

Extracting someone's beliefs from their writings

Tom Rochette <tom.rochette@coreteks.org>

November 2, 2024 — [36c8eb68](#)

1 Question

Is it possible to extract someone's beliefs by reading their writings?

2 Answer

Yes.

If someone has written a blog for example, it is possible to figure out a variety of information about them simply based on what is written in those articles.

If someone writes about a specific software, we may not be able to infer right away what they think about said software, but we know that they've spent enough time to learn a bit about this specific software. If we know all the other alternatives in this category of software, we could infer that the user believes that this software is possibly better than the alternatives, otherwise why would they have picked it?

If the writer writes a lot on a topic, that is also a clue about their beliefs. They probably think that this topic is important, hence why they write about it. Maybe they write about this topic because it is lucrative to them.

What the writer doesn't write about is also informative. If they write mostly about technology, maybe they don't care about politics or sports?

It is possible to extract if-then rules from their writings, which generally expresses some form of believe that **if** something **then** something else. There are other variants where only **if** is provided (**if** something, something else, or something else, **if** something).

A writer may use certain adjectives to describe things as "easy", "simple", "straightforward", "difficult", "impossible", "hard", etc. Those are also useful of indicators of the writer's beliefs.

Nowadays with LLMs it's easy to provide a prompt such as "What beliefs are expressed in the following text" or "Extract the beliefs in the following text". With a bit of scripting we can extract everything available on the blog and feed it to a LLM to let it extract the beliefs for us.