Beyond "Beanbag" Models of Cultural Transmission:
A Proposal for Exploring the Effect of Semantic Relationships on Trait Distributions

Mark E. Madsen*
Department of Anthropology, Box 353100, University of Washington, Seattle WA, 98195 USA

Carl P. Lipo
Department of Anthropology and IIRMES, 1250 Bellflower Blvd,
California State University at Long Beach, Long Beach CA, 90840 USA†

(Dated: April 19, 2007)

Keywords: cultural transmission, semantic models, description logics

This note originated in 2003 as a set of rough notes as Madsen learned description logics, RDF, and semantic web technologies for a (now defunct) startup company. He wrote it up in 2007 as a formal proposal, which was sidelined until 2013, when it became a focus of work for Madsen and Lipo as a collaboration.

In detail, this note has been somewhat deprecated by years of developments in semantic information systems and knowledge representation, but it remains an artifact of the our first explorations of the connection between CT models and semantic models. Such models continue to be essential today, to allow us to understand how various social learning processes are conditioned not just by space, population structure, or cognitive bias, but by the actual content of the information being transmitted. Discussion of CT models that have semantic content between traits have remained verbal and heuristic, and formal models of transmission continue to locate all of the deep content in CT processes within the copying rules or cognitive biases which shuffle "tokens" within a frequency space.

Real evolutionary processes are not token shufflers; the things whose frequency are changed by population processes have real effects and meaning. In modern terms, this proposal is our first entry into an "evolutionary developmental biology" or "evo-devo" of cultural transmission, which we argue is badly needed. Its first expression is our joint work on the "Semantic Axelrod Model."

I. INTRODUCTION

Contemporary research on cultural transmission is split into several approaches, which today seem to communicate little and occasionally appear competitive. One approach, led by Robert Boyd and Peter Richerson, builds and analyzes mathematical models of simple "modes" of transmission: unbiased or random copying, as well as various types of bias in the choice of "models" to imitate. Such models follow the population genetic convention of defining the frequency distribution of traits within a population and examining the effect on this frequency distribution as models of evolutionary processes or population structure are added. The second approach is less cohesive, but is concerned with theoretical description of the relationship between cognition, language, and cultural behavior. Dan Sperber is perhaps the best example of this group, although certain of Liane Gabora's recent papers probably fit this category. These approaches appear to communicate very little, which is perhaps understandable given the differences in approach and the research questions to which each approach is applied.

But it would be wrong to see these as *competing* approaches, although these differences in approach are occasionally portrayed as alternatives (see essays in Aunger XXXX). The value of population-level models is clear: once we get beyond questions of the representation, modification, and expression of cultural information within the bounds of a single phenotype, we inevitably face the need to understand how the "inner" processes affect the "outer" distribution of traits and their evolutionary trajectories. Equally clear, however, is the value of understanding how cultural information is represented, modified, filtered, and expressed into behavior. Essentially, each approach addresses "half" of what we might consider the full range of cultural transmission phenomena. By analogy to the biological realm, it is not sufficient to have *either* molecular biology or population genetics: both are required to allow scientists to address the full range of phenomena we observe.[1] Although there is much work remaining within the normal "comfort zone" of each approach, we be-

---

* Contact Information: `madsenm@uw.edu`; Website: `http://notebook.madsenlab.org`;

† Contact Information: `Carl.Lipo@csulb.edu`; Website: `http://lipolab.org`

[1] Naturally, this is a rough analogy. Actually, addressing the full range of biological phenomena will require not just molecular genetics and population genetics, but a host of other approaches such as proteomics, ethology and behavioral ecology, etc. A similar structure applies to human cultural

lieve in bridging the gap between approaches, and seeing what disparate models are capable of when combined.

What follows is a proposal for exploring a combined approach, which inserts a simple model of the semantic relationships between discrete traits into existing population-level models of cultural transmission. Our approach is to represent the semantics of cultural traits using currently available technologies for ontology engineering and logical inference, and to employ well understood models of trait transmission to examine the population-level effects of "structured" cultural traits. This paper is a proposal for how to proceed, rather than a report of already completed research.

II. SEMANTICS AND CULTURAL TRAITS

In nearly all population-level models of cultural transmission, traits are "featureless" tokens which bear no formal relationship to each other. Even in situations where continuous or quantitative traits are modeled as real-valued variables, the values taken by these variables carry no further information beyond their magnitude and difference. The question thus arises, how can we expand our formal models of cultural transmission, to encompass the idea that traits are not simply tokens, but also have "meaning," in order that we may study the ways in which the semantics of traits can generate patterning in the transmission of such traits—patterns that are in addition to the qualitative and quantitative patterns we see when considering biased and unbiased forms of imitation and social learning for simple, "token"-style cultural traits.

At a high level, we can appreciate two ways of "attaching" semantics to variables which represent cultural traits. One source of semantic content is, of course, the "definitions" or "denotata" of a concept itself. The types of concepts we call "nouns" in grammar are of this type, and gain their semantics because of a denotation relationship between the concept and some aspect of the sensory world. In a less intuitive way, the types of concepts we call "verbs" are also denotational, but of time-transgressive "processes" rather than "objects." Again, this type of high-level description glosses over (indeed, buries) decades of research in cognitive science and linguistics about how such semantics arise and are learned during cognitive development.

Sidestepping this history for a moment, let us simply accept that some concepts have denotative semantics by virtue of "standing for" some part of the sensory world, or a process by which that sensory world changes over time. A second source of semantic content is relational, deriving from naming the relationships between two or more concepts. Immediately, we can see "synonym" and its inverse, "antonym," as relational semantics. But so are relationships which indicate subclass/superclass relationships, and those which create new classes via subset restriction (for example, applying the concept "red" to the concept "apple" creates a new concept, "red apple," which can be viewed as a restricted subclass of the original concept.

Denotative semantics are difficult to model for quantitative purposes, but relational semantics are not. Thus, in our initial explorations, we will act "as if" concepts have arbitrary and unknown denotative meaning, and focus our attention upon the effects of explicit relational semantics on the temporal and spatial patterns of traits.

A. Representing Relational Semantics as Ontologies

[TODO: Here, I want to introduce description logics as frame-based concept reasoning languages, define an ontology, give a small example of the relational semantics defined in such ontologies, and describe how ontologies allow various types of concept deduction and reasoning. The point of using a formal system for this, rather than just building arbitrary networks, is that we want to investigate the consequences of types of transmission bias that involve conceptual fit and reasoning, and a lot is known about the formal properties of various classes of logical reasoning systems.]

B. Examples of Concept Trees

$$\begin{aligned}&\mathsf{hasParent} \sqsubseteq \mathsf{hasAncestor}\\&\mathsf{hasChildren} \sqsubseteq \mathsf{hasDescendant}\\&\mathsf{Parent} \equiv \mathsf{Human} \sqcap \exists \mathsf{hasChildren}\end{aligned}$$

Note that this kind of representation is probably too low level. See Section V for details.

III. RESEARCH QUESTION NOTES

Clearly, we need some rough research questions to focus the exploration.

1. "Simple Effects": The first experiment we should run is to seed a population with randomly constructed semantic networks (possibly random draws from Cyc, rewritten into OWL, so that the kinds of relationships between concepts, their frequency, and density of relationships are all realistic), and then examine the effects on the distribution of relationships, relationship density, etc, if we run a simple unbiased transmission (i.e., random copying) and conformist bias process within the population, where what is copied is simply the concepts, with relationships only remaining if the target individual already has the "other" end of a particular relationship. At least naively, I would assume that unbiased transmission would cause the randomization of concept networks, with relationship R between concepts A and B declining in frequency due to pure drift. I'm not sure what I think conformist transmission would do, but I guess I'd expect less loss of relationships between concepts.

---

behavior, but our point is that all are necessary and none are necessarily alternatives to one another.

2. “Relationship Imitation”: What are plausible models of imitation for relationships between concepts? Perhaps we run an unbiased transmission process, but when a target individual selects a model individual to imitate, we include relationships in the pool of things that could be imitated. In this model, I’d expect we’d lose both concepts and relationships to drift, if we don’t have model mutation of both concepts and relationships. In simple conformist models, we’d see fairly fast convergence to a “worldview”: a set of concepts and relationships that tend to travel together, even though the underlying process is not replicating a whole worldview, but instead just the component parts. The “worldview” in this type of example, if I’m right, is a purely epiphenomenal (or “emergent,” but I tend to agree with Epstein on not using this term) pattern.

3. “Worldview Similarity Bias": One experiment we could run would be to select a focal individual each tick in a Moran process, and then examine the set of candidate models, finding the model(s) which have the most number of traits in common *which share the same set of semantic relationships*, and select one of their traits randomly to copy. Do we end up with persistent associations of concepts? How does the turnover rate of traits change? What level of variation do we see in the population - is it lower, higher, or the same as we’d see in a pure neutral model? We’re essentially doing unbiased selection of traits from a biased selection of individuals.

4. “Worldview Consistency Bias": We can also determine whether to copy a trait from a random individual based on whether the trait selected for copying is logically consistent with the rest of the focal individual’s ontology. If the concept (and its attendant relationships from the target individual) cannot be integrated into the focal individual’s knowledge base without violating KB consistency given DL rules, we don’t copy it. This models the situation were people have absolute bias against inconsistent worldviews. What happens in the population as a whole? Does it split into non-communicating subpopulations rapidly? Clearly that would happen if the rule was about not taking any concepts if the focal individual’s KB has *any* inconsistencies with the target individual, but this is a softer rule, which allows us to take some traits, but not others, from folks we have differing belief structures from...

IV. RAW NOTES

1. Modeling relationships between traits as semantic relationships between concepts
    (a) Note that I’m not claiming that traits are concepts, or that traits are related in the same way concepts are related, but that we’re going to model them this way and see if the parallel modeling is useful.
    (b) Clearly, there is some relationship between concepts and blobs of information that get communicated, but for the moment let’s gloss over that and treat them as isomorphic. Later, we’ll come back and assume that they’re not, and that traits are instead nested sets of concepts of varying size and scope, and see how this assumption changes our model.
    (c) Either way, we care about having a simple and formal way of modeling the varying semantic relationships two concepts bear towards one another.
    (d) Fortunately, the knowledge engineering and semantic web folks, in cooperation with AI workers and logicians, have figured out the formal properties of several classes of “description logics” that allow us to formally describe concept relationships.
    (e) When I refer to an “ontology,” in this work, I’m describing a set of concepts and a set of relationships between those concepts, with the relationships being drawn strictly from those available in a particular named description logic.
    (f) The reason we have to specify the exact description logic we’re using for modeling is that DL’s vary strongly in their descriptive power (i.e., the catalog of relationship types they model within the DL), and thus in the computational complexity of reasoning and inference tasks within that DL. Some description logics are fairly limited in their catalog of relationship types, and possess fairly manageable computational complexity; others have deeper catalogs of relationships, but may require exponential time/space to return answers to inferences. All are simplified subsets of standard first-order predicate logic (FOPL), which is descriptively rich but provably undecidable.
    (g) For examining simple models of transmission with semantically structured traits, we don’t care much about the richness of relationships available (although we might in trying to apply this type of model to actual empirical cases down the road). So I’ll keep it simple and use the DL “dialect” favored by the W3C for Semantic Web use, given its decidability, decent computational complexity, and most especially because of the tool and software support for it. In formal terms (Horrocks XXXX) this dialect is referred to as **SHIQ**, while in the software and semantic web community it is instantiated by OWL-DL which embeds RDF and RDF-Schema.

V. 2013 RETROSPECTIVE NOTES

The research question section here is still right on the money. We would add interests in whether conformism and homophily, construct meaningful clusters of concepts and knowledge relations within a population, which can be described by terms like "occupational specialization" and whether such “coevolution” of meaning and poplation, rather than purely demography, is the reason for the difference in apparent toolkit diversity between the Middle and Upper Paleolithic, for example.

Also, the focus on pure description logics here is somewhat misplaced in 2013, although it made sense given the research and technology available in 2003–2007. Today, we’re exploring the rich databases of the ConceptNet project for ways of describing various types of semantic relationships which are “higher level” than the logical entailments of DL’s.