

A Beginner's Guide to Reasoning: How to Reason Your Way to Better Data*

Valerio Cocchi^{1,*}, Thomas Vout¹

¹*Oxford Semantic Technologies, Oxford, UK*

Abstract

Reasoning has become an increasingly valued tool in the semantic web space, and yet to many it's still a black box solution. Perhaps more tragically, despite the explosion of its development in recent years, many in the space still perceive it as a slow, cumbersome, and ultimately impractical technology, which is far from true today. Whether you're looking to harness reasoning for your own goals, or to peek behind the curtains of someone else's solution, now is your time to learn. Get hands on with a reasoning engine in this interactive walkthrough: A Beginner's Guide to Reasoning. You'll come away understanding the power of reasoning, what it can add to your data, and the fundamentals of how to apply it yourself. With technology in this space running away, there's never been a better time to learn! This tutorial will touch on the basics of SPARQL, OWL, and Datalog, before diving into reasoning at a technical level. Each participant will come away having built a reasoning solution for themselves, guided along the way by knowledge engineers and subject experts. No prior knowledge is required.

1. Motivation

Reasoning is one of the most powerful tools available to those utilizing the semantic web but it has been consistently misunderstood and undervalued as a result of a difficult history. For many years it struggled to bear fruit as a slow technology that failed to meet expectations, and understandably garnered an appropriate reputation. Today however, this has all changed.

Last year Prof. Ian Horrocks was awarded the Lovelace Medal for his work on reasoning systems, work that propelled reasoning out from being a fanciful daydream and into reality, now capable of delivering on the promises of the past. This is the technology that we bring you. Fast, stable, and with absolute correctness, reasoning is being adopted more and more in industry. The time has come for this tool to be widely understood by the community at large, not just as a shallow awareness, but as a deep technical familiarity.

2. Description

In this tutorial, attendees will learn what reasoning is, what it can do to transform data, and crucially, how to do that themselves. After a brief theoretical introduction to the subject,

ESWC'23: The 20th Extended Semantic Web Conference, May 28–June 2, 2023, Hersonissos, Greece

*Corresponding author.


†These authors contributed equally.

✉ valerio.cocchi@oxfordsemantic.tech (V. Cocchi); tom.vout@oxfordsemantic.tech (T. Vout)

🌐 <https://www.oxfordsemantic.tech/> (V. Cocchi); <https://www.oxfordsemantic.tech/> (T. Vout)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

attendees will be expected to get hands on with a reasoning engine in order to learn how to setup a reasoning- ready datastore, and how to reason over it using Datalog rules. No prior experience is required as we will run through the process step-by-step, start to finish.

Over the course of the tutorial attendees will learn:

- How to write and run a SPARQL query
- The importance of reasoning and its application
- How to write a Datalog rule
- How to apply and verify the rules they write
- The extent opportunities with reasoning
- How to create a solution that relies on reasoning

Each topic will first be demonstrated to the students so they can copy an ideal example and see the intended results in an informal teacher-student format. Then they will be given the opportunity to apply their new learned skills without immediate direction, writing rules and queries by themselves. If at any point a participant requires some assistance, the lecturers will be on hand to help, whether that requires a minor hint, a refresh of the material, or gentle guidance. Anyone of any skill level should leave this class knowing what reasoning is and how to implement it, so individual support is flexible depending on the needs of the group.

3. Tutorial material

- A slide deck, including educational content, exercises, and instructions, all made available in advance.
- Data, made available on the day of the tutorial.

4. Audience

We have run similar tutorials before, having between 10 and 20 attendees each time. They have been praised widely with past attendees giving us an average of 4.5 stars of a possible 5. These tutorials scale well and are appropriate regardless of the group size. They are aimed at a spectrum of users, from the non-technical to technically minded but unfamiliar. We cover the basics so that everyone has the tools to follow along with the rest of the tutorial but move quickly onto the more advanced sections. Experts in the field will not find this useful.

5. Software requirements

- Windows 8+, Mac OS 10.14+, Centos 7+, Ubuntu 16.04+
- RDFox downloaded
- An RDFox license, free for attendees