


Applied Scaling & Classification Techniques in Political Science

Lecture 6 – part 2



Supervised classification methods
with human tagging: an introduction



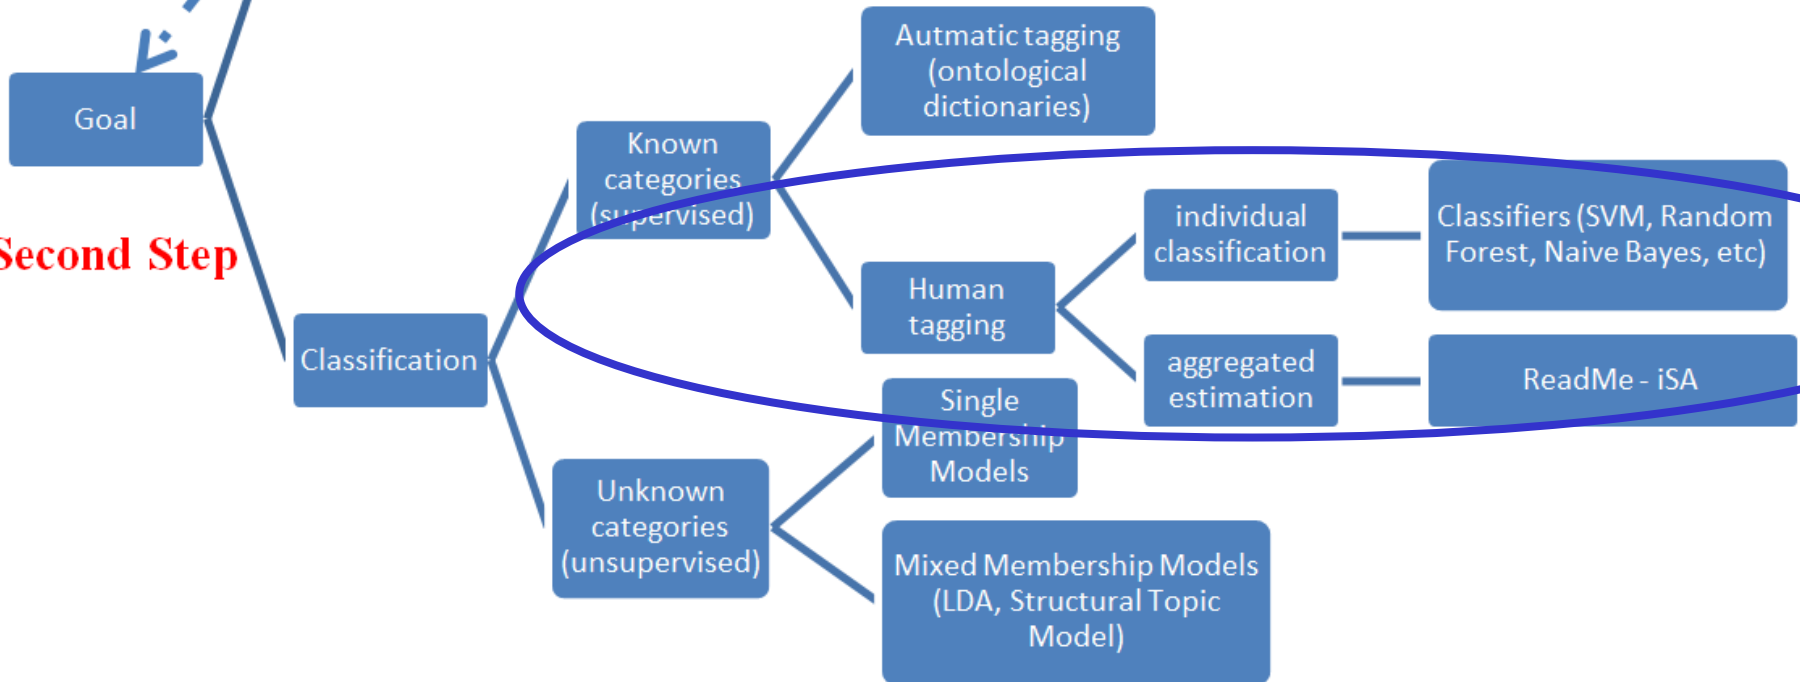
Our Course Map



First Step



Second Step






References

- ✓ Grimmer, Justin, and Stewart, Brandon M. (2013). Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political Analysis*, 21(3): 267-297
- ✓ Curini, Luigi, and Robert Fahey (2020). Sentiment Analysis and Social Media. In Luigi Curini and Robert Franzese (eds.), *SAGE Handbook of Research Methods in Political Science & International Relations*, London, Sage, chapter 29
- ✓ Barberá, Pablo et al. (2020) Automated Text Classification of News Articles: A Practical Guide, *Political Analysis*, DOI: 10.1017/pan.2020.8

Supervised Learning (classification) Methods




The idea of supervised learning is simple: human coders categorize a set of documents (the “**training-set**” or “**labelled-set**”) by hand


The algorithm “learns” how to sort the documents into categories using the **training set and words**

Then, it classifies the remaining set of document not classified by hand (the “**test-set**” or “**unlabelled set**”) using the characteristics of the documents to place them into the categories

A three-steps procedure





1. Data filtering




2. Human classification



A three-steps procedure



Supervised Learning vs. Dictionary methods



This approach to classification has **four major advantages** over dictionary methods:

First, it is **necessarily domain specific** and therefore avoids the problems of applying dictionaries outside of their intended area of use

Applying supervised learning methods **requires scholars to develop coding rules for the particular quantities of interest** under study when working on the training-set

This also forces scholars to develop coherent definitions of concepts for particular applications, which leads to clarity in what researchers are measuring and studying

Supervised Learning vs. Dictionary methods



Second, when using a dictionary you are relying **only** on the words (and on their relative weight: +1 or -1 for example) as defined in the dictionary

That is...the dictionary comes with a hard-wired set of parameter values for the importance of a predetermined set of features

In contrast, when using a supervised approach the relevant features of the text and their weights are estimated directly from the data (in the training set). The feature space is thus likely to be both larger and more comprehensive than that used in a dictionary

The end result is that much more information drives the subsequent classification of text

Supervised Learning vs. Dictionary methods




In particular, the supervised learning model will estimate parameter values optimized to minimize error of the classifier on the training dataset


Thus, they will necessarily outperform the dictionary on that **same sample**, as long as training sample is large enough

And indeed...

Supervised Learning vs. Dictionary methods



Lexicons' Accuracy in Document Classification Compared to Machine-Learning Approach



Source: Gonzalez-Bailn and Paltoglou (2015)

Supervised Learning vs. Dictionary methods



But which does better **out-of-sample**?

Here, too, since a supervised learning model is trained on a sample of the data, it is guaranteed to do better than a dictionary as long as it is trained on a **large enough random sample**

Indeed, as the sample converges to the population - or as the training dataset contains an ever increasing proportion of words encountered - a supervised learning model has to do better than a dictionary, as the estimated parameter values will converge to the true parameter values


Supervised Learning vs. Dictionary methods



Third, human involvement is crucial to understand the correct meaning of a text (double meaning sentences, specific jargons, neologisms, irony)

Fourth, supervised learning methods are much easier to validate, with clear statistics that summarize model performance (more on this later on)

Supervised Learning vs. Dictionary methods



While a dictionary cannot compete with a classifier trained on a **representative and large enough training dataset**, in any given task dictionaries may however outperform a supervised learning model **if these conditions are not met**

Dictionaries bring rich prior information to the classification task: humans may produce a topic-specific dictionary that would require a large training dataset to outperform it

Supervised Learning vs. Dictionary methods



Summing up:

Dictionaries:

- ✓ Can be off the-shelf
- ✓ no creation of a training dataset required
- ✓ easy to apply to a given corpus
- ✓ built by humans who can bring domain expertise to bear

Supervised Learning vs. Dictionary methods



Summing up:

Supervised machine learning:

- ✓ optimized for current research question
- ✓ more comprehensive set of features used to classify text
- ✓ mathematically, SML necessarily outperforms dictionary methods given a large enough training dataset
- ✓ built-in validation mechanism vis a vis human coding: by construction, the analyst knows the performance of the classifier based on multiple measures of fit (i.e, how closely the labels generated correspond to human coding)

Supervised Learning vs. Dictionary methods



Summing up:

Supervised learning can be therefore conceptualized as a **generalization of dictionary methods**, where features associated with each categories (and their relative weight) are learned from the data **via human intervention**



Constructing a training set

For supervised problems, the researcher is aiming to classify documents into a set of known or assumed categories based upon rules or information that can be learned from the training set

This requires **labels** in the training set from which to infer categories in the test set

The most important step in applying a supervised learning algorithm is therefore constructing a **reliable training set**, because no statistical model can repair a poorly constructed training set!

If the training set is **poorly constructed**, the supervised algorithms will simply replicate such poorly construction!

Constructing a training set



(1) creating and executing a coding scheme:

Best practice is to **iteratively develop coding schemes**

Initially, a **concise codebook** is written to guide coders, who then apply the codebook to an initial set of documents

When using the codebook, particularly at first, coders are likely to **identify ambiguities** in the coding scheme or overlooked categories

Constructing a training set



(1) creating and executing a coding scheme:

This subsequently leads to a **revision of the codebook**, which then needs to be applied to a new set of documents to ensure that the ambiguities have been sufficiently addressed

Only after coders apply the coding scheme to documents without noticing ambiguities is a “final” scheme ready to be applied to the data set

Constructing a training set



(2) sampling documents:

Almost all (but not all...) classification methods implicitly assume that the **training set is a random sample** from the population of documents to be coded

This is because Supervised learning methods **use the relationship** between the features in the training set to classify the remaining documents in the test set (out-of-sample predictions)

Constructing a training set



(2) sampling documents:

This presents particular difficulty when...

...**all the data are not available at the time of coding:**

either because it will be produced in the future or because it has yet to be digitized

Constructing a training set



This is due to ***semantic change***, which is the difference in the meaning of language between the labeled and unlabeled sets

For example, we can have **emergent discourse**, where new words and phrases, or the meanings of existing words and phrases, appear in the unlabeled set but not the labeled set, and **vanishing discourse**, where the words, phrases, and their meanings exist in the labeled set but not the unlabeled set

Constructing a training set



Russian election hacking after 2016 US Presidential elections is an example of emergent discourse, language which did not exist a few years ago, whereas *Russian Communism* is an example of vanishing discourse, with language that has largely vanished from ongoing conversations over time

How to face this risk?

Keep updating the training-set (if your test-set is still to come...)!

Constructing a training set



(2) sampling documents:

Moreover, Supervised methods need **enough information** to learn the relationship **between words and documents in each category of a coding scheme**

Hopkins and King (2010) offer **five hundred** as a rule of thumb with one hundred documents probably being enough

Constructing a training set



(2) sampling documents:

Still the number necessary will depend upon **the specific application of interest**. For example, as the number of categories in a coding scheme increases, the number of documents needed in the training set also increases

Moreover, if a category does not occur, or occurs extremely rarely, in the training set, there is insufficient opportunity to “learn” about this category and its properties, which will in turn interfere with the process of classifying test-set documents into this category correctly

Constructing a training set



(2) sampling documents:

When attempting to detect **small changes or rare categories**, *therefore*, increasing the probability that they are observed in the training set often means increasing the size of the training set relative to the test set

Constructing a training set



(3) **checking human-tagging reliability:**

Manually generating the initial set of labels can prove arduous and time-consuming, but is also fraught with concerns about consistency and accuracy

That is, while labeling training data often requires the use of human coders to sort texts into desired categories, human coding lacks consistency and reliability both within and across individuals, above and beyond the time and expense required to complete the task

Therefore always run an **inter-coder reliability test!!!**


Constructing a training set



Given the likely existence of a budget constraint sometimes we will need to make a **choice between more coders per object and more objects coded**. What is better?

The literature shows that while increasing the number of coders for each document can improve the accuracy of the classifier, the informational gains from increasing the number of documents coded are greater than from increasing the number of codings of a given document

Constructing a training set



This does not obviate the need to have multiple coders code at least a subset of documents, namely to determine coder quality and to select the best set of coders to use for the task at hand

But once the better coders are identified, the optimal strategy is to proceed with one coder per document