parabole.ai

Enterprise knowledge at machine scale

SAI ANVESH DURVASULA

Contents

Introduction

Does your enterprise knowledge graph truly represent the entire knowledge within your organization? No, not yet.

With the advent and widespread acceptance of Knowledge Graphs amongst large enterprises, expert level knowledge has become critical for the successful rollout of any AI-powered application. This requirement has made ontology development pivotal for all learning-based solutions that, necessarily, must capture and leverage the knowledge possessed by Subject Matter Experts (SME's). If the true scope of this sea-change is to be understood, currently implemented manual processes of ontology development can never successfully contend with a massive and ever-expanding data universe.

Ontology for building an enterprise scale knowledge base

An ontology is a set of concepts and categories in a subject area or domain that possesses the properties and relations between them. If we were to create a set of concepts as defined for, let's say, a small grocery store, we would include concepts including: vendor name, item, category, customer name, POS sale data, etc. We would also introduce further detailed concepts including; apples, lettuce, ground beef, paper plates, cereal, bread, etc., by their respective vendor and brand names. Based on the grocer's purchases, the connections with his vendors, his customers preferences, the relationships between sales and days of the week, connections between items purchased by individual customers, and other data, we could begin to build an ontology for our grocer. It becomes quite apparent, even at this small scale, that manually building truly comprehensive ontologies is a herculean task.

To think that ontologies can be created for even small or mid-size financial institutions without automation seems rather optimistic. For large banks, it is simply not possible within any reasonable time/cost constraint. *Moreover, if we imagine that such a manually constructed ontology could be created, it would have a fatal flaw; it would be static and unable to scale.*

In stark contrast to manual processing of information, a dynamically generated ontology brings much desired scale and relevance to learning applications, where an automated and self-learned fission process would continuously disambiguate incoming knowledge into the

formal ontology and enable optimal inference (decisions). This ontology will continue to build upon new concepts, attributes, hierarchies, and behavioral relationship rules.

The below-mentioned ontology creation process follows 3 key principles:

1. There is no such thing as “adequate knowledge” as more can be found with every tick of the clock, even in the most unassuming and informal of formats (internal/external documents).
2. Discovery and knowledge disambiguation processes must be free of structural bias and agnostic as to format.
3. Ontology must have multiple layers, with each layer forming a symmetric relationship among knowledge elements.

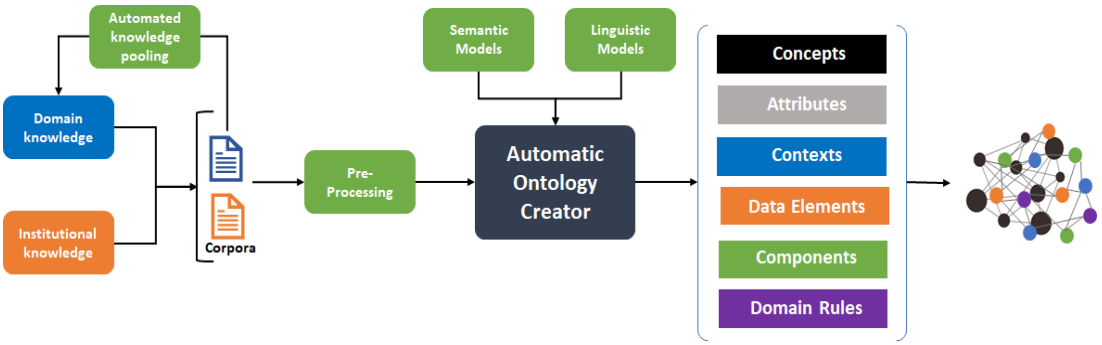

Fig: Process for creating automated ontology

The answer to scale and relevance? Automated Ontology

Static ontology frameworks are an impediment when attempting to launch new products or to comply with a new regulation at the speed of business today. However, creation of self-generative ontology enhances the knowledge discovery exercise to an unprecedented new level. Why?

Automated Ontology:

1. Allows for creation of a continuously evolving Enterprise Knowledge Graph, inclusive of all logical concepts.
2. Increases the accuracy of ML and AI algorithms by providing a dynamic knowledge-base that would be far superior to static frameworks.
3. Assists in the evaluation of generic knowledge alongside the knowledge derived from an institution’s corpora, thereby enriching institution specific data with definitions and their dependencies.

4. Associates business concepts derived from the knowledge graph to the physical data for solving various entity linking challenges.

Now that we have built our complete and ever-refined ontology, what can be done with it?

Applying Ontology to solve real life problems

Applications and queries that can be built utilizing ontology are numerous. Consider this; the risk or modeling team of a bank is tasked with defining the process required to estimate the expected loss for a customer. As a part of this process, the modeler requires specific data to estimate the of Probability of Default (PD), Loss Given Default (LGD) and Exposure at Default (EAD) values. Based on the approach discussed in the previous section, an ontology is created from the bank's internal credit-risk estimation model documents and other regulatory documents which define the guidelines to estimate credit-loss. Now, a modeler can get actionable output from the automatically generated ontology, which has the knowledge relevant to various concepts, data elements and components involved in credit-loss estimation along with semantic clustering and classification of these entities.

Conclusion

In this article we endeavor to provide the reader with a brief introduction to automated ontology and how it tackles the challenge of enterprise scale and how it is applicable within mainstream business use cases. Even though it is impossible to fully describe each of the different methods and algorithms thoroughly, it should provide an initial overview of our progress in the field of automated ontology creation.

About us

Parabole is a Princeton, NJ based cognitive analytics company, automating the creation of enterprise knowledge from unstructured sources of information. We provide our clients with platform to solve risk, finance and regulatory compliance-related challenges that depend on the knowledge-data interchange. We accomplish this by delivering a range of bespoke applications in the areas of risk, compliance, treasury, marketing, legal and data governance domains.

To learn more, visit www.parabole.ai OR reach out to info@parabole.ai