

My AGI story

Tom Rochette <tom.rochette@coreteks.org>

February 6, 2021 — [9686e64b](#)

0.1 Context

0.2 Learned in this study

0.3 Things to explore

1 Overview

In January 2014, I decided I wanted to read *Artificial Intelligence: A Modern Approach* from cover to cover. I don't really remember why I decided I wanted to do that, specifically because I didn't really have a deep interest in artificial intelligence. I did however end up reading the whole book within the month of January, taking about 57h of my time. For someone holding a full time job at the time, that was a lot of time dedicated to a given topic.

I wanted at the time want to start a business. I think I was mostly interested in the idea of intelligence amplification, something I had read about through Douglas Engelbart articles.

My shift toward AGI started around the summer of 2015. At the time, I joined the #ai Freenode channel on IRC and uttered my first comment there. I started reading "The Essential Turing", specifically its article "On Computable Numbers, with an Application to the Entscheidungsproblem", which at the time I mostly did not understand. For a few months I was obsessed about set theory because I thought you could link anything to a number, and then using set theory, associate things with one another. I was thinking about things such as how one could feed in data and have this data compressed in some way such that, for instance, a video of a man kicking a ball could be linked to its textual description.

From summer of 2015 to the end of 2016 I was mostly interested in AGI-related literature. I rarely read papers on the topic of machine learning, except the very popular ones such as "Playing Atari with Deep Reinforcement Learning", "Mastering the Game of Go with Deep Neural Networks and Tree Search" and "Neural Turing Machines" which I did not really understand at the time.

"The Society of Mind" (12h+ (no precise measurement))

"Cognitive Science - An Introduction to the Science of the Mind" (~32h)

In the summer of 2016, and until the end of the year, I spent a good amount of my time reading the book "Neuroscience: Exploring the Brain". The book took me approximately 48h to read, but compared to AI:AMA, the reading was spread over a much longer period of time (5 months vs 1 month). Reading that book was an extremely valuable experience. It allowed me to learn a lot about how the brain works, how neurons work, how different senses transduce environmental signals into neural signals and so forth.

At the same time I was reading "A Concise Introduction to Logic" (~34h). I was hoping that by reading this book I'd have

"Introduction to Automata Theory, Languages, and Computation" (~19h)

"Theory of self-reproducing automata" (18h+)

"Software Testing: A Craftsman's Approach" (~14h)

“Programming the Semantic Web” (~9h)

“Deep Learning” (~32h)

By the time I am writing this (July 2017), I have dedicated about 190 workdays (~8h/day) to the topics of artificial intelligence, artificial general intelligence and machine learning.

240h of theory (mostly book reading)

202h of machine learning (reading part of the ML book, the whole DL book, watching David Silver reinforcement learning course as well as Hugo Larochelle Traitement automatique des langues)

193h of research

134h on v1 (59h) and v2 (75h) of my AGI prototypes (in PHP and C# respectively)

123h of mathematics (Project Euler and reading A Concise Introduction to Logic)

95h on automated research

92h of machine learning on handwriting recognition

92h on a list manager application

86h on code analysis

70h on a knowledge base application

65h on neuroscience

40h on computer science

35h on cognition

29h on automated testing

2 See also

- [My path to AGI](#)

3 References