

Entities, Concepts, and Propositions

a computational approach

Edward O'Connor

31 December 2000

1 Before Starting

My training and background has been in science, not philosophy. My writing is expository, not argumentative. I am not attempting to prove important philosophical points. The point of this work is to posit initial, working hypotheses.

Many of the ideas here are taken for granted as very basic by at least several of my potential readers. So I apologize to anyone who becomes frustrated by the pace of presentation here.

I have found that this framework offers a very simple description of disparate issues in epistemology, but this simplicity does not come from a lack of robustness. The system appears to have great explanatory power for a variety of phenomena, and helps to highlight those areas where work is more urgently needed. That is, in a nutshell, why I consider getting these ideas to a wider audience to be important.

Much of the following may sound strangely familiar, if in an odd form. I've found that, over the course of thinking about these issues, that I've almost come full circle. Initially, I was rather naive, too blinded by Rand's writing prowess to find any flaws in IOE. Following this, and in the beginning stages of working on this work, I was very critical of Rand's work in IOE, and saw what I thought to be gaping holes in her epistemological system. And now, having constructed this framework for thinking about entities and concepts, I've found that I've somewhat accidentally reconstructed many of Rand's conclusions that I had previously found wanting.

*... the end of all our exploring
will be to arrive where we started
And know the place for the first time.*

T.S. Eliot

2 Goals and Provisions

This project is to present a mathematical, computational framework for modeling some of the functions of conscious agents, namely entities, concepts, and propositions. The framework developed should be at the same time simple enough to be easy to understand, and also not so simple that it lacks usefulness.

If possible, the framework should be applicable to disparate forms of consciousness. Ideally, it should be useful in artificial intelligence, cognitive science, and philosophy. It should help us to understand the mental functioning of animals and humans, as well as various forms of agents (connectionist or otherwise).

Additionally, it would be wonderful if possessing such a framework simplified our general philosophical understanding of entities and concepts, or provided new insights into the same.

Some notes on terminology before proceeding: I see no useful distinction between the terms “entity” and “percept.” I will almost always use the first term, though quotes may contain the latter. Also, I will typically use the term “agent” as a generalization for a naturally- or artificially-conscious entity.

Throughout the above and in what follows, I have very much tried to avoid using the word “implicit,” as I have found that it is used in very many different ways by different authors, Rand included, and at this point I have trouble understanding what is meant by it. So you may find that some of the processes described below are “implicit” or “implicitly performed,” or perhaps they are not implicit. I do not know, nor do I think that it matters, given the intended generality of the model.

3 Entities

In order to develop a model for representing entities, we should recall what they are and how they come about.

Radcliffe and Ray define ENTITY* as “a mind-dependent creation produced by a conscious subject’s focusing on some portion of reality in such a way as to proscribe an edge.” [Radcliffe and Ray] For the purposes of this work, I will only focus on those entities which arise due to an agent’s attention to discontinuities experienced by the agent’s sensory apparatus.[†]

Rand describes her notion of percept in “The Objectivist Ethics:”

“A ‘perception’ is a group of sensations automatically retained and integrated by the brain of a living organism, which gives it the ability to be aware, not of a single stimuli, but of *entities*, of things.”

[Rand 1964, p. 20]

* Here and elsewhere in this paper, I follow the philosopher’s convention of writing FOO to refer to the concept, “foo” to refer to the word itself, and foo when using the word normally. Unfortunately, this convention is not followed in most of the works I quote.

[†] Put another way, I am not touching the issue of the block of air in front of you with a sixty foot pole. Suffice it to say that, while I agree with Radcliffe and Ray’s conclusions, it would be counterproductive for me to get bogged down in arguing for their conclusions here.

So an entity is a grouping of attributes which exists mind-dependently.

If entities are groupings of attributes, what are these attributes which are being grouped? I do not have anything resembling a decent answer to this question. So for this work, I restrict myself to those kinds of attributes directly available to the sensory apparatus of the agents in question. In examples, I try to limit myself to those attributes which appear to be directly available in human awareness, if for no other reason than the familiarity the reader has with such modalities as color, shape, sound, texture, smell, and taste.

Entities come into existence in an objective relationship between an agent and agent-independent reality. This does not occur without an active process on the part of the agent. The agent must perform some variety of task or tasks in order to form an entity from its sensations. This process of grouping attributes, the process of entity-formation, is carried on in some way that is heavily dependent on two things: the identity of the external world and the identity of the agent. The three most important elements of the agent's identity in this matter are the agent's scale of observation, its modes of awareness, and its purpose. What this process is is unimportant. What is important, is that there is some particular process or set of processes for any particular kind of agent. I hope to arrive at a model for the output of these processes, without knowing much if anything about the processes themselves, as they most probably differ between sufficiently different agents.

Entities exist in the mind, and everything that exists, exists in some way. So entities must exist in the mind in some manner. Clearly, it doesn't even make sense to entertain the notion that they exist, but in no manner. Another way of stating this basic fact is that, when agents form entities, they *encode* them in some manner. In fact, the formation process and the encoding process are one in the same.

So, in our project to develop this framework, a crucial element, or perhaps even *the* crucial element, is a general model of the encoding of entities. Our immediate task: to find an encoding that is suitable for our purposes, in that it is general enough to work with any kind of conscious agent, and at the same time specified enough such that it possesses useful properties.

Luckily, such an encoding not only exists, it is very straightforward: the vector. A vector is simply an ordered collection of values. For instance,

$$\begin{pmatrix} 7 \\ 9 \\ 4 \end{pmatrix}$$

is a three-dimensional vector with integer components 7, 9, and 4. In modeling entities with vectors, each vector dimension corresponds to some particular kind of attribute. Recall Rand's description: "A 'perception' is a group of sensations. . . integrated by the brain of a living organism. . ." [Rand 1964, p. 20] The vector offers us a simple way to perform the desired grouping.

Researchers in artificial intelligence have developed all sorts of interesting and complicated encodings for enti-

ties, but it can be shown that all more complicated encodings can be reduced to vector encodings.* So this way of representing entities will be able to do anything that we can do by other means.

It should be noted that the objective account of entity-hood allows for different agents to form radically different entities. Because human beings have very similar scales of observations, sensory modalities, and purposes, we tend to arrange the world in very similar ways. But we must avoid the realist leap and impose our division of the world onto other agents. This means that, insofar as the goal of artificial intelligence is to produce an artificial, conscious being, researchers shouldn't be very discouraged when their creations see the world very differently than we do. This is to be expected. Dreyfus rightly found research in what he called "Grand Old-Fashioned AI" to be lacking in this respect: "One needs a learning device that shares enough human concerns and human structure to learn to generalize the way humans do." [Dreyfus, p. *xlvi*] So while the expectations that many AI researchers have regarding the creation of human-like AI in the near-future are perhaps misguided, I do not consider this to be reason for lamentations.

4 From Entities to Concepts

For a variety of reasons, we experience change in our sensory experience, and yet we identify entities that persist over time. How this happens deserves a paper in its own right, and I do not presume to have much of an idea on how to tackle it.

To handle some of this variation, it seems quite natural to say that entity-vectors have some tolerance for experiential variance. In other words, the entity-vectors are a little fuzzy. For instance, consider the fact that we rarely experience an area in our visual spectrum that is almost totally devoid of color variation, and yet I am comfortable with considering the computer in front of me to be blue, without much regard to the relative location of the light sources in this room and all of the other relevant factors.

Such attribute tolerance can be built into our representation by defining each value in the entity-vector as a range and not a specific value. It specifies the acceptable range in which this attribute may appear. Consider this point in light of Radcliffe and Ray's development of the concept UNITY over the course of their paper. A unity is an area in the sensory field which appears continuous relative to its surroundings. The continuity here does not preclude variation; it actively embraces it. The stipulation is just that the variation be considerably less than the discontinuity which surrounds the unity.

Recall Rand's process of concept formation by measurement omission. She makes the rather bold claim in IOE that, not only do adults explicitly form concepts via this method, but also that children use this method, without being consciously aware of it. After describing measurement omission for forming the concept LENGTH, Rand claims that

* Such a demonstration of this claim would be rather technical and thus out of place here. Consider this common-sensical explanation: all of those more complicated encodings are used on computers. Computer memory is a big vector. Therefore, all more complicated encodings can be handled by a vector, because they already *are* being handled in such a manner all the time.

while “the child does not think in such words. . . *that* is the nature of the process which his mind performs wordlessly.” [Rand 1990, p. 11] There is a striking similarity here, between the some-but-any variance of Rand’s concepts and the variance tolerance of these entity-vectors. Indeed, one is left with the impression that, to represent concepts in this system, one need only use vectors “fuzzier” than entity-vectors. I should note that I am restricting myself to considering sortal concepts: those concepts which demark a similarity class of entities.*

There doesn’t appear to be any need to encode entity-vectors and concept-vectors differently. Indeed, a vector of acceptable attribute-ranges would serve quite well as an encoding of a concept, even if we weren’t also encoding entities in that way. So, allow me to introduce a useful bit of terminology which will help during the rest of this paper: an “*ec*-vector” is a vector encoding of some entity or some concept.

Concepts refer to an open-ended collection of entities. Which entities? Those that fall within the area proscribed by their *ec*-vector in the vector space. Concepts are able to subsume unexperienced entities as well as experienced ones because, were the unexperienced entities to be experienced, they too would fall within the concept’s area in vector space.

In some sense, concepts can be formed simply by the experience of an entity (the formation of an *ec*-vector) because of the equivalence of representation. When an agent experiences another entity sufficiently like the first, it can recognize it as “another one of those.” That is, it is able to perform some perceptual categorization on the new entity, simply because it experienced one similar to it previously. It should be noted that, in many systems, explicit memory of the previous experience is *not* necessary to make the later perceptual categorization: the influence on the agent’s state by the previous experience is often enough.

Because of these issues, I find that there are many cases in which Rand’s requirement of two or more exemplars for concept-formation to be unnecessary. While I think that, all other things being equal, using multiple entities tends to yield more useful concepts, I don’t think that concept formation *requires* two or more.

Also, in order for an agent to utilize a concept, such as in the “another one of those” example above, a word is not required. Perceptual categorization is something that we are able to do well before we develop language. Consider:

It has been shown, that, by three months of age, infants have begun to apprehend categories of events and objects. They can sometimes treat an item not seen earlier as an instance of a familiar category while recognizing nonetheless that the item is distinct from ones previously seen. [Boydston, p. 14]

Rand focuses on the deliberate, volitional nature of her process of concept formation in IOE. I think that concept formation is most probably more analogous to breathing: while we can perform the process with direct, conscious attention, we are also able to just “let it happen,” that is, we are able to let our conscious subsystems handle the task.

* Now, it may very well be the case that this kind of concept is the only kind of concept, but I would just assume not worry about that right now. That is another topic for another time.

What requires deliberate attention from the upper levels of consciousness is the *validation* of concepts which have been formed; for this, the agent must ask Rand's Question.

Rand speaks of "the faculty of sensation," "the faculty of perception," and "the conceptual faculty" as though these functions of consciousness were very vividly separated. But one of the deficiencies of this strictly-separated, three-layer model of consciousness (or at least her presentation of it in "The Objectivist Ethics", IOE and elsewhere) is that, using it, it is somewhat difficult to understand how the upper layers of consciousness could have come about from the lower layers during the evolutionary process. In contrast to this, this vector model offers a clear and straightforward path for agents to bootstrap conceptual awareness out of their perceptual ability, because of the clear similarity of representation.

5 Language and Propositions

For our purposes here, I restrict claims about propositions to those proposition which take this general form: "This entity or class of entities possesses or does not possess this attribute or set of attributes in this range or these specified ranges." Propositions of this form can be easily represented as vector equations in this system.

Propositions are formed deliberately by the agent. In order to do this, the agent needs to possess a method to "call out" or "activate" the appropriate *ec*-vectors at the appropriate times. This is because all of the *ec*-vectors possessed by the agent cannot be kept in active attention at the same time (the principle of the crow).

One way to conjure up the appropriate *ec*-vector is to explicitly experience one of its referents. Unfortunately for the agent, the world is such that this is often not possible, so another method must be found. Remember that, in the ancestral environment, the world was much the same in this respect, so living things needed to evolve another method.

If the agent can't cause itself to explicitly experience the appropriate sort of thing, perhaps it can do the next best thing: the agent can conjure up some other entity that it, for one reason or another, associates with the desired *ec*-vector.

This is, roughly, the role of language.* "Words remind us of thoughts. Common nouns are the first conceptual words we learn. . . , and they evidently cue us to think of similarity groups. . ." [Boydston, p. 29] The agent is able to affect some change upon the world (the speech act) which causes itself to experience an entity (the word). The agent has built up some kind of mapping between the word and the *ec*-vector, such that when it experiences the word, the *ec*-vector to which it is bound is activated.

Rand claims that, "in order to be used as a single unit, the enormous sum integrated by a concept has to be given the form of a single, specific, *perceptual* concrete, which will differentiate it from all other concretes and from all other

* I am restricting myself to spoken language, for various reasons, though the broad thrust of this argument is one that I think may also apply to written language.

concepts.” [Rand 1990, p. 10] In one sense, I concur, but in another, I disagree. I find her statement to be ambiguous because of the multiple readings of the verb “to use” in this case. If she means that, in order for an agent to be able to consistently call up a specific concept, the agent needs to have some symbol bound to it in some way, then I agree. But, if she means that, in order for an agent to ever utilize a concept, the agent needs to have a symbol bound to it, then I disagree. There is a clear difference between these two sorts of use.

She continues:

Language is a code of visual-auditory symbols that serves the psycho-epistemological function of converting concepts into the mental equivalents of concretes. Language is the exclusive domain and too of concepts. Every word we use (with the exception of proper names) is a symbol that denotes a concept, i.e., that stands for an unlimited number of concretes of a certain kind. [Rand 1990, p. 10]

Given my stance on the rough equivalence of entities and concepts, at least insofar as their internal representations go, it should come as no surprise that I don't think Rand needs to make an exception for proper names as she does here. On this model, it makes perfect sense to treat the *ec*-vectors denoted by proper names no differently than we treat the *ec*-vectors denoted by other words. This view better explains the common notion of personal identity which allows for people to act in variance to their expected behavior. “That wasn't like you” amounts to “that wasn't like my concept of you,” on this view.

Since they are vectors, *ec*-vectors can be manipulated together for various effects. There are many such possible manipulations, but only a few will be described here, as they will turn out to be quite useful later on.

Radcliffe and Ray explain that the concept of SIMILARITY “arises from our awareness of degrees of difference — that some things are less different from a given object than others.” [Radcliffe and Ray] This vector model suggests a very clear understanding of the concept SIMILARITY: two entity-vectors \vec{a} and \vec{b} are similar along some dimension or some dimensions if the magnitude of their vector difference taken on those dimensions, $|\vec{a} - \vec{b}|$, is regarded as small, relative to some greater difference.

Note that this notion of similarity allows us to consider two things as similar even when we have no concept of the attributes over which they are similar. I follow Boydston in holding that “To detect an attribute, one need not already have formed a concept of it. . . Perceptual pickup of attributes is sufficient for similarity grouping of objects.” [Boydston, p. 28] To consider two entity-vectors as similar, the agent does not need to know much if anything about the attributes which are bundled up into the entities, it merely needs to be able to calculate the vector difference. Indeed, “the course of speech development suggests that concepts of entities are formed before concepts of attributes.” [Boydston, p. 28]

One *ec*-vector can be used to “mask out” certain attributes of another *ec*-vector. Masking along these lines has been used in computer science and mathematics for a very long time. For instance, consider some *ec*-vector \vec{a} , and another

ec -vector \overrightarrow{color} . The \overrightarrow{color} ec -vector is a concept of an attribute, namely color. When we take the componentwise product of these two vectors, we are left with a new ec -vector, whose only specified term is the color range of \vec{d} . It is in such a manner that we can use ec -vectors to isolate certain attributes of other ec -vectors.

Keep in mind that each component of an ec -vector is a ranged value. Given this, set operations such as union and intersection can be performed with ec -vectors. By such means, words which modify other words — adjectives, adverbs, and the like — can be modeled in our system. Consider the ec -vector \overrightarrow{ball} . By intersecting this ec -vector with \overrightarrow{blue} , $\overrightarrow{blue} \cap \overrightarrow{ball}$, we are able to mentally isolate the area of vector space in which all blue balls lie.

This leads to the interesting conclusion that the ec -vectors $\overrightarrow{swiftly}$, \overrightarrow{swift} , and $\overrightarrow{swiftness}$ are all essentially the same, as are all of the analogously-related ec -vectors. The difference between these words is apparently only grammatical; having the different words helps us to better distinguish their use in propositions.

6 Definitions, the Crow, and the Essential

A definition is a proposition which is intended to identify the referents of a concept. In other words, a definition is a proposition which identifies the area in vector space which the concept encircles.

“Words transform concepts into (mental) entities; *definitions* provide them with *identity*. (Words without definitions are not language but inarticulate sounds.)” [Rand 1990, p. 11] While I agree with Rand that a definition *specifies* the identity of the defined term, her parenthetical remark raises several concerns.

One one hand, I must emphatically disagree with her claim that “words without definitions are not language but inarticulate sounds.” We go most of our lives using many concepts we haven’t defined, and yet it is certainly not the case that these concepts lack identity.

On the other hand, she is correct if she means “words without concept bindings” when she says “words without definitions.” In such cases, the word really is just a sound, as it lacks any connection to ideas.

We can in some sense “unpack” an ec -vector into a set of ec -vectors, perhaps by performing several mask operations as described earlier. These ec -vectors, when intersected with one another, produce the original. Such a process is potentially useful for several reasons. For the moment, consider the case of unpacking an ec -vector into many ec -vectors, one for each attribute of the original:

$$\begin{array}{c} \text{Some } ec\text{-vector} \\ \overbrace{\left(\begin{array}{c} x_{1,1} \dots x_{1,2} \\ x_{2,1} \dots x_{2,2} \\ \vdots \\ x_{n,1} \dots x_{n,2} \end{array} \right)} \\ \\ \end{array} = \overbrace{\left(\begin{array}{c} x_{1,1} \dots x_{1,2} \\ 0 \\ \vdots \\ 0 \end{array} \right) \cup \left(\begin{array}{c} 0 \\ x_{2,1} \dots x_{2,2} \\ 0 \\ \vdots \end{array} \right) \cup \dots \cup \left(\begin{array}{c} 0 \\ \vdots \\ 0 \\ x_{n,1} \dots x_{n,2} \end{array} \right)} \\ \text{Each of its attributes, isolated} \end{array}$$

(Here, I use the symbol \emptyset to denote some attribute value which signifies that the attribute is unspecified.)

The above unpacking yields an exhaustive definition of the *ec*-vector in question. Now, consider above definition in light of the principle of the crow. Such a definition is, for most entities, unwieldy at best. We have a tradeoff situation between, on one hand, accurately describing the area in vector space that this *ec*-vector surrounds, and on the other hand, being able to use the description given our limited capabilities.

So the goal is to find some description of the *ec*-vector that is both accurate and concise. Given the n different attributes that appear above, the agent must select some number of the attributes to include in the definition. There must be no more included attributes than the agent can reasonably handle at one time given the crow.

What I have just described amounts to a linear optimization problem which can be straightforwardly solved using the tools of operations research. The general result of such optimization is a small set of attributes which accounts for the *ec*-vector better than any other such small set.

Stepping back to Rand, we find that, in her writings on the process of definition, she emphasizes the role of essential characteristics when defining concepts. To find essential characteristics, Rand proposes what she calls the “rule of fundamentality.” A fundamental characteristic, to Rand, is “the [characteristic] which explains the greatest number of others.” [Rand 1990, p. 45] Or, in our terms, a fundamental characteristic is one which is included in the result of optimizing characteristic-inclusion in the definition of an *ec*-vector. We have arrived at a computational understanding of what makes up the essential. This is one point about which I am particularly pleased, as I do not think it was at all obvious before how to determine essential features of entities.

I previously claimed that this optimization is straightforward to solve. This was hyperbole. In actual fact, while this problem is not impossible to solve, it is NP-complete. This is the computer scientist’s way of saying that it is very difficult. In order to find a decent solution in a fairly short amount of time, the agent needs to employ a good heuristic.

Rand advocates the use of the genus-differentia form in defining terms. This form of definition falls nicely out of the above optimization approach, because it is an *excellent* heuristic for quickly cutting the optimization problem down to something much easier to handle. By specifying the genus in the definition, the agent has done most of the work in locating the concept in vector space. All the agent needs to do now is to perform the optimization over only those attributes which distinguish the concept from others in the genus. Hence, the differentia.

7 Concepts, Referents, and Hierarchy

In IOE and elsewhere, Rand emphasizes and insists upon a conceptual hierarchy of her own devising. While I have found her hierarchy to be generally nice, much of her argumentation involving it has struck me as contrived at best. Consider this excerpt from Chapter 3:

Observe that the concept “furniture” is an abstraction one step further removed from perceptual reality

than any of its constituent concepts. “Table” is an abstraction, since it designates *any* table, but its meaning can be conveyed simply by pointing to one or two perceptual objects. There is no such perceptual object as “furniture”; there are only tables, chairs, beds, etc. The meaning of “furniture” cannot be grasped unless one has first grasped the meaning of its constituent concepts; there are its link to reality. [Rand 1990, p. 22]

Now, consider the concept FURNITURE: FURNITURE does not refer to TABLE, CHAIR, BED, and so on. Its referents are not concepts. It refers to furniture, some of which are the same entities to which TABLE refers, some of which are referred to by CHAIR, etc. So, while FURNITURE is a *wider* concept than say, TABLE, it is not any further removed from perceptual reality than TABLE, as both concepts refer to entities.

Rand’s claim hinges on the apparent necessity to invoke concepts such as TABLE and CHAIR when attempting to define FURNITURE. This is where we part ways. For while I would agree with Rand that it is awfully nice to have the concepts CHAIR and TABLE around when attempting such a definition, because you can use “chair” and “table” to shorten your definition, I disagree with her claim that these concepts are somehow necessary.

Certainly, as concepts get wider, the difficulty in succinctly defining them without such narrower concepts increases greatly. This is a demonstration of the Objectivist tenet that we form concepts to reduce cognitive overhead. It is so much easier to use the narrower concepts in our definitions that we are tempted to overlook this fact. So I find that Rand’s hierarchy of definition is at minimum not as rigid as she would have liked. I actually take this much farther, and have begun to use the above framework to develop an understanding of what I call “conceptual taxonomy.” A conceptual taxonomy is a proposed hierarchical relationship between concepts.

At a certain level of abstraction, concepts are “flat”: each and every concept has referents, and this is the extent to which the hierarchy goes. You have a level for concepts and a level for referents. On this view, how are concepts who have some concepts as referents handled? Consider in the most extreme case the concept CONCEPT: each and every one of its referents is a concept. But since all concepts are entities, such impredicative concepts do not need to be treated specially.

There are many, many dimensions along which we abstract. Any of these dimensions, or any set of them, may be chosen as a base for forming a conceptual hierarchy. But this is not a *carte blanche* to do as you will: depending on your purpose, some hierarchical relationships are more useful than others. For instance, Radcliffe and Ray point out that “within biology there are at least three different concepts of the taxonomic boundary of a biological species — cladistics, numerical phenetics and evolutionary taxonomy — all of which appear to be equally valid.” [Radcliffe and Ray] This is a clear example of the wider principle: conceptual taxonomy heavily depends on the agent’s purpose, and no one taxonomy should be elevated above others in some a priori way, such as Rand elevates her hierarchy.

8 Concluding Remarks

While the Objectivist epistemology offers a wonderful foundation for work in artificial intelligence (as it does for most every other endeavor), Objectivists can also learn quite a bit from AI.

Before entertaining many of the ideas above, I lacked a decent understanding of how to determine the essential characteristics of an entity or concept. I can now not only determine these characteristics, but I can do so in a straightforward fashion.

Much of Rand's writing on her conceptual hierarchy bothered me, but I had much trouble pinpointing what specifically it was about it that irritated me so. I am no longer bothered by the hierarchy issue, and have begun to develop a more robust understanding of how our knowledge can be regarded as hierarchical.

I spent my formative philosophical years immersed in Rand, and have read some, but not much, of other philosophers, and thus my writing here may not thrill my more critical readers. This is a good thing. I want philosophical objections to be raised, because I want to know what I have to work on.

I expect that this work has left you with many unanswered questions. I know that I, for one, have many.

Works Cited

- [Boydston] Boydston, Stephen. "Capturing Concepts." in *Objectivity*, 1:1, pp. 13 – 41.
- [Dreyfus] Dreyfus, Hubert. *What Computers Still Can't Do*. Cambridge: MIT Press, 1992.
- [Radcliffe and Ray] Radcliffe, Tom, and Ray, Carolyn. "Edges, Entities, and Existence: an epistemological excursion." Presented at The Objectivist Center's Advanced Seminar, 1 June 2000. <URL:<http://enlightenment.supersaturated.com/essays/text/carolynray/-RayRadcliffe.html>>.
- [Rand 1964] Rand, Ayn. "The Objectivist Ethics." in *The Virtue of Selfishness*. New York: Signet, 1964. pp. 13 – 39.
- [Rand 1990] —. *Introduction to Objectivist Epistemology*. 2nd ed., Binswanger, Harry, and Peikoff, Leonard, eds. New York: Meridian, 1990.