

# DEBRA: On the Unsupervised Learning of Concept Hierarchies from (Literary) Text

Peter J. Worth<sup>1</sup>, Domagoj Doresic<sup>2</sup>

<sup>1</sup>Department of Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, USA

<sup>2</sup>Interdisciplinary Research Unit Mathematics and Life Sciences, University of Bonn, Bonn, Germany

Email: [pworth2022@fau.edu](mailto:pworth2022@fau.edu), [ddoresic@uni-bonn.de](mailto:ddoresic@uni-bonn.de)

**How to cite this paper:** Worth, P.J. and Doresic, D. (2023) DEBRA: On the Unsupervised Learning of Concept Hierarchies from (Literary) Text. *International Journal of Intelligence Science*, **13**, 81-130. <https://doi.org/10.4236/ijis.2023.134006>

**Received:** August 23, 2023

**Accepted:** October 20, 2023

**Published:** October 23, 2023

Copyright © 2023 by author(s) and Scientific Research Publishing Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

---

## Abstract

With this work, we introduce a novel method for the unsupervised learning of conceptual hierarchies, or concept maps as they are sometimes called, which is aimed specifically for use with literary texts, as such distinguishing itself from the majority of research literature on the topic which is primarily focused on building ontologies from a vast array of different types of data sources, both structured and unstructured, to support various forms of AI, in particular, the Semantic Web as envisioned by Tim Berners-Lee. We first elaborate on mutually informing disciplines of philosophy and computer science, or more specifically the relationship between metaphysics, epistemology, ontology, computing and AI, followed by a technically in-depth discussion of DEBRA, our dependency tree based concept hierarchy constructor, which as its name alludes to, constructs a conceptual map in the form of a directed graph which illustrates the concepts, their respective relations, and the implied ontological structure of the concepts as encoded in the text, decoded with standard Python NLP libraries such as spaCy and NLTK. With this work we hope to both augment the Knowledge Representation literature with opportunities for intellectual advancement in AI with more intuitive, less analytical, and well-known forms of knowledge representation from the cognitive science community, as well as open up new areas of research between Computer Science and the Humanities with respect to the application of the latest in NLP tools and techniques upon literature of cultural significance, shedding light on existing methods of computation with respect to documents in semantic space that effectively allows for, at the very least, the comparison and evolution of texts through time, using vector space math.

## Keywords

Ontology Learning, Ontology Engineering, Concept Hierarchies, Concept Mapping, Concept Maps, Artificial Intelligence, Philosophy, Natural

---

---

Language Processing, Knowledge Representation, Knowledge Representation and Reasoning, Machine Learning, Natural Language Processing, NLP, Computer Science, Theoretical Computer Science, Epistemology, Metaphysics, Philosophy, Logic, Computing, Ontology, First Order Logic, Predicate Calculus

---

## 1. Introduction

This work originates out of ongoing research into the potential application of modern machine learning and artificial intelligence paradigms with respect to metaphysics. It follows on the heels of a paper on semantic geometry which studies the underlying data representation layer for most modern large language models (Word Vectors essentially) [1], which itself follows on the heels of research into the creation of metaphysical reference architecture that applies the principles of quantum probability theory to human decision making through ideological (intentional spelling) space [2]. This paper takes an additional step forward in that we develop a prototype conceptual hierarchical constructor which looks to generate conceptual maps, effectively directed graphs of conceptual nodes tied together by actions (subject-object pairs), to see what might be possible with respect to metaphysical inquiry when applying ontology learning techniques onto specific literary texts of cultural significance, taking the Bible for example.

From an enterprise application development perspective, the data modeling space has taken a significant leap forward technologically with the advent of UML and the ability to both forward and backward engineer models, effectively forms of entity relation diagrams but other UML models as well, both from data base and other data persistence layers (like hibernate for example), as well as more standard database management systems, *i.e.* SQL. Enterprise Architect is a tool for example that supports some of these functions and this type of modeling, and self-documentation you could call it. It is now ubiquitous in the enterprise application development domain. These advancements are all in the last two decades or so.

In our research into NLP, we landed upon a similar type of modeling problem, with a more generic solution that was arguably still in its early stages of development, at least relative to the type of modeling capabilities that existed in the enterprise application development domain. This was Ontology Engineering, which was fundamental to this notion of the Semantic Web, or Web 3.0, which looks to create inter-operable, non-proprietary, decentralized, linked data structures referred to as *ontologies*, which facilitate machine understanding. This is a vision of Tim Berners Lee and has been an active area of research for some two decades, and yet the tools to facilitate this semi-automated, labor intensive and domain-specific process, still had a ways to go and had most certainly not been applied to literary works within the humanities beyond a research paper or two.

While DEBRA doesn't necessarily fill this gap entirely, the general problem of

ontology learning is more difficult and challenging than the authors who had time or resources to fully invest in, with DEBRA however, we did have an opportunity to explore the core conceptual mapping piece of the ontology engineering problem as it were, and as such we most definitely come to a greater understanding of both what is possible with respect to the (unsupervised) learning of conceptual hierarchies from literary text, the overall complexity of the problem of ontology learning itself writ large, and also a greater understanding of the differences between various forms of knowledge representation which facilitate machine learning, or ontology learning more specifically, in general.

In this paper, we look to establish the shared intellectual heritage of computer science, AI, Logic, epistemology, metaphysics and ontology, all of which come together to form some of the foundations of a significant portion of theoretical computer science in areas such as AI and Computing Theory, and we look to inform this subject area with a slightly different approach to knowledge representation, one that comes from the cognitive science world but one that in many respects reflects the most fundamental of forms of knowledge representation with perhaps the widest application across the most disparate fields, namely concept graphs, or maps<sup>1</sup>.

This introductory material is then followed by a technical deep dive into the prototype that we developed that we affectionately call DEBRA, short for dependency tree based concept hierarchy constructor, looking at the precise way we extract concepts from the texts in question, and how we establish the relationships between the extracted concepts, as well as how we impose an ontological (hierarchical) structure on said concepts. This detail is then followed by a brief review of how the transformation, really TF-IDF vectorization, of texts in general supports the ability to perform analysis against these literary texts in vector space, after which we briefly summarize the work and provide some concluding remarks, which include the places where we think we've broken new ground as well as areas of potential further research.

Essentially with DEBRA, we look to establish an intellectual beachhead between Computer Science and the Humanities more broadly for the analysis of literary texts as artifacts of knowledge in and of themselves, as opposed to the analysis of massive digital asset libraries for the purpose of searching and sorting, the latter being the primary goal of much of the NLP research today (understandably so). We think this area of research is ripe for innovation and discovery right now and we hope this paper helps move this branch of knowledge along accordingly.

## 2. On Knowledge Representation, AI & Ontology

Any technically accurate modern definition of AI must include both the notion of *thinking*, as Turing proposes it in 1950 [3], as well as this is a more modern

<sup>1</sup>Arguably neural networks are a form of concept graphs, with a well-defined structure and geometry but ultimately the structure is the same just at a different level of intellectual, or metaphysical, resolution.

addition to the field and definition—*learning* as it is understood from a Machine Learning (ML) perspective. AI is intended to denote applications that “learn”, which distinguish themselves from applications that perform first level analysis on data sets by anticipating future outcomes based on (data from) past results as well as applications that are capable of performing relatively straightforward decisions or forms of reasoning such as *if x then y*.

This area of research however, despite having its roots and foundations laid down some 70 years ago<sup>2</sup>, has been enjoying quite a renaissance in the last few years, primarily driven by the availability of massive computing resources at a relatively reasonable price, along with technical innovations in the field itself (various language models, neural networks, etc.), the combination of which have come together to provide the driving force for a massive wave of technological innovation that rivals the advance of the Internet itself. The technological ramifications are in fact so great that earlier this year (2023) over 1000 AI researchers from around the globe signed a letter that proposed a “pause” on AI development given that the field poses a “profound risk to society and humanity”<sup>3</sup>.

Ontology Engineering, or Ontology Learning, as a branch of NLP and AI, sprung from the need for the construction of domain-specific ontologies (formally defined below) for the Semantic Web at the turn of the millennium more or less. These frameworks, the specifics and details of which we get into in the section on Ontology Learning and DEBRA below, are in turn predicated on mathematical, statistical and linguistic theoretical foundations, all of which come together to allow for the crystallization and synthesis of the information in the form of structured, semi-structured or unstructured data into data structures, called Knowledge Representation in the literature on AI, that not only facilitate the understanding, or meaning, embedded or encoded in the text itself but also allow for—and this is the thrust of Ontology Learning ultimately—the creation of applications and programs which adhere to the taxonomic, semantic and axiomatic structure which is gleaned from the textual corpus that is used as input to the ontology learning process.

The problem of knowledge representation in AI however, manifests as something very specific, as the problem of developing rational agents that are able to successfully navigate their environment effectively, and while solutions to this problem most certainly borrow techniques and tools from what might be considered classical software development or software engineering disciplines, in AI it nonetheless presents unique challenges because the representation of knowledge in this case must allow for the management and evolution of rules or axioms that provide the boundaries and scope for AI behavior, loosely speaking. These rules or axioms also must be tractable in the sense that they can grow and evolve at run time in a real environment, and these rules must be structured in such a way

<sup>2</sup>The research seminar at Dartmouth—Dartmouth Summer Research Project on Artificial Intelligence—that typically marks the origins of the field was held in 1956. See <https://home.dartmouth.edu/about/artificial-intelligence-ai-coined-dartmouth>

<sup>3</sup><https://futureoflife.org/open-letter/pause-giant-ai-experiments/>

that they lend themselves to both querying, in the sense of is this operation or this relationship possible or feasible, as well as reasoning, in the sense of what is the best possible action given the set of rules and axioms that govern the system in question as well as the data input it has from its current environment and context.

Knowledge Representation then, as a specific discipline within AI (and Computer Science), is focused on developing methods and techniques to both *represent* and structure knowledge in a way that computers can understand and at the same time provide capabilities which allow for the rational processing of said information, *i.e. reasoning*. How these types of systems are developed is a critical part of enabling all sorts of AI systems to perform tasks that require what we would consider being *intelligent* behavior, behavior such as problem-solving, deductive or inductive reasoning, decision making, and the perhaps most complicated task of natural language understanding which requires its own unique forms of knowledge representation (typically neural networks). It's fair to say that the field of Knowledge Representation (and Reasoning) is about nothing less than bridging the gap between the way humans think and communicate and the way computers process information.

The formulation of the field of knowledge representation as it relates to the Semantic Web, from which the discipline of Ontology Engineering/Learning has arisen, is driven primarily by the desire to transform the vast amount of human knowledge available in digital form into a structure that computers, machines, can effectively understand and utilize. Classic application development environments that had evolved to solve all sorts of data and transaction processing problems are not well suited for this task necessarily because 1) they require a lot of system and resource overhead to support both a run time environment as well as a database management system, and 2) the systems are typically rigidly architected, and proprietary, and are not easy to modify at run time in terms of what entities were available to the system and how they are related to each other. Effectively these transaction processing systems, despite their power and scalability, required too much hand-crafted design to evolve and adapt to different environments. A new type of rules-based system that has this type of adaptive capability is required and as such a new type of knowledge representation becomes necessary to facilitate its development.

Knowledge Representation as a field of study emerged as researchers realized that traditional programming approaches were not well-suited to capturing the complexity, uncertainty, and contextuality of human knowledge. As the need evolved and the technology requirements became more advanced, the field developed and was influenced by various other disciplines, which of course included logic as a branch of philosophy, but also cognitive psychology and linguistics, which also informed Deep Learning and NLP fields more broadly. Advancements across all of these interdependent fields have persistently informed the research and development into Knowledge Representation for AI practition-

ers and ultimately this is what the Semantic Web is trying to solve, *i.e.* the transformation of the digital assets of the Internet into something that AI, on a global scale, can understand and work with to solve all sorts of problems related to both automation and learning.

One might think then, given the level of integration of the pursuit of knowledge as a scientific endeavor within the field of Computer Science and AI, that the nature of knowledge is very well understood and mutually agreed upon. Certainly, while this may be true, relatively speaking, for the field of Logic from which the Reasoning part of Knowledge Representation and Reasoning rests fundamentally, the study of what knowledge is, what its boundaries are and how it is to be approached and understood, is not at all an agreed upon discipline, with arguments about the nature of knowledge again going back to the very roots of the Western philosophical tradition and persisting to present day.

The study of knowledge as a particular discipline and area of research in philosophy lies in some sense at the very core of the Western philosophical tradition, in what we now call *epistemology*, a term coined by the Scottish philosopher James Ferrier in the middle of the 19<sup>th</sup> century but which in fact comes from Aristotle himself, considered by many to be the father of Western philosophy. The word *epistemology*, meaning the “study of knowledge”, comes from the Greek word for knowledge *epistēmē*, which in turn has been handed down to us through the Latin *sciencia*, a derivation of the Latin verb *scire*, which means “to know”, which is the word that was used to translate the Greek word that was fundamental to Aristotle’s system of philosophy, establishing the framework for the study of knowledge itself.

What types of knowledge are there? How is it to be studied? What is the relationship between knowledge and existence or being? Answers to these questions drive not only significant portions of Aristotle’s system of philosophy (and Plato’s for that matter), but also arguably the bulk of the tradition of Western philosophy since Ancient Greece. And virtually all of this intellectual edifice rests upon how it is we conceive of, and categorize (*i.e.* read *ontology*) knowledge, *epistēmē* [4] [5] [6].

While these principles no doubt underpin Western philosophy, providing the very framework and language that the discipline even to this day operates under, the terminology, and the underlying principles along with it, has also bled into Computer Science as well, and more specifically and more pertinently with respect to this paper, the field of Artificial Intelligence (AI). In AI, we find specific allusions, as well as not surprisingly profound intellectual dependencies, in the research area that has come to be known as Knowledge Representation and Reasoning. In brief, it is this area, or branch, within AI that explores the nature of the structure of knowledge that is necessary for the development of intelligent agents, and it rests quite squarely on development in Logic, which is a field which provides much of the intellectual infrastructure for many areas of Computer Science, with perhaps Computing Theory (Turing, Church, Gödel, etc.) and AI

being the most prominent.

This paper concerns itself with the construction of conceptual structures, hierarchies ultimately, that can be gleaned from literary texts, a surprisingly underserved area of research due primarily to its lack of commercial value. The ability to extract conceptual hierarchies, forms of knowledge representation ultimately, from say Kant's *Critique of Pure Reason* is perhaps of interest to those in the philosophical community, but to the Computer Science community which is primarily focused on research that pushes forward the boundaries of computing at scale, both at the large end of the spectrum as well as the small (e.g. quantum computing), there is not much value gained from exploring the various representations of knowledge of certain types of literary texts than say the analysis of medical documents to facilitate the diagnosis of patients for example, the latter being a significant area of research and development not surprisingly.

What we do have however, given the interest in what is known as the Semantic Web, called Web 3.0 in some circles, is the exploration into tools and techniques as to how to construct what have come to be known as *ontologies*, organized conceptual hierarchical structures which are designed for the purpose of effectively solving the Knowledge Representation and Learning problem for the world wide web writ large [7] [8]. This word ontology of course also is borrowed from philosophy, the term denoting the inquiry into the nature of existence, or being, itself. While ontology as a branch of Philosophy is relatively recent, in the last century or two, the term itself, as with its counterpart *epistēmē* has great relevance and importance within the very heart of the Western philosophical tradition, the word in its root form *ontos*, which is derived from the present participle of the Greek verb “to be”, or *ὄντος*, meaning “being”, or “that which is”. This word is fundamental to both Aristotle and Plato's philosophical systems, as well as even the Pre-Socratic Parmenides [6] [9], and by inheritance within the Hellenic philosophical tradition from which Western science emerges post Enlightenment with the Scientific Revolution. For it is the study of *being*, more specifically what Aristotle calls the study of *being qua being*<sup>4</sup> around which his philosophical system, and in turn Western philosophy, itself is organized which of course from which Science itself ultimately emerges [4] [6].

In modern philosophy, *ontology* is that branch of metaphysics that deals with the study of existence, or again the nature of *being*, and while again its use in philosophical circles is a relatively recent, only introduced in the early 17<sup>th</sup> and 18<sup>th</sup> centuries, it has nonetheless come to form a specific branch of study within philosophy, and in particular metaphysics, itself in the 20<sup>th</sup> century.<sup>5</sup> In Com-

<sup>4</sup>“being qua being” is the English transliteration of the Greek phrase *ousia qua ousia*, a phrase Aristotle uses in his *Metaphysics* to describe his *theoretical philosophy* more or less, or what he calls *first philosophy* which we generally consider metaphysics from a more modern philosophical point of view. It rests on his definition of what *ousia*—ὀυσία in Greek—is, or more precisely, as the phrase *ousia qua ousia* is getting at, what the very nature of *ousia*, meaning “substance”, or “essence”, is. For the nature of substance, or essence, is not fundamental to Aristotle's metaphysics and epistemology, but also his ontology which he outlines in the *Categories* from which this notion of ontology and of course Ontology Learning, ultimately derives. See [4] [5] for more on this concept in Aristotle.



puter Science, the term has been co-opted within a branch of AI that stems from research into the development of the Semantic Web, with the underlying data structures that are requisite for its functioning having come to be known again as ontologies, and as such the discipline of Ontology Learning was born, the study of the creation of conceptual hierarchical structures from both structured unstructured and semi-structured data that facilitates both the knowledge representation as well as learning process associated with (more generally) the Semantic Web and (more specifically) particular intellectual domains in and of themselves [7] [8].

*Learning*, as understood more broadly within AI, fundamentally rests upon our ability to represent knowledge in way that machines can understand it, and as such work with it, process it, and generate something that looks intelligent on the other side. This notion, however obvious it might appear to those in the AI field, nonetheless sits at the very heart of the AI revolution, and more specifically sits at the very heart of modern Machine Learning (ML) algorithms that lay at the heart of Large Language Models, like ChatGPT from OpenAI<sup>6</sup> for example. These algorithmic advancements, which again rely on these various developments in Knowledge Representation, for example Vector Space Models or Word Vectors [1] or more generally neural networks, when married to the availability of computing resources at scale along, technological advancements that fall under the umbrella of Deep Learning more generally, are in fact what is driving the advancements in AI which are transforming the modern technological landscape.

It is the “reasoning” part of Knowledge Representation and Reasoning that is more specific to AI, reflecting the need in the AI space specifically for the ability for intelligent agents to respond to their environment in an “intelligent” way, and by intelligent we mean of course specifically human intelligence. In other words, we look to develop intelligent agents that behave like humans, and one of the fundamental characteristics, from a systems and data structure perspective, that facilitate this “reasoning” is the creation of, at least within the context of the Semantic Web, of ontological structures, *i.e. ontologies*, which encode conceptual relationships as well as logical rules which altogether facilitate solutions to the problem of reasoning, specific to the Semantic Web yes but applicable to problems of AI reasoning more generally.

What we present here with DEBRA effectively is a form of lightweight ontology except our focus is on single, literary texts of cultural significance with the intent of deriving meaning, or understanding from the text itself that is not necessarily rooted in a fully logical or rational basis for the intent of machine un-

---

<sup>5</sup>The first reference to the term *ontology* as a branch of metaphysics can be traced back to 1606, in the Latin as *ontologia*, in Jacob Lorhard’s work on Christian and Scholastic theology (and philosophy) entitled *Ogdoas scholastica* [10], but it really takes root in the Western philosophical tradition with the work of the German rationalist Christian Wolff in the early 18<sup>th</sup> century who published *Philosophia Prima sive Ontologia*, or *First Philosophy or Ontology*, in 1730. See <https://www.ontology.co/history.htm>.

<sup>6</sup>See <https://openai.com/>



derstanding or processing, in search for a meaning, or more specifically a form of knowledge representation through which meaning can be derived—beyond the purely rational one that underpins the technologies that underpin the Semantic Web (ontologies), or more generally AI (first order logic), and by doing so open up a new field of research in the humanities that leverages AI, ML and NLP techniques for the analysis of philosophical and theological literature so that these disciplines can leverage the output of such tools for a more sound mathematical and statistical grounding that at the very least can complement the work of researchers in these fields.

### 3. On the Distinction between Intuitive & Analytical Reasoning

Before *ontology* came into vogue as a philosophical discipline, the two core pillars of Western philosophical inquiry, rooted in the discipline as it was formulated under Aristotle, were *metaphysics* and *epistemology*—the former denoting the study of the theoretical foundations of philosophical inquiry itself, what Aristotle called *first philosophy* or the *theoretical sciences*, and the latter being the inquiry into the nature and boundaries of knowledge itself [11]. The importance of these two areas of research in the history of Western philosophical development, as well as Aristotle’s influence on the definition and scope of these fundamental areas of philosophy, cannot be overstated<sup>7</sup>.

What is typically lost, or left discounted, in this intellectual development cycle, one which culminates in the establishment of the theoretical foundations of logic and computing in the first half of the twentieth century upon which the research in this paper is predicated, is the distinction that is made at the very root of epistemological inquiry into what we will call here *intuitive* versus *analytic* reasoning. Intuitive reasoning involves arriving at conclusions or insights without explicit and systematic reasoning. It relies on instinct, feelings, gut reactions, and subconscious processes generally speaking but it also has a relationship to what some scholars refer to as the paranormal, very much akin to what scientists who studied Quantum Mechanics in the first half of the twentieth century “spooky action at a distance”, or “hidden variables”, or the “implicit order”. It often occurs spontaneously and is based on accumulated experiences and implicit knowledge. On the other hand, analytic reasoning involves a deliberate and systematic approach to problem-solving and understanding. It relies on logical thinking, deductive or inductive reasoning, and critical analysis of evidence. Analytic reasoning aims to

<sup>7</sup>The very words we use to describe the disciplines themselves, *epistemology*, and *metaphysics* specifically, but also *ontology* as well as even the word *science* itself—which is the English translation of the Latin word meaning the same *sciencia*, which was the Latin translation of *epistēmē* which was of course was one of the cornerstones of Aristotle’s entire system of philosophy) – are rooted in the very core of the Hellenic philosophical tradition as established by Plato and Aristotle and others in classical Greece. Even Kant’s Critiques, considered by many to be the very height of Western philosophy, in particular his *Critique of Pure Reason* which he wrote toward the end of the 18<sup>th</sup> century, paid homage to the rational and metaphysical foundations that were laid down by Aristotle some two thousand years prior [12] [13].

minimize biases and errors by following a structured and methodical process.

These are two very distinct ways of knowledge acquisition that were certainly recognized at the very root of philosophical inquiry, arguably a good cognitive science representation of Aristotle's substantial form, but yet the intuitive part of reasoning, of understanding, is left out entirely from the discussion of "reasoning" in AI, and yet with modern language models which have grown so powerful and "intelligent", there is arguably an element of intuitive understanding that these models bring to their discourse despite it not having been programmed for such a thing. It's almost as if the deep learning methods that are used via leveraging various neural networking technology and algorithms almost resolves to a state of intuitive understanding from its purely rational—mostly mathematical and algebraic—foundations.

However, one of the fundamental problems of metaphysics as understood by Kant at least, perhaps the most influential philosopher in the Western tradition after Aristotle, is the very idea of what metaphysics is fundamentally, and what—as a byproduct of this question—could be considered *a priori knowledge* independent of "objective reality" which should provide the basis for any discussion, according to Kant at least, about metaphysics itself. Hence the title of his seminal work on metaphysics, *Critique of Pure Reason*. While this problem in and of itself may seem quite far removed from the notion of ontology engineering and computer science, it nonetheless sits—even still—at the heart of epistemological questions to which this idea of Knowledge Representation (and Reasoning) itself at some level must be constructed upon.

The answer to this question, as explained by Kant himself and what he presents in his Critique and refers to (in humble fashion) as a Copernican revolution of philosophy, is what ultimately we understand as, as he himself refers to his position as, *transcendental idealism* [14] [15], a sort of intellectual middle ground between the idealist and the materialist epistemological positions. In his Preface to the Second Edition in 1787, he describes the problem, and his insight upon which he believes to have solved it, thusly:

*It has hitherto been assumed that our cognition must conform to the objects, but all attempts to ascertain anything about these objects à priori, by means of conceptions, and thus to extend the range of our knowledge, have been rendered abortive by this assumption. Let us then make the experiment whether we may not be more successful in metaphysics, if we assume that the objects must conform to our cognition. This appears, at all events, to accord better with the possibility of our gaining the end we have in view, that is to say, of arriving at the cognition of objects à priori, of determining something with respect to these objects, before they are given to us.*

*We here propose to do just what COPERNICUS did in attempting to explain the celestial movements. When he found that he could make no progress by assuming that all the heavenly bodies revolved round the spectator, he reversed the process, and tried the experiment of assuming that the spectator revolved, while*

*the stars remained at rest. We may make the same experiment with regard to the intuition of objects. If the intuition must conform to the nature of the objects, I do not see how we can know anything of them à priori. If, on the other hand, the object conforms to the nature of our faculty of intuition, I can then easily conceive the possibility of such an à priori knowledge.*

*Now as I cannot rest in the mere intuitions, but—if they are to become cognitions—must refer them, as representations, to something, as object, and must determine the latter by means of the former, here again there are two courses open to me. Either, first, I may assume that the conceptions, by which I effect this determination, conform to the object—and in this case I am reduced to the same perplexity as before, or secondly, I may assume that the objects, or, which is the same thing, that experience, in which alone as given objects they are cognized, conform to my conceptions—and then I am at no loss how to proceed.*

*For experience itself is a mode of cognition which requires understanding. Before objects, are given to me, that is, à priori, I must presuppose in myself laws of the understanding which are expressed in conceptions à priori. To these conceptions, then, all the objects of experience must necessarily conform. Now there are objects which reason thinks, and that necessarily, but which cannot be given in experience, or, at least, cannot be given so as reason thinks them. The attempt to think these objects will hereafter furnish an excellent test of the new method of thought which we have adopted, and which is based on the principle that we only cognize in things à priori that which we ourselves place in them<sup>8</sup>.*

The German word which Kant uses which is translated into English as “intuition”—a fundamental concept in Kant’s epistemology, particularly in his discussion of the distinction between *sensibility* and *understanding*, two of the core mental faculties which facilitate the process of understanding in his system, is *anschauung*, which is more often than not translated as “intuition” but given Kant’s use of the word throughout his work and its important place in it, as well as Kant’s influence on the post Enlightenment Western philosophical tradition which cannot be overstated, the word has also taken on a more subtle, and quite specific definition within philosophy that can be interpreted as something like “*an element of knowledge that is directly given in sense awareness that is perceived a priori, that is to say prior to, the apprehension, or in turn cognition or understanding, of objects of sense perception.*” This is certainly (more or less) what Kant intends it mean, and the implications for metaphysical inquiry, presuming he’s right that objects do conform to our cognition and not the other way around, are indeed no less than Copernican in their revolutionary impact on how it is that we can come to understand anything, which in turn effects quite directly how it is we derive meaning from anything we experience or think about, which collectively can be conceived of as “objects of our understanding” [14] [15].

From an epistemological standpoint then, at least according to Kant, one

<sup>8</sup>*Critique of Pure Reason*, by Immanuel Kant. Second Edition 1787. Translated by J. M. D. Meiklejohn 2018. From the Preface to Second Edition (1787).

could say that *understanding* is a sort of marriage between *cognition*, as a mental faculty, and *objective reality* that is perceived through the senses through experience, through an individual's (a subject's) perception of the world, which is the container so to speak of objective reality. This is Kant's metaphysics in a nutshell, which embeds his epistemological position which is not entirely "subjective" but most certainly not entirely "objective" either. Importantly, Kant does not see this is not an (entirely) rational process, it's something sort of *super-rational* and its proper functioning rests upon this faculty of *intuition*, which is again inherent to the process of *understanding*, or again *verstand* in the German<sup>9</sup>. This is the middle way that Kant constructs to reconcile the idealist and the materialist positions, which in his day were called the rationalists and the empiricists respectively [14] [15].

These two different, and yet complimentary—or perhaps better put, *interdependent*—modes, or again faculties, of the understanding can be seen as representing two seemingly disparate intellectual poles of how it is that we come to any sort of understanding of the world itself, not as it truly is necessarily but as we know it to be, a slight and yet very important distinction in Kant's metaphysics. From Kant's perspective then, it's fair to say that it is these two complementary mental cognitive capacities come together to yield knowledge, as a function of both subjective experience and objective reality as they come together in the (faculty) of the human mind.

There is a philosophical analogue here, but we must move East, and further back in antiquity, in order to find a similar notion of distinct and yet interrelated aspects, or principles, which come together to form an *undivided whole* [17]. Arguably this holistic perspective on not just experience but metaphysically, and ontologically, is fundamental to Eastern philosophy write large, in particular to the philosophical systems that are predominant in ancient India, namely Vedanta and more broadly Upanishadic philosophy, as well as in ancient China, reflected in Daoist (Taoist) thought primarily [12] (Figure 1).

As we discuss in at length in our work on Eurasian Philosophy, this epistemological divide—reductionism versus holism put simply—is a loose corollary to the comparison between Western philosophy, which is rational, logical and reductive by nature primarily, and Eastern philosophy, which at its roots at least, is just the opposite, a holistic view of the world as seen through the eyes of the subject and reflected in the world around them. This notion of the interconnected whole, or *undivided wholeness* as the great 20<sup>th</sup> century (Quantum) physicist David Bohm calls it [17], is encoded in the famed Taiji (*tàiji*) as shown in the figure above, the metaphysical foundations for what we know in the West as

<sup>9</sup>In the *Critique of Pure Reason*, Kant juxtaposes the faculty of the *understanding* (*verstand*)—which is responsible for concepts, judgments, and logical reasoning and plays a central role in organizing and categorizing sensory data to form coherent knowledge, with the faculty of *reason* itself, or what he calls *vernunft*, which is typically translated into English as "reason", but as Kant uses the word alludes to a higher order intellectual faculty which supersedes, and sits on top of metaphysically so to speak, of the two fundamental aspects of cognition which are central to Kant's metaphysics, namely *understanding* and *sensibility* [16].

*Yin-Yang philosophy*, referred to in the *Ten Wings*, the philosophical compendium to the *I Ching*, or *Classic of Changes (Yijing)* as emergent from “the Great Primal Beginning”.

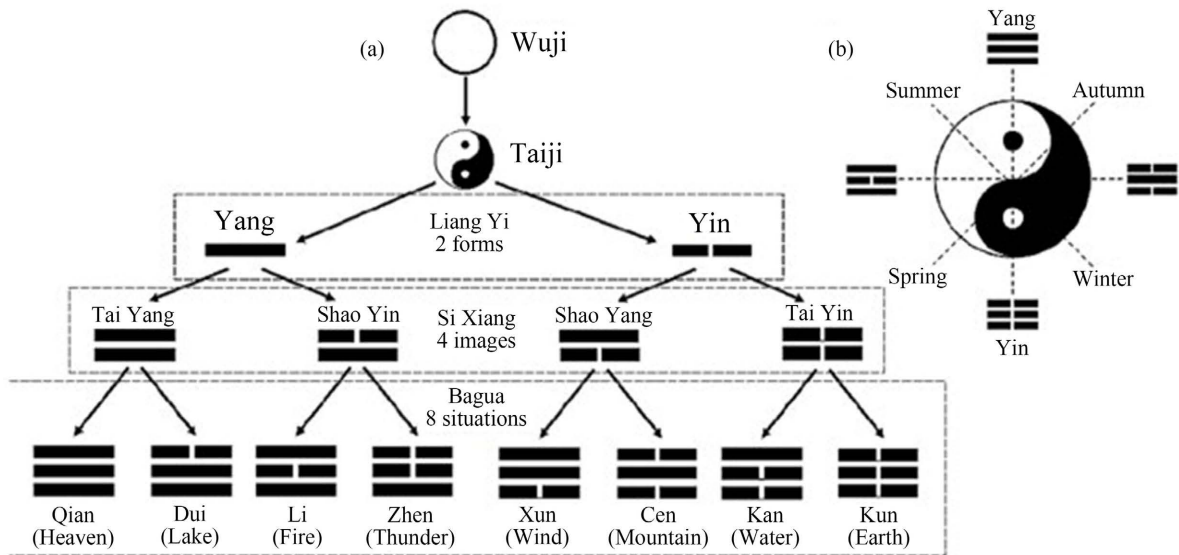
This yin-yang symbol, which is commonly known in the West and typically is associated with Taoism (Daoism), or the Way, *i.e.* the Dao, as a symbolic representation of the philosophy of Lau Tzu, as extant in a text called *Tao Te Ching (Daodejing)*, or the “Way of Virtue”, is from a metaphysical perspective referred to as the *tàijí*, or “great pole”, distinguishing it from the *wújí*, or supreme ultimate (“without ridgepole”), and represents the initial state of differentiated creation from which emerge the core components of the *I Ching* itself – the yin and the yang from which we derive the “four images” from which in turn we derive the “eight trigrams”, or *bāguà*, from which the entire (metaphysical) construction of the 64 symbols of the *I Ching* are formed. This emanation of the many from the one, the *Taiji* from the *Wuji*, is illustrated in the figure below, which also shows the 8 primary bagua from which the *I Ching* symbols are constructed and that make up the core, naturalist and holistic, Daoist philosophy [12] (Figure 2).

In the Indian philosophical tradition (specifically Vedanta which looks to the *Vedas*, the *Upanishads*, and the *Bhagavad Gita* as the primary sources of truth), this fundamental epistemological distinction is called out specifically in the Mundaka Upanishad as the difference between higher, undifferentiated knowledge of Self or Brahman, or *parāvidyā*, and the knowledge of the material, physical world which is regarded as a lower form of knowledge or *aparāvidyā* [18]. Aristotle makes a similar distinction in his corpus, where he distinguishes the basic pursuits of knowledge, or sciences (*epistēmē*), as the pursuit of knowledge for its own sake as belonging to what he calls the *theoretical sciences*, as distinct from the pursuit of productive knowledge in the form of the study of rhetoric, the arts, etc., and the pursuit of and practical knowledge such as ethics, or politics [13].



**Figure 1.** *tàijí*, or “great pole” of Chinese Philosophy<sup>10</sup>.

<sup>10</sup>Image by Klem—This vector image was created with Inkscape by Klem, and then manually edited by Mnmazur, Public Domain, <https://commons.wikimedia.org/w/index.php?curid=3213322>



**Figure 2.** Wuji to Taiji to bagua to eight trigrams (Image from [20]).

This distinction is also drawn in the Chinese philosophical distinction as we note above as the study and knowledge of the Way itself, *i.e.* the Dao, as opposed to the knowledge of, or the existence of, the world as denoted by the term *wan wu*, literally translated as “ten-thousand things” but denoting the myriad and endless nature of physical reality or existence [6] [19]. We see this notion of the emanation of the many from the One, a prominent feature of Christian Gnosticism as well as Neo-Platonic philosophy, in one of the most famous verses of the Tao Te Ching (Verse 42) which reads:

*The Dao produced One, One produced Two,  
Two produced Three, Three produced All things (wan wu).  
All things leave behind them the Obscurity out of which they have come, and go forward to embrace the Brightness into which they have emerged, while they are harmonized by the Breath of Vacancy.  
What men dislike is to be orphans, to have little virtue, to be as carriages without naves, And yet these are the designations which kings and princes use for themselves.  
So it is that some things are increased by being diminished, and others are diminished by being increased.  
What other men thus teach, I also teach.  
The violent and strong do not die their natural death.  
I will make this the basis of my teaching<sup>11</sup>.*

We call attention to this here given that virtually all of the root philosophical systems not only call out this distinction of epistemological perspectives, but that in the Western philosophical system, the intellectual edifice on top of which Ontology Learning and more specifically DEBRA as we present here rests, not just discounts the importance of the whole being just a sum of its parts, but that

<sup>11</sup>*Tao Te Ching (Daodejing)*, or *Way of Virtue*. Translated by James Legge, 1891, Chapter 42. Multiple translations available online at <https://www.egreenway.com/taoism/ttclz42.htm>.



the perspective of the whole is not only distinct from the perspective of the many, but that it in turn provides valuable and unique insight into the understanding of the thing in question. These two vantage points can be looked upon as analogies of the two different types of understanding—namely intuitive (holistic) and analytical (reductive).

What's been lost with the advent of more and more detailed and rationalistic Western philosophy is this notion of the many emerging from, and ultimately indistinct from, the one, again this idea of the undivided whole from which the basic bifurcation of all knowledge stems from. It is this lost understanding one could say, which represents the main problem with classical, orthodox, interpretations of Quantum Theory, from which the “quantum spookiness”, as Einstein referred to it as, emerges. But a metaphysical turn back to the whole, what Bohm and Hiley call the *Ontological interpretation* of quantum mechanics [17], provides a perspective that allows for an emergence of the physical from a “meta” physical, a higher order reality that embeds this interconnectedness.

Of course, it is the *I Ching (Yijing)*, and the underlying system of knowledge that it represents (which underpins all Chinese philosophy essentially) is the very same system that Leibniz studied [21], which in turn inspired the binary system upon which all computing is based. What's been left out from this knowledge transfer as it was adopted in Computer Science however, is the fundamental relation of the whole with its constituent parts, the emergence of the Two from the One, the 0 and 1 as they relate to the very idea of a *bit*, or binary digit, the understanding of the Two as a representation of the One in a dual, opposing form<sup>12</sup>.

This is, at some level at least, what we are attempting to recover in our development of conceptual hierarchies from literary texts with DEBRA—a sort of intuitive, of wholistic understanding of a text by means of an extrapolation of concepts and their relationships from said text as a representation of the whole, rather than (exclusively) a reductionist sum of its parts and their relations. The two forms of representation should complement each other.

#### 4. On the History of Concept Maps, Ontologies & DEBRA

Now we are in a position to describe the field of Knowledge Representation (KR), or alternatively Knowledge Representation and Reasoning (KRR) within AI and how it is that DEBRA both fits into this discipline as well as distinguishes itself from current, more orthodox areas of research within this area.

KRR is primarily focused on the development of formal methods and structures for representing and organizing knowledge in a way that computers can understand and manipulate. It's an essential aspect of AI systems that enables


<sup>12</sup>In some sense this is what quantum computing represents in relation to the fundamentally dualistic approach that has been characteristic of computing since its inception with the transistor and bits, or binary digits. In quantum computing there exists the notion of a third state, or what's called the superposition state, which represents the idea that the value of the “qubit” (or quantum bit) could be either 0, 1, or either (or both really). This is the fundamental distinction between classical computing and quantum computing and the reason why it represents a such a significant breakthrough from a computing and processing power perspective.



them to reason, make inferences, and solve complex problems, and it is from within this field of study that Ontology Learning, or more technically Ontology Engineering, emerges as a sort of best practice in the field, or as a standard terminology within which KRR is typically taught and understood [22] [23].

As a field of study, KR can be traced back to the early days of AI research as in order to build machines that “reason”, which ultimately is the goal of AI after all, and therefore it was necessary to develop data structures to represent logical and symbolic information which could be effectively searched and leveraged to facilitate “machine intelligence”. Early work included the use of formal languages and rules to represent knowledge, which included the introduction of Frames by Marvin Minsky in 1975, allowing for the representation of structured information by organizing knowledge into hierarchical categories and slots for attributes [24].

In the 1980s, Semantic Networks, gained popularity, emphasizing the representation of knowledge using nodes and links to express relationships between concepts [25], supplanted in the 1990s and 2000s with the rise of ontologies and the Semantic Web initiative. Ontologies provide a formal way to define concepts, relationships, and properties in a domain, enabling machines to understand and reason about the semantics of information [7] [8]. In the last two decades, advances in machine learning and deep learning have influenced the field of KR as researchers have explored methods to integrate probabilistic reasoning, uncertainty handling, and neural network approaches into knowledge representation systems [23]. **Figure 3** depicts the diverse landscape when it comes to Knowledge Representation.



**Figure 3.** Forms of Knowledge Representation in AI.

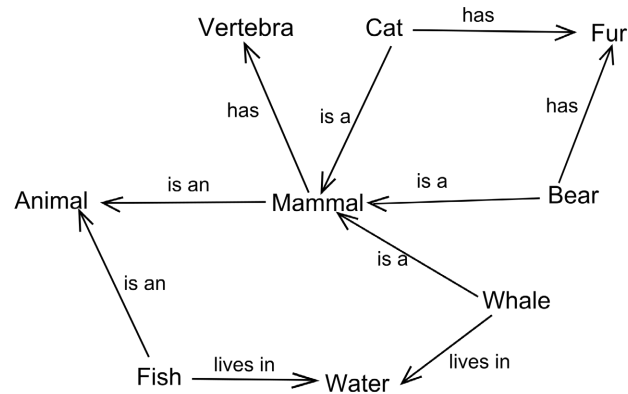
In particular with respect to DEBRA, as an unsupervised method of extracting concept hierarchies from text, *semantic networks* hold a special significance. Semantic Networks are a graphical form of knowledge representation which represents entities as nodes connected by directed edges which represent relationships between the entities (nodes). Semantic networks are often used to represent hierarchical or associative relationships between concepts, making it easier to understand the connections between different pieces of information. They are relatively simple structures and can be useful for organizing and visualizing knowledge without formal specification requirements, loosely corresponding to the notion of intuitive understanding from a philosophical context that we describe above. An example of a basic semantic network representing mammals is depicted in **Figure 4**.

Semantic networks have roots that trace back to various fields, including psychology and linguistics along with AI. One of the earliest and notable developments in this area is the work of Ross Quillian in the 1960s. In 1968, Ross Quillian [26], a cognitive psychologist, introduced a model called the “Semantic Memory Model,” which is considered one of the pioneering works in the field of semantic networks. Quillian’s model aimed to represent human semantic memory and how concepts are organized in the human mind. He proposed the use of hierarchical networks to represent semantic relationships between concepts, introducing the idea of nodes representing concepts and links representing the relationships between them, including “is-a” relationships and other forms of associative links, laying the foundations for subsequent research in the field of artificial intelligence, knowledge representation, as well as cognitive science from which it was born.

Having said that, while the “discovery” of semantic networks is a twentieth century phenomenon, precursors to this mode of intellectual architecture can be found at the very root of the western philosophical tree, alongside logic as conceived of by Aristotle in fact, as reflected in his work *Categories* which was one of the treatises on logic that was included in the *Organon*. The first example of this type of knowledge structure in fact can be found as a compendium to the Porphyry’s highly influential introduction to Aristotle’s *Categories* called the *Isagoge*, or “Introduction”. This work was translated into the Latin in the 6<sup>th</sup> century by Boethius, who added the tree (shown below) which although was not created by him, or Porphyry for that matter, was nonetheless unique in that it was the first example of a visual representation of the metaphysical, or ontological, structure which was expressed in Aristotle’s seminal work on ontology, *i.e. Categories*. The left side is the image crafted by Boethius presumably, with an English translation just beside it (the right side). For details, please find the following link:

[https://en.m.wikipedia.org/wiki/File:Arbor\\_porphyrii\\_%28probably\\_from\\_one\\_of\\_Boethius%27\\_translations%29.png](https://en.m.wikipedia.org/wiki/File:Arbor_porphyrii_%28probably_from_one_of_Boethius%27_translations%29.png)

In contrast to ontologies, as they are formally defined within the context of AI and the Semantic Web at least, represents a much more formally defined and structured form of knowledge representation than semantic networks, or DEBRA



**Figure 4.** Example of Semantic Network of Mammals<sup>13</sup>.

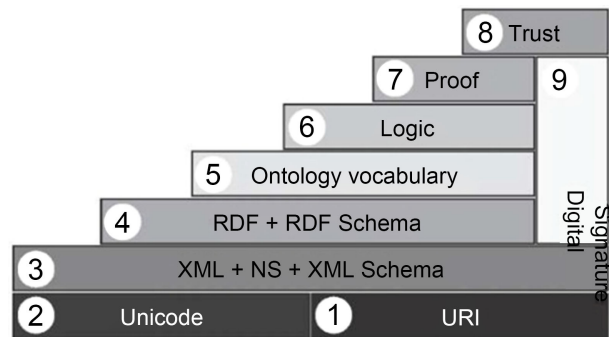
for that matter, reflecting the more analytical, and logical, aspects of knowledge as we describe in the philosophical introduction to epistemology above.

While ontologies do, in fact, define concepts and their relations in a hierarchical structure, as we do with DEBRA in some sense, ontologies structure this hierarchical relationship not only more formally, but in a pre-defined way such that the hierarchy goes from the top down from the more abstract to the more detailed, whereas DEBRA, for example, constructs its hierarchy as a function of the concepts that are inherent in the text that is being analyzed, using the text itself as the ontological framework rather than the ontology which is a) domain-specific and b) is typically defined outside of, or in relation to, the textual corpus that is being analyzed. Furthermore, ontologies, given their fundamentally logical architecture, provides the capability to embed and encode not only the details of specific relationships between concepts in the ontological structure, but also rules and axioms that define both the boundaries and behavior of the concepts and their relations with respect to the underlying ontological structure.

Ontologies also have the capability to capture and store additional information such as attributes, properties, constraints, axioms, and rules that specifically define the behavior, *i.e.* semantic meaning, of the concepts and their respective relations. The Semantic Web in this context can be seen as a standard format(s) of data description that facilitate the creation of applications to process said data, leveraging open-source standards like XML, OWL and RDF rather than proprietary RDBMS which are licensed through companies like Oracle and Microsoft which dominate the enterprise information processing landscape. A depiction of the Semantic web architecture is shown in **Figure 5**.

So ontologies not only encode concepts and their relations, but also linguistic (really logical or semantic) rules about how these concepts can be related to each other, allowing for both a conceptual description of a given domain as well as a semantic and logical description, which then in turn can be used to inform the structure of applications. This is the special sauce of ontologies, and ultimately the Semantic Web in the Berners-Lee vision.

<sup>13</sup>Source: Public Domain, <https://commons.wikimedia.org/w/index.php?curid=1353062>



**Figure 5.** Semantic Web layered architecture<sup>14</sup>.

And while these formal types of Knowledge Representation structures do in fact share many similarities, ultimately from a logical (and more broadly philosophical) perspective, it is worth noting that each of these types of knowledge representation structures can potentially be resolved to systems of first order logic, *i.e.* predicate calculus, albeit perhaps clumsily in many cases hence the development of these alternative forms of representation each with its own strengths and weaknesses and each with its own unique areas of applicability. From the perspective of first order logic however, as a sort of intellectual denominator for these various formal forms of knowledge representation, as an intellectual baseline for establishing the ground rules of concepts and their relations in a given domain, hence the focus on first order logic in AI (versus other forms of logic or ontology) from a graduate and teaching perspective [23].

While this is most certainly not a complete list of all the various forms of Knowledge Representation, and of course to solve various problems in the field of AI researchers have developed hybrid approaches that combine different representation techniques, it should be clear that this is a very unique problem to the domain in question (AI) and while DEBRA doesn't necessarily fit neatly into any of these categories (with Conceptual Graphs being the closest fit arguably), it is necessary to understand the area of research to contextualize DEBRA from a Knowledge Representation perspective.

Related to semantic networks are *concept graphs*, or *concept maps*<sup>15</sup>, which were first introduced by John Sowa in the 1970s as a way of describing the relations between entities in database systems [27], and then more broadly applied toward the mapping of ideological structures to machine processing problems in the 1980s [28]. Concept maps are designed to show how different concepts are connected and how they contribute to a broader understanding of a topic, and are a more generic, and more specifically less formal, representation of the relationship between concepts in a given domain than semantic networks or ontologies. An example concept map for electricity is shown in **Figure 6** below to illustrate the visualization component of concept maps, or hierarchies.

<sup>14</sup><http://what-when-how.com/information-science-and-technology/owl-web-ontology-language-information-science/>

<sup>15</sup>We use "concept graphs" and "concept maps" interchangeably here.

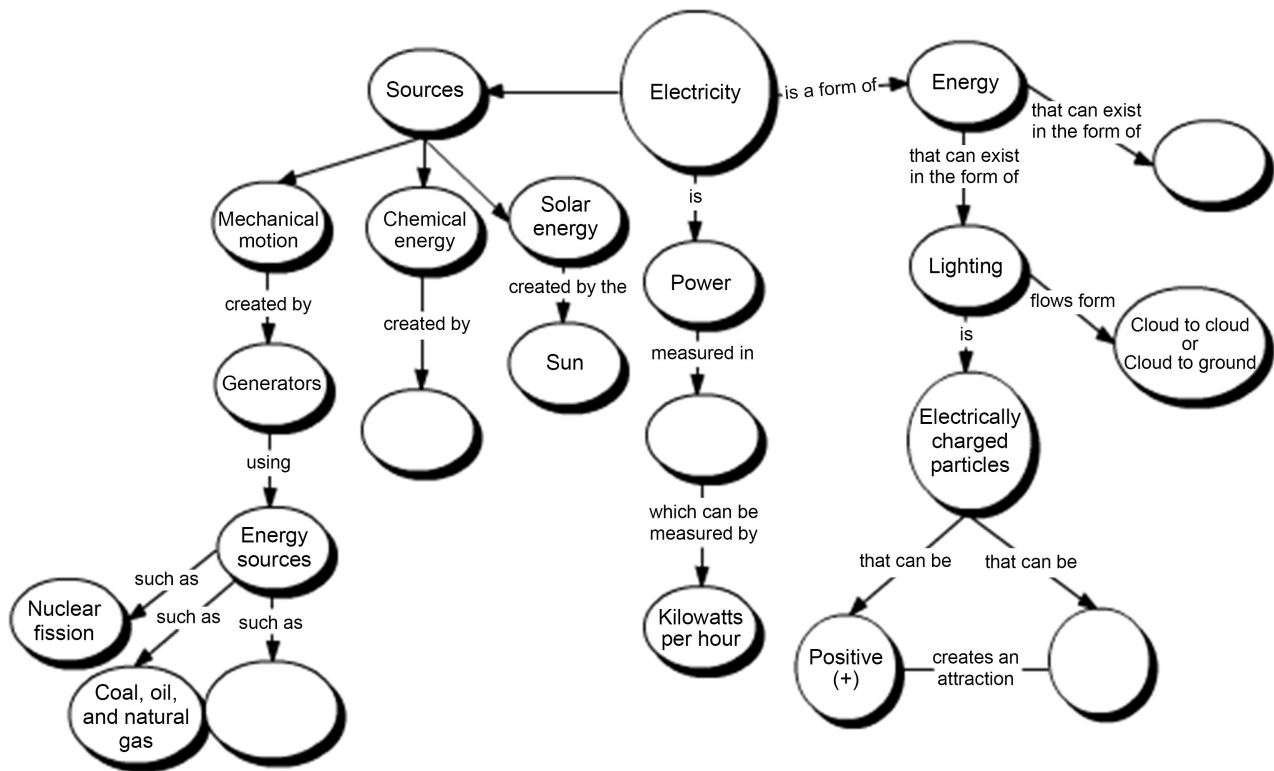


Figure 6. Concept map for electricity<sup>16</sup>.

A concept graph typically consists of nodes, representing concepts or ideas, and labeled arrows or lines which represent the connections, or relationships, between and among the conceptual nodes. With a concept graph, or map, the direction of arrows between concepts conveys the nature of the relationships between said concepts—such as cause-and-effect, part-whole, or sequential relationships. Concept maps are used in various fields such as education, business, and science to concisely convey information, *i.e.* as a form of knowledge representation, serving as a powerful tool for organizing, clarifying, and communicating knowledge across a wide range of fields, ultimately enhancing understanding and facilitating effective learning, communication, and problem-solving.

From one of the most preeminent scholars of concept maps and their relation to both meaning as well as understanding, we have the following quotation which gets to the very heart of the inter relationship of these ideas, both from a knowledge representation perspective as well as from a pure epistemological and ultimately cognitive science perspective as well.

*...we defined concept as a perceived regularity (or pattern) in events or objects, or records of events or objects, designated by label. It is coming to be generally recognized now that the meaningful learning processes described above are the same processes used by scientists and mathematicians, or experts in any discipline, to construct new knowledge. In fact, Novak has argued that new*

<sup>16</sup>Map from

[http://scifiles.larc.nasa.gov/text/educators/activities/2000\\_2001/worksheets/elec\\_concept.html](http://scifiles.larc.nasa.gov/text/educators/activities/2000_2001/worksheets/elec_concept.html)

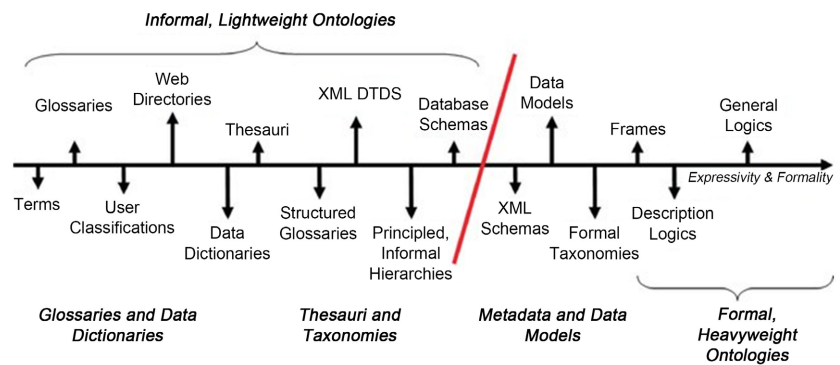
*knowledge creation is nothing more than a relatively high level of meaningful learning accomplished by individuals who have a well organized knowledge structure in the particular area of knowledge, and also a strong emotional commitment to persist in finding new meanings ([29] [30] [31]). Epistemology is that branch of philosophy that deals with the nature of knowledge and new knowledge creation. There is an important relationship between the psychology of learning, as we understand it today, and the growing consensus among philosophers and epistemologists that new knowledge creation is a constructive process involving both our knowledge and our emotions or the drive to create new meanings and new ways to represent these meanings. Learners struggling to create good concept maps are themselves engaged in a creative process, and this can be challenging, especially to learners who have spent most of their life learning by rote. Rote learning contributes very little at best to our knowledge structures, and therefore cannot underlie creative thinking or novel problem solving [32].*

As we can see from the quotation above from Novak and Cañas's work on Concept Maps [32], Novak being the inventor of concept maps as a tool to understand the learning process in general, there is presumed to be a close relationship between the acquisition of knowledge, understood more generally from a philosophical perspective as *epistemological inquiry*, and concepts along with their inter-dependent and inter-connected relations, ultimately, in a transcendental idealistic fashion following Kant in his Critique—mediated by cognitive and rational aspects of the persona, or psyche.

This epistemological structure laid out by Novak is in fact reminiscent of Kant's epistemology, which posits that our understanding of objects, or the world, conforms to our perceptions and mental faculties rather than existing in and of itself independent of such perception which is typically how the world is viewed from a physicalist, or empiricist or materialist, perspective. In other words, it is fair to say that the epistemological underpinnings of concept maps, can be understood as a sort of cognitive psychological interpretation of Kant's transcendental idealism, the two branches of knowledge themselves converging on the same truth as it were.

Also, as you can see illustrated in the figure above, there is typically an implicit hierarchical structure that is intrinsic to the concept graph. So while this structure doesn't embed formal semantic or logical information within the graph itself explicitly, it nonetheless shares the basic structure of an ontology with respect to the delineation of concepts and their relationships in a hierarchical, *i.e.* ontological<sup>17</sup>, framework. In this sense concept maps can be conceived of as initial step in the creation of more formal ontological structures, which is more or less the idea put forward with respect to *lightweight ontologies*, the spectrum of ontological structures laid out quite nicely by [33] which we reproduce here (**Figure 7**).

<sup>17</sup>This is the other meaning of "ontological", *i.e.* an adjective which denotes that which it describes as a hierarchical conceptual structure.



**Figure 7.** Types of ontologies [33].

Generally, we can see from this spectrum of ontologies illustrated above that first and foremost not all ontologies are created equal. Secondly, we can see that there are more formal variants which include the kind of semantic and axiomatic information that is required for ontological structures that underpin the Semantic Web, structures which embed logical statements with respect to the concept in the domain in question, again similar to what we might find in an Entity-Relationship diagram for a data processing application except with a different syntax and method of representation.

On the less formal side of the spectrum however, we do see a sort of conglomeration of different types of structures which, while providing both explanatory and informational type information related to the conceptual structure of a given domain, nonetheless lack the formal logical structure that we would need for example to support the development of intelligent agents in a classic AI scenario. The concept maps we generate from DEBRA sit somewhere on the dividing line although they nonetheless represent something distinct from these types of rational, analytical structures, with a purpose that is fundamentally different—namely the derivation of meaning for the support of *human* understanding from a given textual corpus rather than designed to support *machine* understanding.

So while concept maps are not designed to provide formal, logical semantic information related to concepts necessarily, such as again semantic networks or ontologies proper, or even more fundamentally first order logic knowledge bases—they do however serve a valuable purpose in facilitating the construction of an intellectual baseline for the domain in question. As illustrated with DEBRA, these concept maps then in a very real way encode, or encapsulate, a conceptual structure that while isn't well suited to AI applications per se, can be very valuable in defining the intellectual terrain for a given domain before, or orthogonal, to the development of a more comprehensive and rationally descriptive ontology.

## 5. On DEBRA's Modular System Architecture

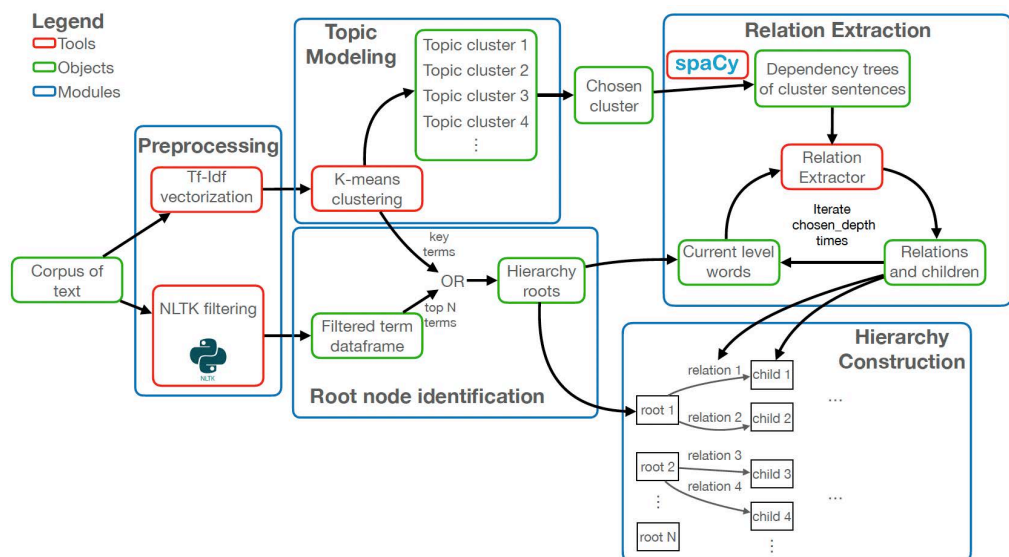
Work on DEBRA, short for dependency tree based concept hierarchy constructor, originated from an experiment to see what kind of concept hierarchies we could elicit from literary texts, specifically theological and philosophical works. When we looked at how best to do this, the most well-formed and documented



path appeared, initially at least, in the Ontology Learning domain. But while that literature is relatively mature—some two decades at least since that research had begun with at least two major texts on the subject [8] [34]—we were looking to do something different, and perhaps most importantly we were evaluating the solution algorithm against a different set of input data, a single literary text as opposed to a disparate source of structured and unstructured data which is typically the case with Ontology Learning generally speaking.

The proof of concept that we developed was built with Python, leveraging the latest in AI and NLP libraries that are readily available with this platform. The core part of the algorithm, the piece that looks at the text and builds the (lightweight) ontology, *i.e.* conceptual hierarchy, is essentially made up of four components, *pre-processing*, *topic modeling*, *relation extraction* and *visualization*, each described in turn, the full architecture of which is depicted in **Figure 8** below.

We have as illustrated in the architectural depiction of DEBRA above, a basic modular architecture which is reflected in the underlying Python code as well<sup>18</sup>. The paradigm follows the standard NLP approach of first pre-processing the text, which is a function of the format of the underlying text—the KJV comes as a set of verses in a CSV file for example whereas the Eurasian Philosophy text we worked with is a raw txt file—as well as a function of what stop words we want to be removed from the analysis, the removal of punctuation and capitalization, etc., all of which is designed to transform the text itself from a literary work to something that a machine can understand, which in our case ends up being a TF-IDF based vector of chapters for each text. We discuss this basic vectorization technique, one that is very common in NLP, more below as it has benefits beyond textual analysis given the nature of the mathematical, really geometrical, form and content of TF-IDF vectors.



**Figure 8.** DEBRA modular architecture.

<sup>18</sup>Python code available in Github at [https://github.com/Doresic/hierarchical\\_clustering](https://github.com/Doresic/hierarchical_clustering)

The next module the pre-processed text is fed into is the Topic Modeling component which uses K-Means clustering to extract N topics from the text, with each topic being represented by N key terms, the output of which is then fed into the Relation Extractor, which, as the name indicates, identifies the relations inherent in the literary work and then passes this on to the core hierarchy construction part of the algorithm. The pseudocode for the algorithm is illustrated in **Figure 9** below.

The first text we used as input to test DEBRA with was the Bible (King James Version) simply to see what kind of conceptual hierarchy could be extracted from the text itself, Genesis in particular, to acquire a basic understanding of what kind of conceptual hierarchical structure, or map, could be in a sense “recovered” from the text using (at least at first) fairly standard and well understood NLP techniques that were widely available and well-studied in the research literature. The books of the Old Testament are so well understood that, heuristically at least, “correctness” could be evaluated simply by looking at the derived conceptual map, rather than again conforming the output toward some sort of

---

**Algorithm 1** DEBRA Relation Extraction Algorithm

---

**Input:** Set of key terms  $K$   
Set of cluster sentences  $S$   
Relation ordering measure  $M$   
Hierarchy depth  $h_d$   
Hierarchy width  $h_w$

**Output:** Hierarchy relations  $H$

```

1:  $H \leftarrow []$  ▷ Initialize hierarchy list.
2:  $C \leftarrow K$  ▷ Initialize current layer terms.
3:  $R \leftarrow \text{GET\_RELATIONS}(K, S, 0)$  ▷ Get relations between key terms.
4: for  $(Subject, Verb, Object)$  in  $R$  do
5:    $C \leftarrow C \setminus \{Object\}$  ▷ Remove all inferior terms, as they will appear in the next layers.
6: end for
7:  $U \leftarrow C$  ▷ Initialize set of used terms.
8: for  $d$  in  $\text{range}(0, h_d)$  do
9:    $R \leftarrow \text{GET\_RELATIONS}(C, S, 1)$  ▷ Look for all children of current level words.
10:  for  $(Subject, Verb, Object)$  in  $R$  do
11:    if  $Object$  in  $U$  then
12:       $R.\text{remove}((Subject, Verb, Object))$  ▷ Remove all used terms.
13:    end if
14:  end for
15:   $R \leftarrow \text{ORDER\_RELATIONS}(R, M)$  ▷ Order relations with respect to the given measure.
16:  if  $h_w$  is not None then ▷ Restrict number of children to a maximum of  $h_w$ .
17:     $counter \leftarrow \{\}$ 
18:    for  $(Subject, Verb, Object)$  in  $R$  do
19:      if  $Subject$  in  $counter.\text{KEYS}$  and  $counter[Subject] > h_w$  then
20:         $R.\text{remove}((Subject, Verb, Object))$ 
21:      else if  $Subject$  in  $counter.\text{KEYS}$  then
22:         $counter[Subject] = counter[Subject] + 1$ 
23:      else if  $Subject$  not in  $counter.\text{KEYS}$  then
24:         $counter[Subject] = 1$ 
25:      end if
26:    end for
27:  end if
28:   $C \leftarrow \emptyset$ 
29:  for  $(Subject, Verb, Object)$  in  $R$  do ▷ Update used terms and current layer terms.
30:     $U \leftarrow U \cup \{Object\}$ 
31:     $C \leftarrow C \cup \{Object\}$ 
32:  end for
33:   $H.\text{EXTEND}(R)$  ▷ Extend the hierarchy with the next layer.
34: end for
35: return  $H$ 

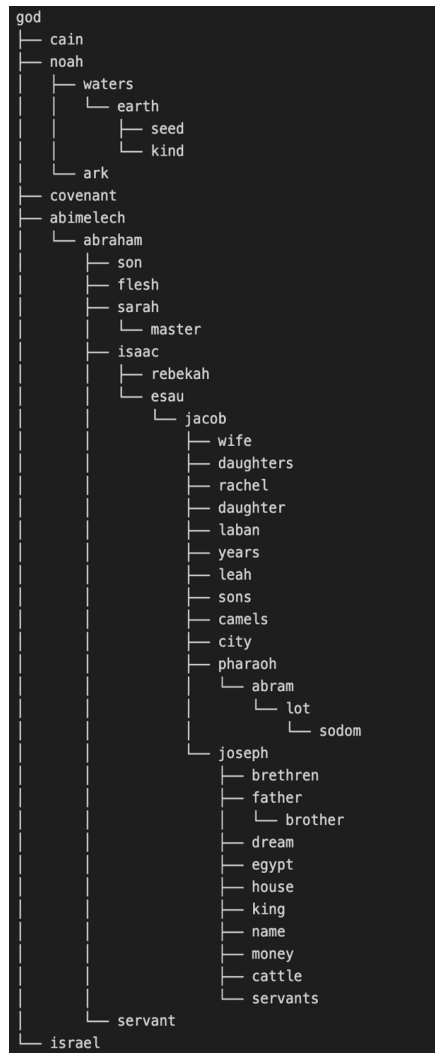
```

---

**Figure 9.** DEBRA algorithmic structure in Pseudo-Code.

measurement criteria which is typically the case in the Ontology Learning domain [35] [36]. The text itself is the guide for DEBRA, given the fully automated, *i.e.* unsupervised, nature of the solution we were working with.

In initial versions, we started seeing things like the figure below, images and hierarchies that gave us a good sense as to what core linguistic analysis, using some of the latest libraries available in Python (spaCy, NLTK), could produce when looking at a literary text that in itself was quite unique, both in terms of the language used and of course the subject area, neither of which had been analyzed in this fashion before as far as the authors knew, at least with respect to the lack of research available into the analysis of culturally significant literary works using the latest in ML and NLP tools and techniques. There was certainly plenty of research into the creation of ontologies to support the development of the Semantic Web but as we relate here with this work, while this problem is related to what we are doing with DEBRA its purpose is significantly different and as such it follows different design principles and produces different outputs (**Figure 10**).



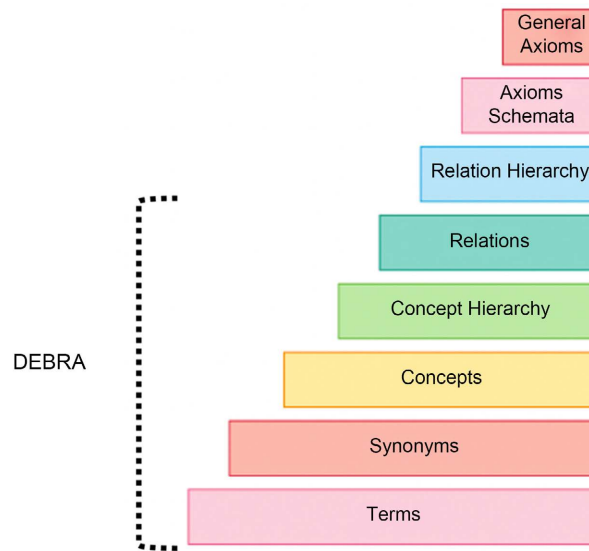
**Figure 10.** Early DEBRA output from Genesis.

What we came to understand much later, as we began studying the literature on Ontology Learning, was that what we were developing shared at least some similarity with the problem that had been given the moniker of *lightweight ontologies* [37], a construct illustrated in The Ontology Type visualization we provide in **Figure 7** above (from [33]), which is presented as a less mature form of ontology along a spectrum of ontology development, with full-fledged ontologies that include semantic and axiomatic structures associated with the concepts as well as their relational, and ontological, structure. The latter of which we are not interested in at all with DEBRA.

To understand how DEBRA fits into the current research landscape within Ontology Learning it is necessary to understand the notion of the Ontology Learning *layer cake* which is the generally accepted way of describing the process by which full-fledged ontologies are developed, process wise, a corpus of textual input—a corpus that is presumed to be in digital format and consisting of many documents in a variety of formats (structured, unstructured and semi-structured) that we are interested in encoding into an ontological structure for the purposes of machine learning and understanding. The methodology presumes a modular, iterative process again consistent with ML in general, by which core terms are extracted from the textual corpus, after which concepts and relations and their respective hierarchies are constructed, after which the more formal axiomatic and semantic structure is established [34].

While DEBRA most certainly follows the same bottom-up construction that we find in the standard layer cake, it nonetheless stops just prior to the creation of the more formal aspects of the ontological structure, as illustrated in **Figure 11**. As we can see from the illustration above, with DEBRA we are effectively focusing on the core building blocks of the ontology layer cake, looking at concepts and terms and their relations primarily, with the subtle and yet important distinction that we are not concerned with creating a taxonomic structure—“type of” or “instance of” type relations, nor are we concerned with identifying synonyms, each term or concept is evaluated independently and any relations that we establish between concepts comes from the text itself. The same cannot be said of Ontology Engineering in the general case which pulls from a variety of resources external to the textual input to guide the creation of the ontological structure.

What we're left with, by design, is a crystalized representation of the (literary) text information in the form of a concept hierarchy, or again map, that is 100% reflective of the underlying conceptual structure inherent in the text itself, rather than a semi-supervised created ontological structure of a domain whose formation is guided by a domain expert or third party. This is not a question of which technique is better or worse, of course, simply a matter of which is more suited for the task at hand. Given that we are looking for encoded conceptual structures, we want the (literary) text to inform us rather than look to a source outside the text to establish the ontological structure. The master source is the (literary) text itself, not the intellectual domain as it has been defined by industry, or again expert, standards.



**Figure 11.** Ontology learning layer Cake & DEBRA.

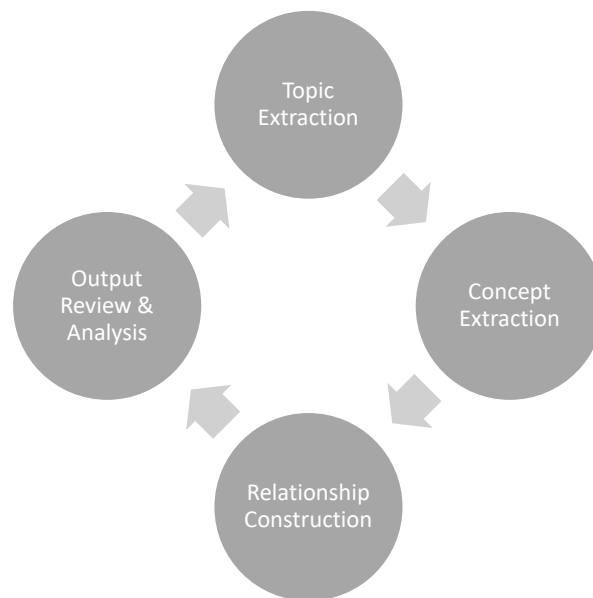
DEBRA does, however, share many of the same functional and mechanical development characteristics of Ontology Learning/Engineering though, with as already pointed out the ground up build methodology that is expressed in the core layers of the ontology learning layer cake, with DEBRA we're just looking for the text itself to be the guide—given the algorithm we have developed to create the concept hierarchies—rather than again a combination of the text and an outside party from which the basic ontological structure of the given domain is pre-established.

Also, like most machine learning algorithms, our approach to eliciting the conceptual hierarchy encoded in a given text is iterative, as we run DEBRA initially against a literary text with a base set of parameters, after which we analyze the output to determine if the expected (concept hierarchical) output is more or less “correct” given what we understand about the nature of the concepts embedded in the text, after which we modify and tune system parameters and then re-run. This iterative process flow is depicted in **Figure 12**.

Tuning parameters include the number of topics, the number of keywords, the depth of the conceptual hierarchical tree as well as other options that are built into the code. We follow this process, this tuning, until we land on a concept hierarchy that resembles what we know about the conceptual structure of the text itself, this is why it's important, at least for the initial version of DEBRA, which we work with texts that we are familiar with and understand.

### 5.1. On Preprocessing & TF-IDF Vectorization

The first thing we must do with the literary text that we wish to use as input to DEBRA, and this is consistent with NLP problems in general, is to transform the input data into a form that can be computationally analyzed, *i.e.* understood by the machine effectively, a process that follows just a few well-ordered and sequential steps after which we have a core underlying data structure which can



**Figure 12.** DEBRA engineering iterative process flow.

then be used for concept and relation extraction. The specific steps and code that we use for this process, the preprocessor, is very specific to the literary text in question and as such are quite different for example for the Eurasian Philosophical work versus the Old Testament Books of the Bible that we use as test and prototype input into DEBRA, given the unique attributes of their respective input format and structure.

Having said that, what we are looking for at the end of the preprocessing component, regardless of the nuances of the text input, is a structure that *represents* the structure of the literary work in question, which from DEBRA's perspective is a set of chapters (or Books in the case of the Bible), each of which consists of a set of sentences (after stop word and punctuation removal) which transforms the words (tokens in NLP parlance) into something that the NLP part of the code can analyze and take apart into constituent pieces. Once in this form, we then can begin the process of concept and relation extraction, a process that in and of itself is very much dependent on the underlying language, or more specifically the underlying linguistic structure of the language which, in the case of all Western European languages at least, is predicated on the idea of a sentence, which in turn typically has a subject and object as well as a myriad of other attributes, the sum total of which is referred to generally as a language's *grammar*, a concept which is shared by both linguistics and computer science, which denotes the set of syntactic and morphological, and sometimes even semantic and phonological, rules that underpin the (proper use) of a given language. It is this structure, and the associated libraries that we use to decode it, which ultimately give us the information necessary to feed DEBRA's (lightweight) ontological structure.

*Preprocessing*—standard fare for all NLP pipelines is the step that takes the raw text itself and creates these TF-IDF [37] and sentence and chapter (book)

data structures so that the core analysis part of DEBRA can be run. Preprocessing, and filtering, code is predicated upon these TF-IDF structures to both filter out (and sort ultimately) the top N words in a given text that in turn feed our conceptual hierarchical output, with N being a core DEBRA system parameter, as well as parts of speech tagging which we perform using Python spaCy libraries that facilitate the deconstruction of sentences into parts that allow us to find concepts, and their relations, within the sentence structure itself, the decoding process as it were. This part of the analysis allows us to distinguish between nouns and verbs for example, giving us the ability to distinguish between nodes in our lightweight ontology (terms or concepts again) and their associative relationships (verbs) for example.

We also use the Python NLTK libraries for the extraction of *bigrams* and *trigrams* from the sentences and chapters within our input text which allows for a conceptual architecture that extends beyond words (or terms and/or tokens as they are sometimes called) into concepts which can cross over into two (bigram) or three (trigram) word pairs, which allows for the distinction of for example “ancient Chinese” as a unique concept that is distinct from both “ancient” and “Chinese” as they appear in the text.

In order to synthesize, or crystalize, the concepts and relations that we are interested in that will provide inputs into the final (lightweight) ontology, or more precisely our concept map, we construct an intermediary structure (pandas data frame to be specific) which consists of all of the words and terms in the text to be used as a filtering structure, a structure which supports for example the exclusion of stop words and other terms in the text that, due to either the specific word or because of how insignificant the term is relative to the text as a whole (more on this below) is filtered out of the analysis and does not make it to the second step of the algorithm. As is standard fare for NLP applications, this (stop-word) list includes like “the” “to”, etc. but with DEBRA it also includes specific verbs to be filtered out of the final node/concept list, determiners, pronouns, adverbs and numbers (with some exceptions). Among other means, we primarily use the term’s *term frequency* (DF), and *document frequency* (DF) measures, in the form of the *TF-IDF measure*, which is again standard, and widely used, in NLP.

Although widely used and well understood in the NLP literature, TF-IDF in brief a way to calculate the “significance” of a given word that is more sophisticated than a simple frequency count in a given text, allowing us to account for how relevant, or unique, a given term is within a given corpus. So for example it gives the ability to discern between a term (token) like “and” or “or”, each of which would have very large count metrics in our text vectors or matrices relative to terms or tokens that we are much more interested in like “philosophy”, or “tradition”, or “deity”, the latter of which would be over shadowed in a count vector (aka bag of words model) by words that are just simply common even if conceptually they are not interesting.



$$tf - idf(t, d) = tf(t, d) \cdot idf(t)$$

$$idf(t) = \log \frac{1+n}{1+df(t)} + 1 \quad (1)$$

-  $tf(t, d)$ , term frequency of term  $t$  in document  $d$ .

-  $n$ , number of total documents, *i.e.* sentences, in the corpus.

-  $df(t)$ , number of documents that contain  $t$ . It ranges from  $[1, \log(1 + n) + 1]$ .

The less frequent terms are multiplied by greater values.

-The resulting vectors are normalized by the Euclidian norm.

Scikit-Learn's TfidfTransformer in default settings.

Equation (1): TF-IDF Formula<sup>19</sup>.

The  $tf$  part stands for *term frequency*, which is a number given (a count), *i.e.*, the frequency, a given term is found in a given document. Added to this however, is the normalization of the count for each document, so each term is divided by the number of tokens in said document. Now we know the relative importance of each of the terms in a given document, but we know for example that the word “the” or the word “or” or “and” (aka “stop-words”) are not terms we are especially interested in if we’re looking for example for the most relevant terms in a given document.

To account for words that may be found frequently in the text but are not necessarily significant to a given chapter, book or document in a given corpus, a problem that is very relevant to Information Retrieval for example, we calculate another statistic, *inverse document frequency* which, as the name implies, is simply the inverse frequency of the number of documents the term is found in, on a log scale so as to smooth out the effect of increase counts. We can then multiply the two metrics together to get a “score” for every term in the “corpus”, giving us a way to rank all of the terms in a given corpus across each other, where the TF-IDF metric increases proportionally to the number of times a word appears in a given document, offset by the (log of the) number of documents in the corpus that contain the word.

It is the TF-IDF vectorization technique that proves most useful as input into our topic modeling aspect of DEBRA perhaps not surprisingly, as it allows for the algorithms to “see” the documents in a form that reflects all of the words in a given document (book or chapter) in a given corpus (literary text), but also includes a “relevance”, or “importance” factor associated with each word or term (or token) in the text that is more sophisticated than a simple count which denotes the importance of that term within its specific context—document or chapter in our case.

So this intermediary data structure then, once it has been populated with the necessary TF-IDF values for each (relevant) term in the documents and chapters we are looking at with DEBRA, serves a couple of purposes:

- *Firstly*, it is used as a filter of terms we are interested in and would like to include in our analysis,

<sup>19</sup>From <https://towardsdatascience.com/an-overview-for-text-representations-in-nlp-311253730af1>

- *secondly*, it is a database of the TF-IDF, TF and DF measures of all terms, later used when ordering extracted relations,

- *lastly*, as it is ordered with respect to the desired measure (default is TF-IDF ), from it we can extract the top N words, where N is a configurable parameter, and pass them onto the relation extractor as an alternative to the key terms generated by the topic modeling module.

## 5.2. On Topic Modeling and Key Concept Identification

Once we had a good working model of concept hierarchy construction, it became evident, at least for the texts we were working with that a delineation of concept hierarchies by topic might be useful as sort of guiding hand for the establishment of the lightweight ontological structure we were looking to build. There were a few reasons for this, but ultimately we looked to Topic Modeling (using K Means clustering as described in more detail below) as a way to establish the top root nodes of our conceptual map given that it provides a very nice, clean, mathematically elegant and straightforward way of identifying the main topics in a given text as well as key words that are most representative of each topic, the latter being what we end up using to “seed” our conceptual model.

In other words, if you want to fully appreciate and understand how a literary text encodes information in the linguistic conceptual hierarchical sense at least—conceptual relations should be understood within the context of their own conceptual sphere so to speak, and the cleanest and most elegant way to divvy up a text into these “conceptual spheres” is by mining the topics from the text itself using standard statistical techniques, a problem that is well understood within NLP.

Note that this design choice, to use Topic Modeling to help us seed the basic ontological structure we are generating, is distinct and unique to the task at hand, and while topic modeling is most certainly one of the methods that is spoken of and leveraged in the creation of various aspects of an ontology within the context of Ontology Learning & Engineering [38], with DEBRA we (uniquely) use it to seed the top layer of the ontological structure, a design choice that would not be appropriate if we were intending to build a more formal ontological structure.

It is important to note a subtle but important distinction here with respect to how we use topic modeling as a means to facilitate the extraction of our conceptual hierarchy, and this distinction is driven from the very specific nature to the problem we are solving for with respect to how it is distinct from the majority of NLP problems that take up the majority of the research and development resources in the field of NLP more broadly, namely in the support of Information Retrieval (Search as it is called in the industry and outside of research circles). For generally the modeling of topics for a given textual corpus (corpus here not a single literary text which again is the case with DEBRA) is primarily done to facilitate the ability to navigate, *i.e.* search or categorize the corpus of digital assets, whereas what we are looking to do with DEBRA is the extraction of topics

from a single literary text, or source, that is divided into chapter, or Books in the case of the Bible. While these two problems are similar enough that they rely on basically the same statistical and mathematical techniques, their different purposes drive different usage even though the output (of this module) is the same.

To identify the best topic modeling algorithm to use we ran a variety of topic modeling algorithms (LDA, LSA, K-Means Clustering, Non-Negative Matrix Factorization) against the literary text(s) we were testing against to evaluate their relative performance, using again a heuristic for “understanding” which was self-evident given that the types of literary texts we were working with and our understanding of said texts. Illustrations of this self-evidence will become more clear as we look at the DEBRA output we got in our testing further into the analysis. The specific method we used to perform the analysis, and this was true of all the topic modeling techniques we tested with, was the creation of the respective topic models from the literary text itself (or some subset if we were only looking at specific Books or Sections of the text in question) and then feeding the text back into the model that we created to see how the model organized the same textual input into topics and their respective key words. Note that this deviates from standard ML practice where typically you will partition your input data into development and test data sets and then build your model against the development data set and then test it against your test data set, but in our case we turn the model against the input data itself and see how it self-evaluates in a way, and for sure not all models handled this relative subtle design change elegantly.

What we found, is that the algorithms that were the most well suited for such a task were not the ones that were most touted in the literature on Topic Modeling such as Latent Semantic Analysis, also known as Singular Value Decomposition which was put forward in 1990 [39], or even Latent Dirichlet Allocation [40], techniques which represent more or less the industry standard(s) to solve again Information Retrieval and organization problems at scale, but that the K-Means clustering algorithm, which is quite long in the tooth as far as AI algorithms go with its mathematical foundations going back to the 1950s and 1960s [41] [42], along with a relatively newer technique known as Non-Negative Matrix Factorization, or NMF for short [43], actually did the best job of organizing the chapters (or books) of our literary texts into their appropriate, optimal, topic sets.

After a host of experimentation, we found that at least for use for the kind of texts and input data that we were working with—the most effective topic modeling technique is (from a mathematical and statistical point of view at least) one of the most straightforward. The technique we landed on was a variant of K-Means clustering which is a statistical technique that was put forward in 1967 [44] and was further elaborated on in 1979 [45] and is one of the most widely used statistical methods for clustering data around this notion of *centroids*, which are computed iteratively based upon a pre-defined objective function that we look to minimize/optimize as we iterate through the data and create our “clusters” around these centroids.

Of course, there are some necessary conditions to getting these topic models to “work” (or “fit” which is the technical term that is used in ML) the most important of which is transforming the data into something the algorithm can understand, *i.e.* a TF-IDF text vector for each chapter such that the chapters themselves can be “processed”, by the algorithm in question. This data structure is a very important intermediary step for DEBRA, both in terms of supporting the process of concept and relation extraction as well as providing the underlying semantic geometry to support the comparing and contrasting of literary texts in general, one of the natural byproducts of the research work herein (more on this below). It is the TF-IDF vectorized form of the text that we use not only to support the filtering process as we describe above, but also as the underlying data structure that we pass into our topic models, as a form of knowledge representation effectively, as input into our (K-Means) topic modeling module. It is the TF-IDF vectorization technique [37] that proves most useful as input into our topic modeling aspect of DEBRA perhaps not surprisingly, as it allows for the algorithms to “see” the documents in a form that reflects all of the words in a given document (book or chapter) in a given corpus (literary text), but also includes a “relevance”, or “importance” factor associated with each word or term (or token) in the text that is more sophisticated than a simple count which denotes the importance of that term within its specific context—document or chapter in our case.

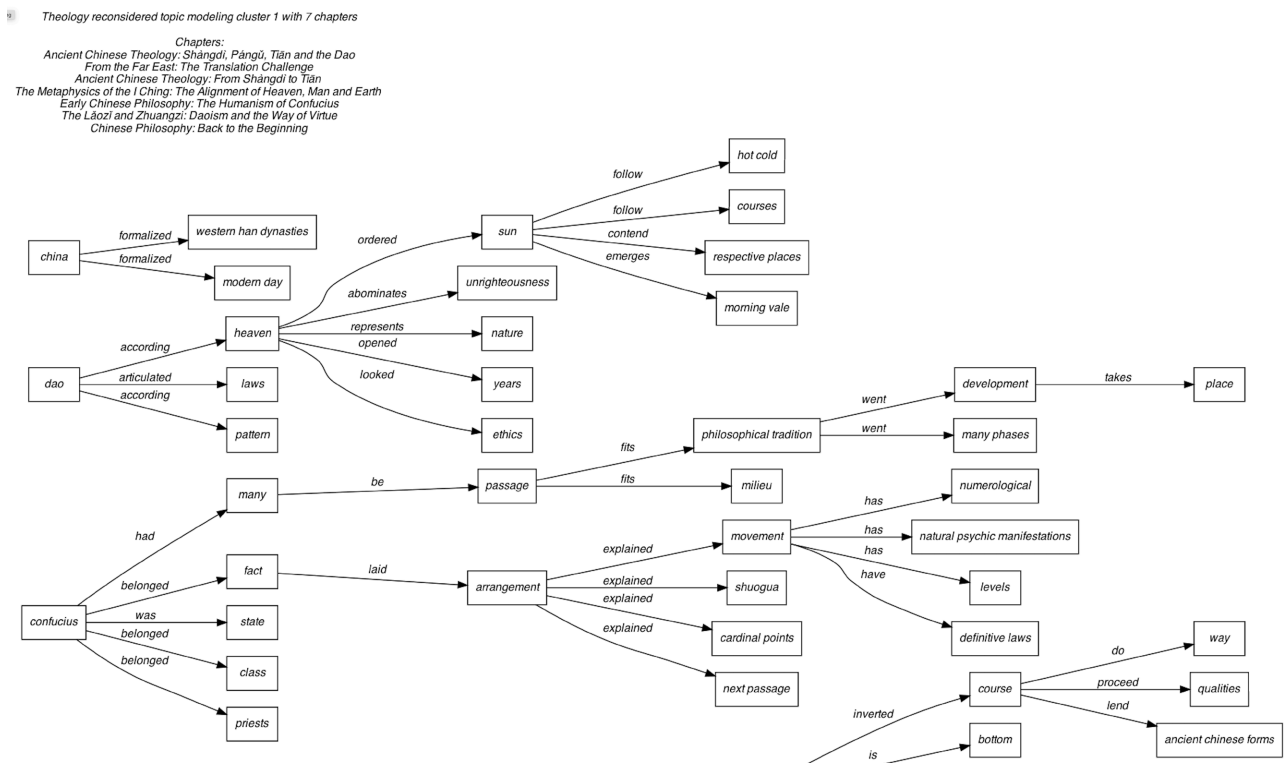
With respect to how it is that the topics, and their related key terms, seed the ontological structure we create with DEBRA, we can get a better idea as to how this works by using an illustrative example from our testing. So for example when we were working with the analysis of the philosophical work related to the emergence of Eurasian Philosophy in antiquity [12], which lends itself to such analysis given that the topic of a given chapter is encoded into the chapter name itself, we see quite clearly (heuristically) that the conceptual hierarchy for ancient Hellenic philosophy is separate and distinct from, conceptually or intellectually, to the topic and related concepts and relations that are associated with Chinese philosophy, or Enlightenment Era/Age of Science philosophical developments for that matter. Three branches from the same tree perhaps, but different branches nonetheless—a branch in this case corresponding quite neatly to the topics that our topic generator yields from its analysis of the text in question.

To take advantage of this correspondence, we generate the topics and then use the output of this module to feed into the subsequent DEBRA modules to (with some exceptions that we outline below) establish these topic driven key terms at the root of our (lightweight) ontological structure, and then run DEBRA’s concept and relation extractor off of that given key word itself to build out the encoded conceptual hierarchy around that specific term or concept. In this way the system has a way to naturally and ultimately statistically and linguistically—split out a text by topics and then yield concept hierarchies for each topic that can be evaluated, and tuned, independently. This is certainly useful when dealing with a text like *Theology Reconsidered* or the Bible for example, each text which covers

a broad range of topics each of which one would expect to have unique embedded conceptual hierarchies of meaning.

In **Figure 13** below for example we see the (partial) output of DEBRA against again the same text dealing with Eurasian Philosophy in antiquity [12] for just one single topic, namely Chinese Philosophy. We see the chapters (upper left corner) that it tagged as belonging to this given topic, all of which, again given the specificity of the chapter titles, clearly belong to said topic. Then the root nodes become the terms that are output by the Topic Modeling piece of DEBRA, after which she fills out the corresponding concept hierarchies for each of the terms that have been identified by the topic modeling module as (most) representative of said topic. This type of analysis allows for a more narrowed down, higher resolution view of the encoded conceptual structure of a specific, topic driven, part of the literary text that we are analyzing.

We use this technique as the starting point of DEBRA, generating the core topics and their related key terms from the literary text itself we are analyzing, first normalizing the text data input via preprocessing steps and then vectorizing the tokens by chapter using the TF-IDF metric, computing the clusters using K-Means with a fixed number of clusters as input. The K-Means cluster model, sometimes referred to as Lloyd’s algorithm named after its founder from Bell Labs, aims to group similar documents, or chapters in this case, together into clusters based on their feature representations, or their TF-IDF vectors, ultimately partitioning the (text) data into K clusters, where K is a user-defined parameter representing the number of desired clusters.



**Figure 13.** DEBRA output (partial): Chinese philosophy topic.

### 5.3. On Relation Extraction, Ontological Structure & Visualization

Now that we have computed the set of key terms around a given topic from a given text, or simply the top N non-filtered words if we are not using the topic modeling option, we then feed these terms into DEBRA's engine for the final hierarchical tree construction. This part of DEBRA generates the bulk of the concepts and their relations, and also organizes them into a hierarchical, or ontological, structure which is unique to DEBRA and one that is significantly different from what one might expect when working with more standard Ontology Engineering tools.

The first thing this engine does is extract all of the directed relations between concepts, which (loosely) correspond to nouns or subjects from a linguistic and *syntactic* perspective, the underlying sentence syntax being broken down by spaCy's sentence/language dependency parser, what is sometimes referred to as *parts of speech tagging*, or POS tagging for short. With this library we find relations by extracting subject-object pairs, identifying negative relations, replacing pronouns, using conjunct subjects, and other techniques that are required in order to establish the relationships between concepts as well as their ontological ordering, or status you might say. If the language of a given text is particularly unique, as is the case with the King James version of the Bible for example, some intelligence must be added to the parser to ensure the relations and their relative ordering in the concept hierarchy more accurately reflects the meaning, or underlying *semantics*, of a given sentence in a given text.

As a relatively simple example, from the spaCy dependency parser for the sentence "*Then Jesus sent the multitude away, and went into the house.*" Depicted in **Figure 14** below, DEBRA will first detect the verbs of the sentence which in this case are "sent" and "went". Next, the algorithm will search for *subjects* connected to these verbs, searching within the sentence for a *nsubj* dependency connection for the specific verb in question. Using this method, we do in fact find the term "Jesus" as a subject connected to the verb "sent" for example.

For the second verb however, "went", the process by which we identify the concepts associated with this action, and their relative ontological status, is a little more involved. The algorithm first looks for concepts which are associated with this action, or verb, by using *conjunct verbs* via the *conj* attribute from spaCy. If we don't find anything meaningful there, we next we look for objects within the sentence that are tagged with *obj*, denoting a direct object dependency on the verb in question in which case the relation is clear. In the case where none can be found, we extend our search through prepositions (*prep* dependency), a process which ultimately leads us to the objects "multitude" and "house", so that ultimately for the sentence, "*Then Jesus sent the multitude away, and went into the house*", the relation extractor will extract two relations: Jesus->multitude and, using a conjunct subject, Jesus->house.

Another interesting example is the dependency tree presented in **Figure 15** for the sentence "*No man hath seen God at any time*". As in the previous example,




Figure 14. spaCy dependency parser, simple example.




Figure 15. spaCy dependency parser, negative determiner example.

the relation extractor will find the relation (man->God), but this time due to the negative determiner “No” attached to the subject “man”, the inverted relation (God->man) will be extracted. One can say that the relation extractor searches for action-oriented relations.

As these relations are mostly based on subject-object pairs, as one might expect, they have an action-oriented nature which is a function of the structure of the English, object and action-oriented language basically. With DEBRA though, we search for higher order entities that are implied by the underlying language structure. For example with the sentence “*Then answered Peter, and said unto Jesus, Lord, it is good for us to be here...*” even though one could easily argue that the superior entity in this relation is “Jesus”, the relation (Peter->Jesus) will be extracted as Peter is the one who is the actor in this case. However, across a whole literary work, we expect that a superior entity will be the actor in more cases, hence we have with our relationship extractor also an ontological paradigm that is also imposed about the sentences of the text as they are parsed and feed into the conceptual hierarchy.

The relation extractor also:

- Replaces pronoun subjects with non-pronoun subjects they refer to,
- uses conjunct subjects if there is no direct one,
- searches for subject or object compounds in order to construct bigrams or trigrams,
- searches for direct and indirect objects through numerous dependency combinations: *dative-pobj*, *prep-pobj*, *prep-conj*, *prep-prep-pobj*, *dobj-conj*, etc., and
- searches for negative verb adverbs, negative subject or object determiners which would invert the relation.

In the pseudo-code in Figure 9, the majority of this functionality is not shown



due to brevity but the bulk of this dependency tree parsing would occur between lines 4 and 5 of the GET\_RELATIONS function.

As one can see, when we analyze a given sentence, we look to glean the underlying conceptual ontological structure embedded in the sentence itself, a structure which typically can be picked up just from the subject->object relations but in many cases, certainly with the King James Bible, some further analysis is required which is coupled with semantic assumptions that are encoded in the structure of the English language. In other words, the conceptual hierarchical structure, *i.e.* the ontological structure, which DEBRA generates originates both the syntactic structure of the text as well as the semantic meaning as well in some cases.

In fact, some of this ordering of relations occurs prior to the hierarchical construction part of DEBRA that occurs during relation extraction. Before feeding relations into the iterative hierarchy construction algorithm, we would like to order the relations so the “more important” ones are given to the algorithm first. The measure with respect to which the relations are ordered can be user-specified but in our experience, we found a mix of relation occurrence and the TF measure of the relation’s superior word to be the best. In the example we provide above for example, this results in the relation (“Jesus->Peter”) being passed before the opposite since it occurs more often, and “Jesus” has a higher TF value across the text than “Peter”. As the hierarchy construction algorithm is iterative, the order of the relations given to it has a big impact on the final structure of the (lightweight) ontology, aka the conceptual hierarchy.

When we use the topic modeling aspect of DEBRA, instead of using the top N words using the standard TF-IDF measure, we identify the topics and key terms for and we use these to seed the top level of the ontological structure, *i.e.* the root nodes. But not all of these key terms end up as root nodes though, given the following constraints:

- *Relationship between key terms*: if two (or more) key terms are related to each other, *i.e.* they have a conceptual relationship, then we remove one—the inferior one with respect to its relation to higher order terms or concepts because it will be “found” on the next step of the algorithm and we do not want to have duplicate concepts within our ontology. So if “Jesus” and “Peter” were both in key terms, and we find a relation “Jesus”->“Peter”, then we remove “Peter” from key terms, because we know “Peter” will again be picked up in the next iteration of the algorithm (when searching for next level terms).

- *Key terms with no children*: in the process of searching for all relations coming in which the key terms are superior, we will sometimes find that a key term has been identified in the K-Means clustering algorithm that does not have any child relations, in this case we will also remove said key term as it would end up being a lone node, concept at the root of our multi-faceted tree.

In essence, there’s only two reasons why a key term won’t show as a root in the found tree: 1) it’s inferior to some other key term, or 2) it doesn’t have any inferior words (there are no relations in which its superior) so it has no children.

Once we have established the ordered, directed relations that are to be used as nodes in our concept map, we then create the structure itself, which is in dictionary form (Python dict) and consists of rows, entries in the dictionary, that contain parent words/terms/concepts and their children. During this step we check each node for cycles, and we construct the set of words/terms/concepts that have parents so that we can ensure the term is placed in the proper place in the ontology, *i.e.* in the bottom most place in the ontological tree effectively. If we are using topic key terms as our ontological root, it is these key terms that are used to drive the rest of the concept hierarchy, subject to the restrictions regarding conceptual relations and leaf nodes as described above.

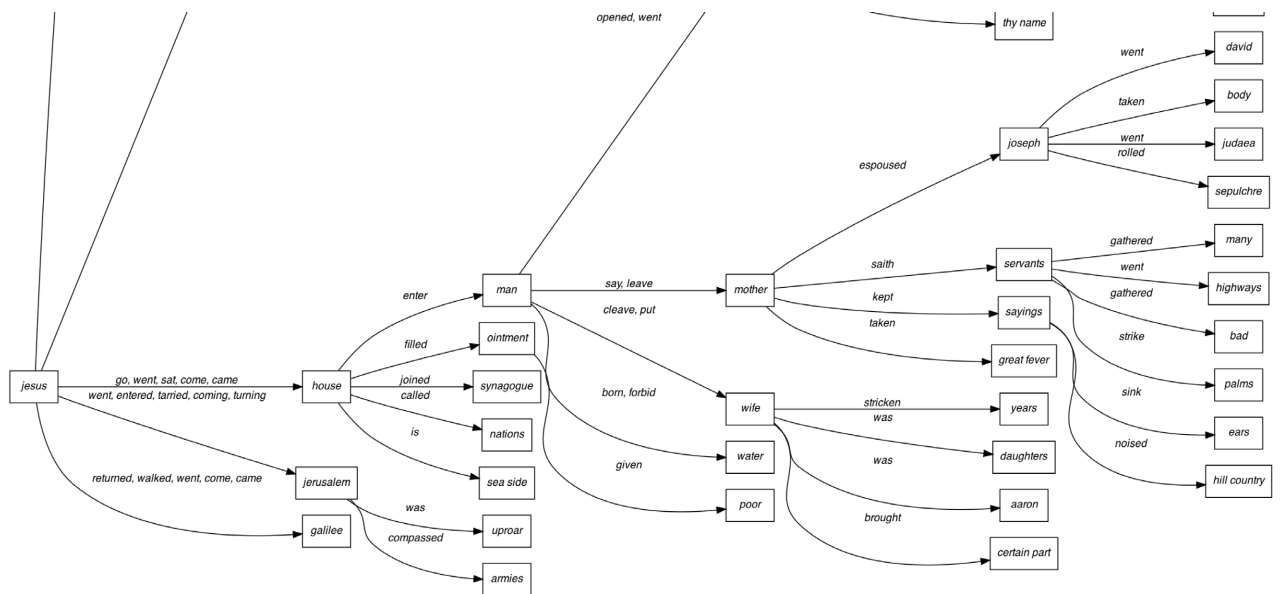
Once established, we then are able to visualize the encoded conceptual map using DEBRA, leveraging the Python *graphviz* library which provides a very clean and concise, and automatically formatted, network graph visualization which is ultimately what we're looking for to be able to see the "recovered" conceptual hierarchy from a given literary work or subset thereof.

Two examples are provided below, one from the (King James) Bible and another from a philosophical work [28] that lends itself to this type of analysis. As we can see from the figure directly below, DEBRA correctly identifies the four gospels and other Biblical chapters into a topic (four was the number of topics used in this particular run), which grouped the following chapters together: [Genesis, Judges, Ruth, Esther, Jonah, Matthew, Mark, Luke, John and Acts], thus establishing a clustered topic which more or less maps to Jesus and his teachings—as can be gleaned from not only the chapters that were associated with the given topic, but also the key terms that were associated with the given cluster, namely: [*said*, *thou*, *jesus*, *shall*, *said unto*, *ye*, *lord*, *came*, *man*, *thee*, *god*, *disciples*].

We can see from the lower extract of the full concept hierarchy (which is too large to show here), that Jesus has been placed at the root part of that conceptual structure and the key concepts or terms that it is mapped to, in the second level of the tree, are: *galilee*, *jerusalem*, *house*, and (not shown) *things*, and *disciples*, with the first three of these (or last three given that it's the bottom of the concept hierarchy) being shown in the figure below (Figure 16) which includes their child concepts as well as the relations (verbs essentially) that connect the parent nodes to their respective children.

Now if we use a more philosophical work as input into DEBRA, using again our text on Eurasian Philosophy in antiquity [12], we can see perhaps an even cleaner example of topic and key term extraction, along with conceptual hierarchical construction, by DEBRA as illustrated in Figure 13 which is an excerpt from the DEBRA output for the same text for the topic "Chinese Philosophy".

We can see here that our topic modeler did indeed correctly identify all of the Chapters that discuss Chinese Philosophy, as is clearly indicated by the Chapter titles, one of the distinctive features of the text, and the key terms associated with this topic, the bulk of which seed the ontological structure that we see above; namely *chinese*, *heaven*, *ancient*, *dynasty*, *confucius*, *ancient chinese*,




**Figure 16.** DEBRA Output Example I: King James Bible, Jesus in the Gospels.

“*dao*”, “*confucian*”, “*daoist*”, “*china*”. The algorithm then weeds out several of these, given the criteria described above in which acceptable root nodes from the key terms list are identified, and then walks through the creation of the conceptual hierarchy, yielding root nodes of *china*, *dao*, *confucius*, *ancient chinese* and *Chinese* in that order. The resultant conceptual structure which DEBRA produces a sort of intellectual roadmap for the topic at hand.

## 6. On Measuring the Similarity of Texts in Vector Space

It is worth noting here that an interesting corollary to the use of TF-IDF vectors, really Vector Space Models [1], which we leverage here to facilitate the generation of topics which feed into our (lightweight) ontology constructor (*i.e.*, DEBRA), is that this TF-IDF transformation of natural language texts, as a form of knowledge representation in and of itself, is extremely useful, and powerful, given that these (data) structures can, and are, used in many search applications as an intermediary format to facilitate search and document retrieval. For these TF-IDF vectors/matrices, by leveraging fairly straightforward linear algebraic mathematical techniques which we describe briefly below, can be used to measure document similarity, which is precisely the way these structures are used in Information Retrieval primarily, to generate a ranked list of possible matches. This same functionality can be used to rank similarity between and among literary texts, providing valuable insights into the nature, and more specifically the relationship, between and among literary texts of the same language (Figure 17).

This is a well-known attribute of the TF-IDF vectorization technique, the ability to identify document similarity in vector space, hence its widespread utility and study in the Information Retrieval literature. This attribute of these TF-IDF intermediary structures, which encode the text itself into Vector Space of size dictionary, could be also applied to the interpretation, and relative influence, of



**Figure 17.** Vector space model<sup>20</sup>.

philosophical and theological works from antiquity, a field which is ripe for some sort of normalized, mathematical foundations<sup>21</sup> In other words, this underlying form of data, or knowledge, representation allows for the computation of document similarity *in general*, not necessarily related to a user search query but as it pertains to a question regarding the relationship of two or more literary texts.

So if you wanted to know whether or not Kant was more of a Platonist or a Peripatetic for example, while certainly we could turn to philosophical experts to answer this question, this type of data representation (presuming all the texts are translated into the same source language) can be used to answer this question mathematically, by looking at the relationship between the TF-IDF based matrices which are computed to represent the texts in question (again assuming they are all parsed in the same source language) using the cosine similarity formula, or some derivation thereof depending upon how the document frequency metric is calculated (more on this below).

Another interesting use case for this natural byproduct of this form of data (really natural language) representation would be to use this technique to evaluate the authenticity of certain texts. So for example, there remains an open question in the field of Classics, really Hellenic philosophy, as to the authenticity of a series of extant letters which are attributed to Plato, in particular the famed *Seventh Letter* which has some bearing on the overall interpretation of Plato's philosophy and the status of the so-called *unwritten teachings* which are ascribed to Plato [6] [12], via the presumption of the authenticity of this letter (among other sources). This is still in fact an open question in the Classics, Ancient Philosophy, and even Comparative Religious academic community with respect to not only Platonism as a philosophical discipline but also Hellenic philosophy more broadly as it relates to the status, and existence, of unwritten doc-

<sup>20</sup>Image from

<https://blog.christianperone.com/2013/09/machine-learning-cosine-similarity-for-vector-space-models-part-iii/>

<sup>21</sup>As is the field of metaphysics as well.

trine in ancient Greek philosophical, and mystery, schools<sup>22</sup>.

The basic system structure setup if one wanted to leverage TF-IDF vectors for document comparison would look, from a preprocessing step at least, very much like what we have done with DEBRA, and in fact very much like what is done with many standard NLP applications. We parse and preprocess the text and put it into a TF-IDF vectorized format for preparation for machine processing. Once we have these literary texts, the “corpus”, in this format we then can query these texts using TF-IDF parsed queries, or alternatively compare the texts directly.

When a user submits a search query, the query goes through the same preprocessing steps as the documents in the corpus, resulting in a query TF-IDF vector, which represents the importance of each term in the query. To perform the actual search, we simply compare the TF-IDF vector of the query to the TF-IDF vectors of all the documents in the corpus using a similarity measure, the most common of which is the *cosine similarity* measure. It calculates the cosine of the angle between the query vector and the document vectors, indicating how similar the query is to each document. The higher the cosine similarity, the more relevant the document is to the query.

$$\text{Cosine Similarity} = \frac{A \cdot B}{\|A\| \cdot \|B\|} \quad (2)$$

Equation (2): Cosine Similarity Formula

where:

-  $A \cdot B$ : The dot product of vectors  $A$  and  $B$ . This represents the sum of the element-wise products of the corresponding components of the vectors.

-  $\|A\|$ : The Euclidean norm (magnitude) of vector  $A$ , which is calculated as the square root of the sum of the squares of its components.

-  $\|B\|$ : The Euclidean norm (magnitude) of vector  $B$ , calculated in the same way as  $\|A\|$ .

The cosine similarity value will range between  $-1$  and  $1$ , and the value is interpreted as follows:


- *Cosine Similarity* =  $1$ : The vectors are identical, and the angle between them is  $0$  degrees. This means the documents are highly similar.

- *Cosine Similarity* =  $0$ : The vectors are orthogonal (perpendicular), and the angle between them is  $90$  degrees. This implies that the documents are dissimilar.

- *Cosine Similarity* =  $-1$ : The vectors are completely opposite in direction, and the angle between them is  $180$  degrees. This indicates that the documents are negatively similar (**Figure 18**).

Once you have computed the cosine similarities between the query and each document, we can then use this information to sort the documents in descending order of similarity, thereby allowing for the ranking of search results, or

<sup>22</sup>This is precisely what was done, or at least some form of this type of analysis was done, by a Harvard student in 2021 [46]. See also [47] for a more conservative and older treatment of the topic of the authenticity of Plato’s Seventh Letter and a more modern, and thorough, treatment of the issue from March of 2023 in the online journal *aeon* entitled *The sage and his foibles* at <https://aeon.co/essays/what-the-controversial-letters-of-plato-reveal-about-us>.



**Figure 18.** Cosine similarity measurements<sup>23</sup>.

again simply comparing and contrasting how similar, or different, two or more documents are from each other, using the same measure as comparison.

In practice of course, a wide variety of optimizations and supplementary techniques are either added to a TF-IDF metric to facilitate better, and faster, Search or Information Retrieval [48] [49] [50], or different techniques not based on TF-IDF vectorization have also been researched and are also effective in various application domains [51] [52] but generally it can be seen how powerful this relatively straightforward form of knowledge representation can be used to solve for problems related to the measuring of the relationship between and among documents in a given digital corpus for standard Information Retrieval applications, which in turn also can be applied to evaluate the textual similarity, and relationships, between literary texts of cultural significance.

## 7. Conclusion and Summary Remarks

Ultimately when we look at the Semantic Web from a philosophical perspective, we see an attempt at the creation of domain-specific reference data structures which embed semantic and syntactic information that facilitate the processing and definition of these domains with a specific bent towards the underlying data elements and their inherent relationships. It's a quest to take the unstructured and disorganized data of the first iteration of the Internet and make sense of it, to consolidate it and construct conceptual hierarchies—ontologies that facilitate both the understanding of the domain as a form of knowledge representation, along with the establishment of the formal relations and overall structure of the underlying concepts to support decision making (reasoning) within that domain.

The Semantic Web does set the stage for a higher order intelligence in this sense, as these disparate knowledge-based systems become more centralized and more standardized. We now can apply intelligence across domains ushering in capabilities related to the aggregation of knowledge itself in a manner that was

<sup>23</sup>Image from

<https://blog.christianperone.com/2013/09/machine-learning-cosine-similarity-for-vector-space-models-part-iii/>

never thought would be possible. DEBRA and associated concept mapping and networking software, all designed to store and manage hierarchies and other forms of conceptual data, should provide us some insights into just how intelligent computers can be, and in turn the next set of questions about not just the nature of computing itself, but the nature of next generation intelligence. The end goal, as with all artificial intelligence efforts really, is to create an intellectual backbone that supports the creation of intelligent agents against a common knowledge base structure at scale. In the Semantic Web then, in its ultimate vision, the Internet itself becomes the operating system on top of which these applications run. A lofty goal has no doubt, but nonetheless feasible given the state of modern technology.

Orthogonal to these efforts however, or at least some sort of intellectual offshoot to them, should be the leveraging of many of these same tools and techniques that have been invented to support Ontology Engineering efforts in order to facilitate the understanding of works that have supported, and do support, what might be called human knowledge writ large which at least from a certain perspective can be viewed as fully reflected in the texts, literary works, that represent the highest goals, or achievements, of the human mind. What ideological or from a philosophical perspective metaphysical or epistemological—structure is represented in these texts? How do these texts relate to each other? Is there a core epistemological or metaphysical structure to these works? Does this reflect bias in some way?

While these questions don't necessarily lend themselves to clear cut, "scientific", answers necessarily, they are relevant and important for the better understanding of the human condition, for a full complement of knowledge once could say, for the good of the state of knowledge and truth more broadly one might say, as intellectual benchmarks that we now have the opportunity to perhaps put on more sound intellectual, specifically statistical or mathematical—footing. It also provides us an opportunity to look at works within the corpus of historically significant theological literature and analyze it from a similar vantage point, to elicit metaphysical structure—ultimately concept maps as we have proposed here with DEBRA to look at these works from a different perspective. Not as harbingers of truth necessarily but as ontological structures that can be looked at and compared with other literary texts of the same genre, in a manner that is at best bereft of bias and at worst at least shares a common mathematical, semantic, and lexical structure.

With DEBRA we look to the current state of Ontology Learning & Engineering to see what kind of conceptual hierarchies we can elicit (unsupervised) from a given literary text of some cultural significance, looking primarily for the evaluation of the potential of these tools to yield a synthesized and crystalized conceptual map of the given work, through which we can come to a better understanding of the text itself within a glance, or somewhat superficial review, of the conceptual relationships that are embedded therein, ultimately serving to pro-



vide a higher level, and unique, form of knowledge representation of the work at hand which in turn serves to facilitate a greater and more multi-faceted understanding of the work, shedding light on extra dimensions of information that would otherwise lay latent in the text itself and would not reveal itself without intense study.

With DEBRA, we automate the creation of conceptual graphs, and while the framework no doubt shares many of the characteristics of Ontology Learning/Engineering. Ultimately DEBRA is designed to elicit knowledge and “meaning” from a single textual corpus rather than the extraction of conceptual relationships across a whole body of textual corpus within a given domain that facilitates “reasoning”, or “querying” capabilities of a large corpus of text, requirements which are typically strong design considerations that underlie the need for the creation of formal ontologies, taking with the Semantic Web for example. Both approaches nonetheless rest upon core natural language processing (NLP) tools and techniques to make sense of the textual corpus inputs that they are working with however, and certainly. DEBRA is no exception here.

Having said that, DEBRA does share many common components from both a process perspective and an architectural perspective with both semantic networks and ontologies in the sense that DEBRA can be understood as a precursor analytical step that creates the core conceptual structure of a given text that in turn has the ability to feed additional downstream layers (if need be) to create a more semantically rigorous, *i.e. formal*, knowledge representation structure that is necessary for the development of AI agents more generally.

Regardless of what is done downstream with DEBRA, we most certainly have developed a proof of concept that establishes a definitive, more intuitive, and crystalized form of knowledge representation which serves to elicit nuanced and subtle meanings and relationships from a given literary text which most certainly can facilitate understanding of a given text, and if nothing else provides an alternative form of the text which encodes core information about the important concepts within a text and specificity with respect to their relations.

### **7.1. Findings Summary**

With our research into the application of Ontology Learning and NLP tools and techniques on literary texts, and arguably the humanities more broadly, we found out a few things:

- Computer science and philosophy, in particular epistemology and logic, are closely intertwined and interdependent fields of study, and they should be understood in this way and should continue to inform each other,
- AI as a field of study takes a definition epistemological position, namely that which has meaning is that which can be ultimately normalized into true or false (logical) statements, hence the reliance on first order logic as a cornerstone to “reasoning” within the field of Knowledge Representation and Reasoning within AI.

- There is room for an alternative, balancing epistemological position that is more like human intuition rather than human reason, those terms themselves being well defined in the Western philosophical tradition post Enlightenment and have a deep rooted, symbiotic and complementary relationship that is well documented at the very root of all of the major Eurasian systems of philosophy from antiquity.

- The state of the current widely available toolset of NLP libraries, in particular in Python, should allow for more research into the utility of some of these algorithms and models (of knowledge representation) with regards to an additional dimension of understanding as it relates to literary tradition development.

- The ability to create conceptual hierarchical structures that encode the core concepts of a given work, or portion of a given work, is not just feasible but also quite flexible for anyone with knowledge of Python and a (perhaps limited) background in NLP.

- These conceptual hierarchies can also be understood as less mature, or more raw, ontologies—with respect to having that term is understood from a formal mathematical context with respect to Ontology Learning & Engineering as a field of study within AI and Computer Science.

- There are a variety of dimensions of understanding in the humanities that could be opened up by applying these techniques that we demonstrate in DEBRA to the study of culturally significant texts from various literary (linguistic) traditions, with the technical limits being established only by what is supported in the standard, widely available NLP libraries for Python (or any other programming language for that matter).

In brief, we show with this work what kind of capabilities a literary text concept hierarchy constructor like DEBRA has for eliciting conceptual graphs, unidirectional networks really, that represent a more abstract and less logical perspective than formal ontologies for a given corpus, facilitating a unique relationship with a text where in a single image you can (potentially) see the core conceptual architecture of a whole document or work, a different and unique form of knowledge representation which yields a more intuitive, and immediate, understanding of the conceptual architecture of a given work, its metaphysics in a word. One potential use of such a concept mapper like DEBRA would be for example to use the output as a sort of visual aid that might be coupled with a text summarization when a text is delivered as a part of search results, like what Google returns for a company search for example. In some sense, this work is very much like mind mapping, or UML reverse engineering design software sets out to achieve except there, like ontology engineering in its present form at least, a significant investment of expertise and design modeling is necessary to create the basic ontological, mind mapping, structure. With DEBRA this is done automatically from the language in the text at hand.

DEBRA then from a certain perspective represents a step along this chain of technological advancement, a step in the direction of leveraging state of the art

AI, ML, NLP and more specifically Ontology Learning techniques in order to come to a better understanding of the literary tradition upon which our cultural and social intellectual edifice rests. It provides us with an opportunity to create structures of comparison, mathematically, from which philosophical works of all ages can be compared and contrasted, and ontological structures be looked at from a purely analytical perspective as outputs from the algorithms in question rather than as interpretations of individuals that come from certain intellectual backgrounds with certain biases that do not necessarily add value to the analysis of a given work within the context of the intellectual tradition as a whole.

## 7.2. Concluding Remarks

What we find here with this endeavor then, the creation of a tool for the extraction of concept hierarchies from a literary text, our DEBRA, is the opening up of the possibilities of leveraging some of the latest developments in Computer Science and AI back into the humanities, and in particular back into the field of Philosophy, and its offshoots metaphysics, ontology and epistemology, all of which have informed the field of AI itself, thus allowing for a sort of returning of the favor of sorts where the AI fields which stemmed from the philosophical disciplines can then be used as analysis tools for the form of the discipline which they were born. Seems like a natural progression of intellectual progress in a way, supporting the continued mutually interdependent evolution of both the philosophical, computer science and even cognitive science aspects of these adjoining fields.

The nature of intelligence is most certainly predicated on logic, and we are finding that even simple building blocks at scale, that are designed to find patterns within data and then store relationships between said patterns across massively scalable neural networks (Deep Learning), converge on what we had considered to be the exclusive domain of human beings, homo sapiens aptly named. And we are now starting to see a glimpse of what the capabilities of artificial intelligence are, and they are astounding. In this context, it can be helpful to see Computer Science as the engineering branch of logic, which in turn of course is (classically) one of the three main branches of philosophy—Hellenic philosophy, Stoicism [5] [53]. This engineering aspect of Logic allows for the true testing ground of what the final capabilities are of Logic itself, what is computable, again computer science.

For at the end of this intellectual enterprise which starts at the very root of the Western philosophical tradition, and Eastern too if we are to accept that Chinese philosophy played some role in informing modern Computer Science (through Leibniz and the binary system itself), we arrive at a place where the boundaries of intelligence itself can be explored not just from a theoretical perspective as Turing, Gödel, Church and others did some one hundred years ago but from an engineering perspective to see what exactly sits at the end of these theoretical boundaries as it were. Certainly to this end, we sit at a very interesting and

unique time in history where many of these questions can be answered, and it most certainly feels appropriate to turn these techniques back upon the intellectual tradition itself to see what they can tell us about how we understand how it is that we got here, intellectually, and what if anything that can tell us about what direction we should be heading in.

## Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

## References

- [1] Worth, P. (2023) Word Embeddings and Semantic Spaces in Natural Language Processing. *International Journal of Intelligence Science*, **13**, 1-21. <https://doi.org/10.4236/ijis.2023.131001>
- [2] Worth, P. and Pande, S. (2022) Ideological Reference Architecture (IRA): An Epistemological Interpretation of Quantum Mechanics. *Open Journal of Philosophy*, **12**, 266-305. <https://doi.org/10.4236/ojpp.2022.123018>
- [3] Turing, A.M. (1950) Computing Machinery and Intelligence. *Mind*, **59**, 433-460. <http://www.jstor.org/stable/2251299> <https://doi.org/10.1093/mind/LIX.236.433>
- [4] Marc Cohen, S. and Reeve, C.D.C. (2021) Aristotle's Metaphysics. The Stanford Encyclopedia of Philosophy (Winter 2021 Edition). <https://plato.stanford.edu/archives/win2021/entries/aristotle-metaphysics/>
- [5] Valdez, J. (2014) The Legacy of Socrates: Ideas, Forms and Knowledge. *Journal of Social Philosophy Research*. <https://www.semanticscholar.org/paper/The-Legacy-of-Socrates%3A-Ideas%2C-Forms-and-Knowledge-Valdez/1161346a4da7d1f418669fd7da29148efdce70c2>
- [6] Valdez, J. (2017) Theology Reconsidered: Volume I Mythos and Logos. Lambert Academic Publishing, Saarbruecken.
- [7] Maedche, A. and Staab, S. (2002) Ontology Learning for the Semantic Web. *IEEE Intelligent Systems*, **16**, 72-79. <https://doi.org/10.1109/5254.920602>
- [8] Maedche, A. (2002) Ontology Learning for the Semantic Web. Kluwer Academic Publishers Group, New York. <https://doi.org/10.1007/978-1-4615-0925-7>
- [9] Palmer, J. (2020) Parmenides. The Stanford Encyclopedia of Philosophy (Winter 2020 Edition). <https://plato.stanford.edu/archives/win2020/entries/parmenides/>
- [10] Øhrstrøm, P. and Uckelman, S.L. (2022) Lorhard, Ramus, and Timpler and "The Birth of Ontology". *Journal of Knowledge Structures and Systems*, **3**, 48-56. <https://philpapers.org/archive/HRSLRA.pdf>
- [11] Valdez, J. (2015) Philosophy in Antiquity: The Greeks. LAP Lambert Academic Publishing, Saarbrücken.
- [12] Valdez, J. (2019) Eurasian Philosophy and Quantum Metaphysics. Dorrance Publishing, Pittsburgh.
- [13] Shields, C. (2022) Aristotle. The Stanford Encyclopedia of Philosophy (Spring 2022 Edition). <https://plato.stanford.edu/archives/spr2022/entries/aristotle/>
- [14] Valdez, J. (2021) Metaphysics Reconsidered: A Gnostic Reading of Kant. Dorrance Publishing, Pittsburgh.

- [15] Stang, N.F. (2022) Kant's Transcendental Idealism. The Stanford Encyclopedia of Philosophy (Winter 2022 Edition). <https://plato.stanford.edu/archives/win2022/entries/kant-transcendental-idealism/>
- [16] Hanna, R. (2022) Kant's Theory of Judgment. The Stanford Encyclopedia of Philosophy (Spring 2022 Edition). <https://plato.stanford.edu/archives/spr2022/entries/kant-judgment/>
- [17] Bohm, D. and Hiley, B.J. (1993) *The Undivided Universe: An Ontological Interpretation of Quantum Theory*. Routledge, New York. <https://doi.org/10.1063/1.2808635>
- [18] Nikhilananda, S. (1949) *The Upanishads, Volume 1*. 5th Edition, Ramakrishna Vivekananda Center of New York, New York.
- [19] Valdez, J. (2016) *Philosophy in Antiquity: The Far East*. LAP Lambert Academic Publishing, Saarbrücken.
- [20] Matos, L., Machado, J., Monteiro, F. and Greten, H. (2021) Understanding Traditional Chinese Medicine Therapeutics: An Overview of the Basics and Clinical Applications. *Healthcare*, **9**, Article No. 257. <https://doi.org/10.3390/healthcare9030257>
- [21] Nelson, E. (2011) The Yijing and Philosophy: From Leibniz to Derrida. *Journal of Chinese Philosophy*, **38**, 377-396. <https://doi.org/10.1111/j.1540-6253.2011.01661.x>
- [22] Sowa, J.F. (2000) *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks/Cole, a Division of Thomson Learning, Pacific Grove, CA 2000.
- [23] Russell, S. and Norvig, P. (2021) *Artificial Intelligence: A Modern Approach*. 4th Edition, Pearson Education, Inc., London.
- [24] Minsky, M. (1975) MIT-AI Laboratory Memo 306, June, 1974. McGraw-Hill, New York.
- [25] Findler, N.V. (1979) *Associative Networks: Representation and Use of Knowledge by Computers*. Academic Press, Cambridge.
- [26] Quillian, M.R. (1968) Semantic Networks. In: Minsky, M.L., Ed., *Semantic Information Processing*, MIT Press, Cambridge, 17.
- [27] Sowa, J.F. (1976) Conceptual Graphs for a Data Base Interface. *IBM Journal of Research and Development*, **20**, 336-357. <https://doi.org/10.1147/rd.204.0336>
- [28] Sowa, J.F. (1984) *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley, Reading.
- [29] Novak, J.D. (1977) *A Theory of Education*. Cornell University Press, Ithaca.
- [30] Novak, J.D. (1993) Human Constructivism: A Unification of Psychological and Epistemological Phenomena in Meaning Making. *International Journal of Personal Construct Psychology*, **6**, 167-193. <https://doi.org/10.1080/08936039308404338>
- [31] Novak, J.D. (1998) *Learning, Creating, and Using Knowledge: Concept Maps as Facilitative Tools in Schools and Corporations*. Lawrence Erlbaum Associates, Mahwah. <https://doi.org/10.4324/9781410601629>
- [32] Novak, J.D. and Cañas, A.J. (2008) *The Theory Underlying Concept Maps and How to Construct and Use Them*. Technical Report IHMC CmapTools 2006-01 Rev 01-2008, Florida Institute for Human and Machine Cognition, Pensacola. <http://cmap.ihmc.us/Publications/ResearchPapers/TheoryUnderlyingConceptMaps.pdf>
- [33] Giunchiglia, F. and Zaihrayeu, I. (2009) Lightweight Ontologies. In: Liu, L. and

- Özsu, M.T., Eds., *Encyclopedia of Database Systems*, Springer, Boston, 1613-1619. [https://doi.org/10.1007/978-0-387-39940-9\\_1314](https://doi.org/10.1007/978-0-387-39940-9_1314)
- [34] Cimiano, P. (2006) *Ontology Learning and Population from Text: Algorithms, Evaluation and Applications* (Vol. 27). Springer Science & Business Media, Berlin.
- [35] Raad, J. and Cruz, C. (2015) A Survey on Ontology Evaluation Methods. *Proceedings of the International Conference on Knowledge Engineering and Ontology Development, Part of the 7th International Joint Conference on Knowledge*. <https://dl.acm.org/doi/10.5220/0005591001790186>
- [36] Wilson, R.S.I., Goonetillake, J.S., Indika, W.A. and Ginige, A. (2021) Analysis of Ontology Quality Dimensions, Criteria and Metrics. In: Gervasi, O., *et al*, Eds., *Computational Science and Its Applications—ICCSA 2021*, Lecture Notes in Computer Science, Vol. 12951, Springer, Cham, 320-337. [https://doi.org/10.1007/978-3-030-86970-0\\_23](https://doi.org/10.1007/978-3-030-86970-0_23)
- [37] Sammut, C. and Webb, G.I. (2011) TF-IDF . In: Sammut, C. and Webb, G.I., Eds., *Encyclopedia of Machine Learning*, Springer, Boston, 986-987. <https://doi.org/10.1007/978-0-387-30164-8>
- [38] Rani, M., Dhar, A.K. and Vyas, O.P. (2017) Semi-Automatic Terminology Ontology Learning Based on Topic Modeling. *Engineering Applications of Artificial Intelligence*, **63**, 108-125. <https://doi.org/10.1016/j.engappai.2017.05.006>
- [39] Deerwester, S.C., Dumais, S.T., Landauer, T.K., *et al*. (1990) Indexing by Latent Semantic Analysis. *Journal of the American Society for Information Science*, **41**, 391-407. [https://doi.org/10.1002/\(SICI\)1097-4571\(199009\)41:6<391::AID-ASI1>3.0.CO;2-9](https://doi.org/10.1002/(SICI)1097-4571(199009)41:6<391::AID-ASI1>3.0.CO;2-9)
- [40] Blei, D.M., Ng, A.Y. and Jordan, M.I. (2003) Latent Dirichlet Allocation. *Journal of Machine Learning Research*, **3**, 993-1022.
- [41] Lloyd, S.P. (1957) Least Square Quantization in PCM. *IEEE Transactions on Information Theory*, **28**, 129-137. <https://doi.org/10.1109/TIT.1982.1056489>
- [42] Forgy, E.W. (1965) Cluster Analysis of Multivariate Data: Efficiency vs Interpretability of Classifications. *Biometrics*, **21**, 768-780.
- [43] Lee, D. and Seung, H. (1999) Learning the Parts of Objects by Non-Negative Matrix Factorization. *Nature*, **401**, 788-791. <https://doi.org/10.1038/44565>
- [44] MacQueen, J.B. (1967) Some Methods for Classification and Analysis of Multivariate Observations. In: *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, Berkeley, 281-297.
- [45] Hartigan, J.A. and Wong, M.A. (1979) Algorithm AS 136: A k-Means Clustering Algorithm. *Journal of the Royal Statistical Society, Series C*, **28**, 100-108. <https://doi.org/10.2307/2346830>
- [46] Perry, J.B. (2021) Examining the Authenticity of Plato's Epistle VII through Deep Learning. Bachelor's Thesis, Harvard College, Cambridge.
- [47] Bluck, R.S. (1949) Plato's Biography: The Seventh Letter. *The Philosophical Review*, **58**, 503-509. <https://doi.org/10.2307/2182043>
- [48] Fautsch, C. and Savoy, J. (2010) Adapting the tf idf Vector-Space Model to Domain Specific Information Retrieval. In: *Proceedings of the 2010 ACM Symposium on Applied Computing (SAC10)*, Association for Computing Machinery, New York, 1708-1712. <https://doi.org/10.1145/1774088.1774454>
- [49] Fu, Y. and Yu, Y. (2020) Research on Text Representation Method Based on Improved TF-IDF. *Journal of Physics: Conference Series*, **1486**, Article ID: 072032. <https://doi.org/10.1088/1742-6596/1486/7/072032>

- [50] Fei, L. (2022) Research on Text Similarity Measurement Hybrid Algorithm with Term Semantic Information and TF-IDF Method. *Advances in Multimedia*, **2022**, Article ID: 7923262. <https://doi.org/10.1155/2022/7923262>
- [51] Albitar, S., Fournier, S. and Espinasse, B. (2014) An Effective TF/IDF-Based Text-to-Text Semantic Similarity Measure for Text Classification. *15th International Conference on Web Information Systems Engineering*, Thessaloniki, 12-14 October 2014, 105-114. [https://doi.org/10.1007/978-3-319-11749-2\\_8](https://doi.org/10.1007/978-3-319-11749-2_8)
- [52] Wang, J. and Dong, Y. (2020) Measurement of Text Similarity: A Survey. *Information*, **11**, Article No. 421. <https://doi.org/10.3390/info11090421>
- [53] Valdez, J. (2014) Stoic Philosophy: Its Origins and Influence. *Journal of Social Philosophy Research*.