

Seminar on: Author Profiling

L715/B659
Dept. of Linguistics, Indiana University
Fall 2016

Author Profiling

Goals
Topics
Where we're going

Goals of this class

Goal:

- ▶ Based (solely) on the linguistic properties of a text, provide some characteristic of the *writer*
 - ▶ ≠ text classification, where the goal is to characterize the *text*
 - ▶ characteristic = age, native language, identity, etc.
- ▶ General approach: input features of the text into a classifier (supervised learning)

Author Profiling

Goals
Topics
Where we're going

◀ ▶ ⏪ ⏩ ⏴ ⏵ ⏶ ⏷ ⏸ ⏹ ⏺ ⏻ ⏼ ⏽ ⏾ ⏿ 🔍 ↺ ↻

1 / 12

◀ ▶ ⏪ ⏩ ⏴ ⏵ ⏶ ⏷ ⏸ ⏹ ⏺ ⏻ ⏼ ⏽ ⏾ ⏿ 🔍 ↺ ↻

2 / 12

Goals of this class

Particular goals of this seminar:

- ▶ See the connections between what are at times disparate topics & different fields/backgrounds
- ▶ Understand the theoretical underpinnings of making such inferences
 - ▶ e.g., explore a range of linguistic features
- ▶ Obtain practice building a practical system

More generally, my goal in a seminar is to see you learn to:

- ▶ Develop as researchers
 - ▶ Collaborate together on fun topics
- ⇒ Think of this more like a research lab than a class

Author Profiling

Goals
Topics
Where we're going

Specific topics

The specific topics will include:

- ▶ Authorship attribution
- ▶ Gender & age prediction
- ▶ Political affiliation identification
- ▶ Personality classification
- ▶ Native language identification
- ▶ Perhaps other topics ...

Part of that will depend upon interest ...

Author Profiling

Goals
Topics
Where we're going

◀ ▶ ⏪ ⏩ ⏴ ⏵ ⏶ ⏷ ⏸ ⏹ ⏺ ⏻ ⏼ ⏽ ⏾ ⏿ 🔍 ↺ ↻

3 / 12

◀ ▶ ⏪ ⏩ ⏴ ⏵ ⏶ ⏷ ⏸ ⏹ ⏺ ⏻ ⏼ ⏽ ⏾ ⏿ 🔍 ↺ ↻

4 / 12

Authorship attribution

Authorship attribution: who is the author of a document?

- ▶ Given a set of authors, which one (if any) wrote a particular work?
- ▶ **Authorship verification:** did a particular person write a particular work or not?
 - ▶ More challenging problem
 - ▶ Connects to forensic linguistics & topics like plagiarism detection

<http://pan.webis.de/clef16/pan16-web/author-identification.html>

Author Profiling

Goals
Topics
Where we're going

Gender & age prediction

Author profiling: categorize authors based on a demographic property, such as:

- ▶ Sex/gender
- ▶ Age
- ▶ Political persuasion
- ▶ ...

Task is to take a document (or set of documents) and determine a category (e.g., "18–24 years old")

This is commonly of interest for work in social media, where demographic information can help with product assessment

- ▶ Gender & age prediction have been the most popular classifications
- ▶ ... and the ones with longer histories in sociolinguistics

<http://pan.webis.de/clef16/pan16-web/author-profiling.html>

Author Profiling

Goals
Topics
Where we're going

◀ ▶ ⏪ ⏩ ⏴ ⏵ ⏶ ⏷ ⏸ ⏹ ⏺ ⏻ ⏼ ⏽ ⏾ ⏿ 🔍 ↺ ↻

5 / 12

◀ ▶ ⏪ ⏩ ⏴ ⏵ ⏶ ⏷ ⏸ ⏹ ⏺ ⏻ ⏼ ⏽ ⏾ ⏿ 🔍 ↺ ↻

6 / 12

Political affiliation identification

- Author Profiling
- Goals
- Topics
- Where we're going

Political affiliation is potentially:

- ▶ more overt
- ▶ more transient
- ▶ more reflective of one's values than one's social categories

The context of use (social media, blog, editorials, congressional proceedings) will have a big impact

Personality classification

- Author Profiling
- Goals
- Topics
- Where we're going

Personality classification also gets into issues of:

- ▶ transience: personality vs. mood?
- ▶ context of use: when do people express personality?
- ▶ defining a gold standard: what are the personality categories to classify?

Can be related to deception detection

see, e.g., <http://www.secretlifeofpronouns.com>

Native language identification

- Author Profiling
- Goals
- Topics
- Where we're going

Native language identification: identify native language (L1) of writer based on second language (L2) writing

- ▶ Topic may be more specific to language usage:
 - ▶ i.e., classifying a property of the writer's **language**, not a non-linguistic demographic
 - ▶ similar in some ways to dialect identification or language proficiency classification

Many language learners have some similar patterns, e.g., (non)usage of articles

- ▶ But the entirety of their patterns tend to differ
- ▶ Data size is thus an issue (as with all these topics)

<https://sites.google.com/site/nlsharedtask2013/home>

Other topics?

- Author Profiling
- Goals
- Topics
- Where we're going

And are there are other topics to examine?

- ▶ I don't know, but any (more or less permanent) demographic is in principle possible
 - ▶ where a demographic property is an inherent or cultural marker of the person
- ▶ e.g., income, (language) disability, education, various preferences, ...
 - ▶ nb: something like mood (or even political affiliation) is more fluid (a spectrum)

Fishing for some of these could lead to accidental correspondence, but we should be curious to find anything where:

- ▶ Input = natural language text(s)
- ▶ Output = classification of the writer, in terms of some demographic property

Where we're going

- Author Profiling
- Goals
- Topics
- Where we're going

For the next two-ish weeks, we'll make sure we're more or less on the same page:

- ▶ Machine learning & text classification techniques
- ▶ Natural Language Processing (NLP) tools
 - ▶ Our focus will be largely practical

After that, it'll be less me & more you leading discussion

- ▶ By Tuesday, September 6, I want you to sign up to lead the discussion on a particular topic.
 - ▶ Lead ≠ Be an expert on
 - ▶ Lead = read papers ahead of time, help determine the interesting areas to explore, come to class with the interesting points identified, etc.
- ▶ I'll give you a specific assignment on this next time

Expectations

- Author Profiling
- Goals
- Topics
- Where we're going

This seminar is going to be:

1. Exploratory & interactive: think of this as collaborative learning in a lab-like environment
2. Demand-driven: I have sketched a syllabus; the contents will be driven in part by your interests
 - ▶ Note also that I have organized the content by task, but you may prefer looking across tasks at methods, features, etc.