained to be in $\mathcal{T}$.

Necessarily, an abstract model of the TAD that aims for explanatory adequacy must contain the following:

**(i)** an enumeration of the class $w_1, w_2, \ldots$ of possible observations

**(ii)** an enumeration of the class $SD_1, SD_2, \ldots$ of possible structural descriptions

**(iii)** an enumeration of the class of $T_1, T_2, \ldots$ of possible theories

**(iv)** specification of a function PARSE such that it may assign a structural description $SD_{f(i,j)}$ to observation $w_i$ by $T_j$ for arbitrary $i, j$

**(v)** specification of a function ACQUIRE such that given primary domain data $\mathcal{X}$ and PARSE chooses $T^*$ from the space of possible theories $\mathcal{T}$

This approach may be taken with any inductive inference problem in any domain in the cognitive sciences; such an approach has been taken in inductive inference problems in grammar and vision. We merely adapt the approach here to intuitive theories.

The general situation the conceptual analyst is faced with is the same as in any induction problem: Given limited observations, the induction problem is to produce a model of the observations that summarizes the observations well. There are two kinds of failure a model can have, shown here:
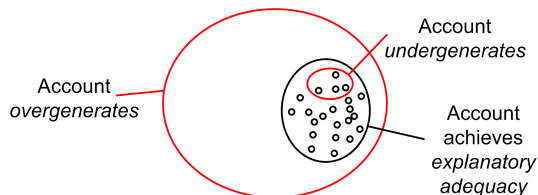


One kind of failure is to have a model that *overgenerates*, and predicts large classes of observations that are not observed. The other kind of failure is to have a model that *undergenerates*, and fails to predict large classes of actually attested observations.

This general situation applies to the induction of a metalanguage of theories ($\mathcal{T}$). The discovery by the conceptual analyst that there are elements of TAD state common to many domains would suggest that there are excellent candidates for a generative model of $\mathcal{T}$. In Niyogi (2005), I sketch a minimal UT model for a inter-related system of elementary spaces, kinds, attributes, relations, and causal laws. This account cannot be considered complete until each component is circumscribed: what is a possible space? kind? attribute? relation? causal law? Very different proposals can be found in Gopnik et al (2003) and Tenenbaum, Griffiths and Niyogi (in press) where probablistic models of causation and abstraction are emphasized. A very different account is proposed by Rogers and McClelland (2004) who modeling theory change in a PDP setting. These early efforts each provide a glimpse into the form of Universal Theory.

Any one proposal for T can be seen as overgenerating or undergenerating the space of possible theories. A candidate UT unable to represent probabilistic causal laws

may be considered as undergenerating. Alternately, a candidate UT that does not constrain the set of spaces and natural paths in those spaces would be overgenerating. Some aspects of UT may require outright stipulation, such as datatypes involving "propositions" or "events", as many in lexical semantics would suggest.

## Accounts of the VAD

An account of $L$ with *descriptive adequacy* must describe the precise internal structure of a particular person P's lexicon $L$ at a particular point in time $t$. Lexical databases may appear to do this, but their accounts of $L$ are differential, not constructive. In the Universal Theory Model of Concepts, a descriptively adequate account of $L$ must map sounds to elements of $G(T^*)$ for a particular person P. Clearly, we cannot have a descriptively adequate of the account of the VAD without also having a descriptively adequate of Person P's TAD state $T^*$, and a model of $G$. Again, it is reasonable for the conceptual analyst to assume homogeneity and consider that $P$ is representative speaker, e.g. of English or Chinese.

In contrast, an account of the VAD with *explanatory adequacy* must demonstrate how the set of Ls that can be hypothesized and evaluated by a VAD is fully determinable by its possible inputs. Note that there is two viewpoints on "possible inputs":

- Viewpoint 1: the set of concepts $G(T_i)$ across all possible $T_i \in \mathcal{T}$

- Viewpoint 2: the set of concepts $G(T^*)$ for a fixed $T^*$

The first viewpoint is appropriate for the species, while the second is appropriate for the individual P at a specific moment in time – c.f. Niyogi (2005) for related discussion.

Likewise, the metalanguage of concepts ($G$) is also faced with a similar induction problem. Mismatches between the actually observed concepts and the concepts generated by the analysts guess at $T^*$ may be caused by several factors, requiring appropriate measures by the analyst.

If a word is observed whose meaning is not adequately modeled by any element of $G(T^*)$, then it may be due to a deficiency of $T^*$ (e.g. missing particular kind, attribute, relation, law) or a general inadequacy of the concept generator $G$.

If $G(T^*)$ contains concepts that are not attested in any word of a particular language (that is, $G$ is overgenerating), this should not be cause for alarm; it is widely known from machine translation that conceptual gaps are systematic, and in general it will be the case the $G(T^*)$ is larger then the number of roots in $L$. Of concern would be the situation where $G(T^*)$ overgenerates across a large sample of languages ("large" being a notoriously subjective decision). An element of $G(T^*)$ unattested in any $L$ may not be observed simply because it is *useless* and has no useful communicative value. If such overgeneration is systematic, however, the conceptual analyst may be required to revise $G$ to not overgenerate these unattested elements.

# Case Study in Light and Visual Perception

The conceptual analyst, like the syntactician who proposes to study Universal Grammar, has different streams of evidence available to him:

1. naive beliefs of homogenous populations: scientists, children at different stages of development, alternative conceptions of students and adults

2. subjective causal knowledge of the analyst

3. observed word meanings in a variety of languages

4. syntactic selectional requirements of lexical items

5. entailment observations and commonsense inferences

6. metaphoric extension and historical change

The first two streams provide evidence for the content of $T^*$; the latter 4 streams provide evidence for the content of $L$. Traditionally, these observations are tabulated for very different purposes by psychologists, lexicographers, and linguists. The conceptual analyst must infer $T^*$ and $L$ based on these streams of observations.

Consider the streams available in the domain of visual perception and light. The analyst may select naturally occuring terms in a variety of languages in this "domain", and choose prototypical and naturally occurring instances of these terms (see Figure 2). Differential lexicons themselves are available to facilitate this process (FrameNet, WordNet, VerbNet), at least for English. There is little reason to expect sharp delimitations of a "domain"; there are body parts and actions (eye, blink), artifacts (telescope), social phenomena (show, hide) that are clearly related but may be considered to be in another arbitrarily-defined domain. Thus many interelated domains may undergo conceptual analysis. The system of kinds, attributes, relations, laws, and parts embedded in $T^*$ must be able to generate concepts $G(T^*)$ so as to be mappable to each of the terms in these domains.

## Naive Beliefs of Children, Adults, Ancient and Modern Science

The analyst may consider the naive beliefs of subjects representative of a larger population of subjects. In a large number of domains, the naive beliefs of students have been tabulated. Across many domains, the beliefs of some populations are similar to those of ancient philosophers and medieval scientists; these beliefs can be differentiated from what modern science can tell us. In the domain of perception, for example, many children and adults, like ancient philosophers such as Plato and Euclid, hold an "extramission" belief, that we see by shooting rays outward to sample an object (Winer et al 2002). Many misconceptions are held of light: that it flows and mixes as if it were a paint-like liquid; that it can be at rest, such as on the surface that it hits; that color and shadows exist only *in* an object. These beliefs provide constraints on $T^*$.

Very different beliefs can be tabulated by comparing multiple populations. Blind children can acquire a concept of *look* and *see* that is not visual per se (Landau and

| English Root | WordNet gloss |
|---|---|
| eye | the organ of sight |
| look | the act of directing the eyes toward something and perceiving it visually |
| see | perceive by sight or have the power to perceive by sight |
| ignore | fail to notice |
| notice | the act of noticing or paying attention |
| perceive | to become aware of through the senses |
| view | the range of the eye |
| search | try to locate or discover, or try to establish the existence of |
| detect | discover or determine the existence, presence, or fact of |
| aware | having or showing realization or perception; alert and fully informed |
| show | make clear and visible, or noticeable |
| reveal | make visible |
| conceal | prevent from being seen or discovered, make imperceptible |
| cover | the act of concealing the existence of something by obstructing the view of it |
| occlude | block passage through |
| appear | come into sight or view |
| disappear | become invisible or unnoticeable |
| clear | transmitting light; able to be seen through with clarity |
| opaque | not clear; not transmitting or reflecting light or radiant energy) |
| visible | perceptible especially by the eye; or open to easy view |
| invisible | impossible or nearly impossible to see; imperceptible by the eye |
| hidden | not accessible to view; covered from view; difficult to find |
| illuminate | make lighter or brighter |
| dark | devoid of or deficient in light or brightness |
| telescope | a magnifier of images of distant objects |
| eyeglasses | optical instrument consisting of a pair of lenses for correcting defective vision |
| blink | a reflex that closes and opens the eyes rapidly |

Figure 2: A collection of related roots in the domain of light and visual perception, from WordNet.

Gleitman 1985); given *look up!*, the blind child will raise his hands to grope for something; given *see the X*, on the other hand, the blind child will haptically detail the X given to him. The blind child's concept of *look* and *see* and the normal seeing child have something in common: that *look* involves an agent's making it *possible* to perceive an entity, whereas *see* involves necessarily detailed modeling of it. The concepts underlying these "perception" terms make reference to somewhat different $T^*$.

## Subjective Analysis

The analyst may combine the above observations with his own subjective beliefs to sketch a set of kinds, attributes, relations, part-whole structures, and causal mechanisms that interrelate these sets. Unlike alternative conceptions, commonsense facts such as physical objects having color, an eye having an attribute of gaze orientation, degree to which an eyelid is open, etc. is rarely the subject of extensive documentation. Figure 3 shows a diagram identifying some of these key mechanisms, in which many of the roots in Figure 2 may hope to be grounded. Some key causal mechanisms diagrammed are:

- the process of seeing is initiated by shooting a ray out from the eyes

- the presence of an entity (1) with attributes (color, reflectance, transmittance) enables a contact point from the ray

- occluders (2) prevent the ray from reaching the entity; occluders themselves are just like the entity.

(1) Entity
Attributes: color, reflectance, transmittance, ...

(2) Occluding Entity
Attributes: color, reflectance, ...

(3) Illumination source
Attributes

(4) Eye
Attributes: orientation, ...

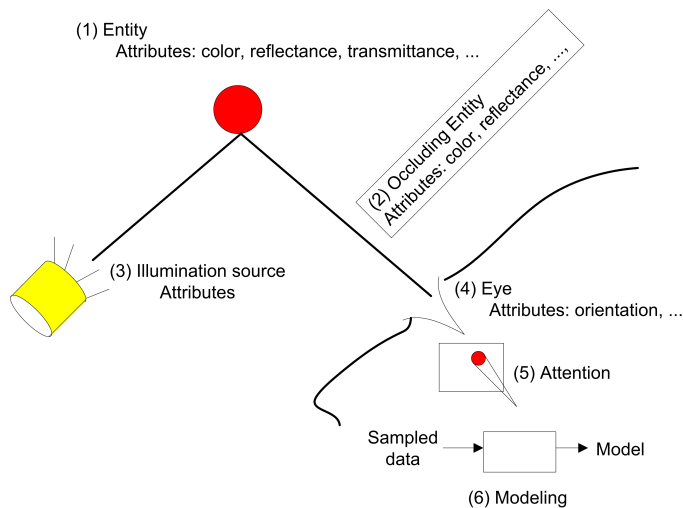(5) Attention

Sampled data → Model

(6) Modeling

Figure 3: Subjective conceptual analysis of the domain of light and visual perception.

- illumination (3) is necessary for sampling to occur
- the eye (4) is controllable by the perceiver through motor actions (4') that control the key attribute of gaze orientation
- attention (5) is required for modeling to occur, selecting a part of what is sampled
- modeling (6) maps what is sampled to a cognitive model

Again, none of these mechanisms have anything to do with what actually is happening. Moreover, these interconnected causal mechanisms do *not* require further reduction; instead, the description merely be shallow and phenomenological (diSessa 1993, Rozenblit and Keil 2003). Each of these causal processes must be described in the theory metalanguage $\mathcal{T}$; the part in $T^*$ concerned with the domain of light and visual perception must be sufficiently complete to generate concepts in $G(T^*)$ that map to naturally occurring terms in the domain.

The analyst may relate the parts of $T^*$ concerning one domain (of visual perception) to parts of related domains, such as that of the mind (e.g. further post-sensation processes), body (how the *eye* works), or artifacts (how *glasses, microscopes, light bulbs, mirrors* work). This web of interconnected domains in $T^*$ does not imply a lack of modularity or some bizarre form of holism. Instead, the interconnectivity is systematic: detailed conceptual analysis may reveal instead that speaking of theories of perception, artifacts, body, mind, etc. are largely gross (yet convenient) pointers to parts of $T^*$.

## Syntactic Phenomena

The analyst may cull surface syntactic phenomena for each of the above lexical items, e.g.

He hid/*showed/*revealed the toy from the baby.
He *hid/showed/revealed the toy to the baby.
He *hid/showed/*revealed the baby the toy.
He revealed the toy *for hours/*in an hour/at 6:00.

The similarity and differences in these terms may be directly encoded through the precise way $G$ maps the sets in $T^*$ concerning mechanisms of visual perception to $G(T^*)$. However, many other factors are known to be at work in the above. The above phenomena may be due to the formal compositional properties of the elements in $G(T^*)$, about plausibility judgements (the ungrammatical sentences possess the same status as *colorless green ideas*), or properties of constructional learning (Goldberg, in press).

## Entailment

The analyst may query native speakers for entailment judgements on the family of related terms, such as:

(1a) John saw the rock. →
(1b)    John looked at the rock.
(2a) John hid the rock from Sally. →
(2b)    John prevented Sally from seeing the rock.
(3a) John showed Sally the rock. →
(3b)    John caused Sally to see the rock

In the above cases, the (a) sentence entails the (b) sentence, but not vice versa, providing constraints on the mapping between $G(T^*)$ and the concepts behind the entailment. The analyst may readily ask native speakers for entailment judgements (whether *looking* entails *seeing*, whether *hiding* entails *occlusion*, whether *showing* entails *seeing*, and so on). In many cases, the entailment judgements are not obvious, and well-known cases (such as "kill" meaning "cause die") have been difficult for analysis by linguists, let alone naive subjects. Naturally, the inference engine required to test whether a model of $L$ could account for these entailment inferences would hinge on the fine structure of $T^*$.

## Metaphoric extension

Finally, in many interesting cases, special syntactic phenomena arise that require more involved explanation: (Sweetser 1990)

| | |
|---|---|
| *I saw the red ball.* | (direct perception) |
| *I saw the boy cry.* | (direct perception) |
| *I saw that the boy cried.* | (indirect inference) |
| *We saw her as guilty* | (mental) |

In this case, that perception verbs can take complementizer arguments may be directly related to how the parts of $T^*$ model the relation between visual seeing and subsequent inferential processes (e.g. *She conceals her anger well*, *They ignored all the warning signs*, *The plan was clear to us.*) Sweetser (1990, ch. 2) observes that verbs of perception systematically undergo the same paths of historical change.

The domain of perception may be related to other domains through other sorts of metaphoric mapping. In *I shot a glance at the clock*, there is a mapping between the causal mechanisms in $T^*$ underlying visual perception and those of a gun. Likewise, in *I felt his glance*, there is a causal mechanism that maps visual contact to a kind of physical contact. Tight interrelations between the internal content of $G(T^*)$ and processes underlying conceptual metaphor are likely to exist.

# Towards Constructive Conceptual Analysis

The conceptual analyst can harness these diverse streams of evidence to induce $T^*$ and $L$, aiming not for a differential lexicon but a constructive lexicon. This challenge, as Miller et al (1990) says, is not easily met, but the conceptual analyst who seeks to build a constructive lexicon would appear to require nothing less than a generative model of theories and a generative model of concepts. This paper described aspects of the logical structure of conceptual analysis, in the context of the Universal Theory Model of Concepts, where $\mathcal{T}$ and $G$ represent the required generative models.

Outside of lexicography, conceptual analysis is a very active tradition in cognitive linguistics and lexical semantics (Lakoff and Johnson 1999, Jackendoff 1990, Rappaport-Hovav and Levin 1998), although with very different models of concepts, and not with the intent of building constructive lexicons. Representative work in cognitive linguistics have undertaken the task of describe the most elementary physics verbs such as *grasp* (Bailey et al 1997) and how it may be used in different domains. Barsalou et al (in press) outlines the causal structure underlying intuitive theories of artifacts, while Lawvere (1993) and Atran (1995) describe the fine structure of kinship. Gopnik et al (2003) develop computational accounts of a child's concept of a *blicket* as an entity that has certain causal powers. These diverse works are windows into the candidate models of $\mathcal{T}, T^*, G$ and $L$, can be seen as work toward accounts of the TAD and VAD with descriptive and explanatory adequacy.

Historically, mainstream generative linguistics has assumed a position where the study of general systems of knowledge and belief is considered separate (by definition) from the study of syntax. This separation need not be an uncrossable chasm. Instead, there is a *bridge* between the two systems: the concept generator $G$, a map a theory $T^*$ ("knowledge and belief") to a set of lexicalizable concepts $(G(T^*))$. We may inherit the vocabulary of early generative linguistics and use it the task of conceptual analysis, while casting any syntacto-centrism aside.

The project to assemble these models in a constructive lexicon is something that has never been undertaken. Numerous web-based applications make it far easier to manage the interrelated systems embedded in $T^*$ and collaborate on them in a decentralized way. It is natural to use the differential lexicons such as WordNet, FrameNet, and VerbNet as a rich starting point. One may optimistically hope that there exist migration paths from these differential lexicons to a constructive lexicon. Serious innovations in our models of concepts will likely be needed to traverse this path and integrate the many streams of evidence that make this task tractable. Unlike the ad hoc construction of differential lexicons, there is a logical structure to conceptual analysis. We may enlarge our experimental apparatus and assume the role of the conceptual analyst.

## References

[1] S. Atran. Classifying nature across cultures. In E. Smith and D. Osherson, editors, *Thinking: An invitation to cognitive science*, Cambridge, MA, 1995. MIT Press.

[2] D. Bailey, J. Feldman, S. Narayanan, and G. Lakoff. Modeling embodied lexical development. In *Proceedings of the Annual Cognitive Science Society*, 1997.

[3] L. Barsalou, S. Sloman, and S. Chaigneau. The HIPE theory of function. In L. Carlson and E. van der Zee (Eds.), *Representing functional features for language and space*, New York, 2005, to appear. Oxford University Press.

[4] S. Carey. Knowledge acquisition: Enrichment or conceptual change? In *The Epigenesis of Mind: Essays on Biology and Cognition*, 1991.

[5] N. Chomsky. *Aspects of the Theory of Syntax*. MIT Press, Cambridge, MA, 1965.

[6] N. Chomsky. *Lectures on Government and Binding*. Dordrecht, Foris, 1981.

[7] A. diSessa. Toward an epistemology of physics. *Cognition and Instruction*, 10:105–225, 1993.

[8] C. Fillmore, C. Wooters, and C. Baker. Building a large lexical databank which provides deep semantics. In *Proceedings of the Pacific Asian Conference on Language, Information and Computation*, Hong Kong, 2001.

[9] J. Fodor. *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press, New York, 1998.

[10] D. Gildea and D. Jurafsky. Automatic labeling of semantic roles. *Psychological Review*, 28:245–288, 2002.

[11] A. Goldberg. *Constructions at Work: the nature of generalization in language*. Oxford University Press, Oxford, to appear.

[12] N. Goodman. *The Structure of Appearance*. Harvard University Press, Cambridge, 1951.

[13] A. Gopnik, C. Glymour, D. Sobel, L. Schultz, and T. Kushnir. Theory formation and causal learning in children: Causal maps and bayes nets. *Psychological Review*, in press.

[14] A. Gopnik and A. Meltzoff. *Words, thoughts and theories*. MIT Press, Cambridge, MA, 1997.

[15] R. S. Jackendoff. *Semantic Structures*. MIT Press, Cambridge, MA, 1990.

[16] F. Keil. *Concepts, kinds, and cognitive development*. MIT Press, Cambridge, MA, 1989.

[17] G. Lakoff and M. Johnson. *Philosophy In The Flesh*. Basic Books, 1999.

[18] B. Landau and L. R. Gleitman. *Language and experience: Evidence from the blind child*. Harvard University Press, Cambridge, MA, 1985.

[19] W. Lawvere. Kinship and mathematical categories. In R. Jackendoff, P. Bloom, and K. Wynn, editors, *Language, Logic, and Conceptual Representation*, Cambridge, MA, 1993. MIT Press.

[20] G. Miller, R. Beckwith, C. Fellbaum, D. Gross, K. Miller. Introduction to WordNet: An On-line Lexical Database. *International Journal of Lexicography*, 3(4):235–244, 1990.

[21] G. Miller, M. Palmer, C. Fillmore, D. Lenat, and P. Hayes. Large-scale knowledge representation resources for cognitive science research. In *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*, 2004.

[22] F. Newmeyer. *Language Form and Language Function*. MIT Press, Cambridge, MA, 1998.

[23] S. Niyogi. The universal theory model of concepts and the dissolution of the puzzle of concept acquisition. *Proceedings of the 27th Annual Meeting of the Cognitive Science Society*, 2005.

[24] M. Rappaport Hovav and B. Levin. Building verb meanings. In M. Butt and W. Geuder, editors, *The Projection of Arguments: Lexical and Compositional Factors*, pages 97–134, Stanford, CA, 1998. CSLI Publications.

[25] T. Rogers and J. McClelland. *Semantic Cognition: A parallel distributed Processing approach*. MIT Press, Cambridge, MA, 2004.

[26] L. Rozenblit and F. Keil. The misunderstood limits of folk science: an illusion of explanatory depth. *Cognitive Science*, 26:521–562, 2002.

[27] E. Sweetser. *From Etymology to Pragmatics*. Cambridge University Press, New York, 1990.

[28] J. Tenenbaum, T. Griffiths, and S. Niyogi. Intuitive theories and rational causal inference. In A. Gopnik and L. Schulz, editors, *Causal Learning: Psychology, Philosophy, and Computation*, Oxford University Press, Oxford, to appear.

[29] J. Tenenbaum and S. Niyogi. Learning causal laws. In *Proceedings of the 25th Annual Meeting of the Cognitive Science Society*, 2003, submitted.

[30] G. Winer, J Cottrell, V. Gregg, J Fournier, and L. Bica. Fundamentally misunderstanding visual perception. *American Psychologist*, 57:417–424, 2002.