

Text Classification

What does it do?

Text classification algorithms let you take a small labeled set of data and extrapolate those labels to a larger

What's so great?

It can help you find a needle in a haystack, theoretically; classify texts as protestant sermons and find more protestant sermons out there.

Authorship attribution is one subset of this that has sometimes been a particularly active area of research.

What decisions do you have to make?

What are the “features” the algorithm will use? Usually, it's words; but you can throw other elements into the mix.

What implementation should you use?

There are two choices to make here: the algorithm, and the software.

In terms of algorithms: if you want great performance, talk to a computer scientist. Unlike many other things digital humanists do, this directly maps to their areas of research. If you can't find one, there are several sensible approaches to try out. “K-nearest-neighbor” approaches find the most similar documents, and logistic regression uses words to adjust probabilities of being in class or another.

In terms of software; you will have to use a computer language to do this well, because the first run won't work. Python or R might both work: for certain goals like authorship attribution, the `stylo` package in R already includes a number of these algorithms specifically for text.