

Text Processing, Tasks, and Corpora

Nathan Schneider ~ 14 January 2025

<https://people.cs.georgetown.edu/nshneid/cosc5402/>

Text Processing

Basics of working with text

- Computers represent text as strings of **characters**
 - Letters, numbers, punctuation, symbols, emojis, diacritics/combining characters, spaces
- At the machine level, a string is stored as a sequence of **bytes**, which are interpreted as integers → code points → characters in some **encoding**
 - e.g. Unicode defines 0x0B90 to be “ஐ” (Tamil letter Ai)
 - UTF-8 is a way of encoding Unicode characters as bytes
- A group of characters used to write in a language is called a **script**
 - e.g., Latin characters for European languages; Hangul for Korean; Devanagari for Hindi; CJK Ideographs
- **Text editors** provide functionality for working with (plain) text files

Find & replace

- In a text editor, or in Python, there are ways to find/replace exact string matches
- If you don't know the exact string you're looking for, you can use a **regular expression**


- In a text exact s

WHENEVER I LEARN A NEW SKILL I CONCOCT ELABORATE FANTASY SCENARIOS WHERE IT LETS ME SAVE THE DAY.





replace

- If you don't use a regular



you can



Find & replace

- In a text editor, or in Python, there are ways to find/replace exact string matches
- If you don't know the exact string you're looking for, you can use a **regular expression**
- Regular expression patterns allow for **disjunctions** and **wildcards**, so the pattern may match more than one particular string
 - ▶ Pattern to find all English articles ("a", "an", "the")?
 - ▶ Pattern to check whether a string is an email address?

Regular expressions: (pro|con)s

- Simple pattern language.
 - (though some differences by implementation, e.g. Python vs. Perl)
- Deterministic—the expression either matches the string or it doesn't!
 - No ambiguity in the correct behavior (of course your expression may have a bug)
- Requires a very precise statement in terms of characters/substrings.
 - “names of people”, “sentences about chemistry” would be hard to express as a regex

Documents/sentences/words

- For more sophisticated NLP, we often want to break documents down into sentences, and sentences into words.
 - (In linguistics, “word” is very hard concept to define precisely across languages. But for now, think roughly “groups of characters separated by spaces in English text”.)
- (Word) tokenization and sentence tokenization
 - The split units are called **tokens**.
 - *Token* also refers to an instance of a word, in contrast to its **type** (the word in general).

Tokenization

Daniels made several appearances as C-3PO on numerous TV shows and commercials, notably on a Star Wars-themed episode of The Donny and Marie Show in 1977, Disneyland's 35th Anniversary.



Daniels made several appearances as C-3PO on numerous TV shows and commercials, notably on a Star Wars-themed episode of The Donny and Marie Show in 1977, Disneyland's 35th Anniversary.

- For most kinds of writing, regular expressions can do this pretty well, most of the time.
 - Language/regional/stylistic differences in use of spaces, punctuation, etc.
 - Python library for English: `nltk.word_tokenize(string)`

Sentence Tokenization/ Splitting: Not always easy

In April 1938, Bernarr A. Macfadden, publisher of *Liberty* magazine stepped in, offering a prize of \$1,000 to the winning composer, stipulating that the song must be of simple “harmonic structure”, “within the limits of [an] untrained voice”, and its beat in “march tempo of military pattern”.

The contest rules required the winner to submit his entry in written form, and Crawford immediately complied. However his original title, *What Do You think of the Air Corps Now?*, was soon officially changed to *The Army Air Corps*.

[https://en.wikipedia.org/wiki/The_U.S._Air_Force_\(song\)](https://en.wikipedia.org/wiki/The_U.S._Air_Force_(song))

Preprocessing

- Tokenization is a kind of **preprocessing** that can be done to a piece of text.
- Depending on the data source and application, it may help to perform
 - **data cleaning:** remove irrelevant material like page numbers in an OCR'd text, or irrelevant documents like spam/advertising
 - **normalization:** collapsing different types of spelling variation such as capitalization (colour → color; U.S. → US → us)
 - **lemmatization:** mapping words of a language to a canonical form (the “dictionary entry” or **lemma**): {eat, eats, eating, ate, eaten} → eat

Tasks


- Recall that NLP covers a range of problems, ranging from core linguistic analysis to practical applications:

Applications & Core Tasks

- Machine Translation
 - Information Retrieval
 - Question Answering
 - Dialogue Systems
 - Information Extraction
 - Summarization
 - Sentiment Analysis
 - ...
- Language modeling/text generation
 - Part-of-speech tagging
 - Syntactic parsing
 - Named entity recognition
 - Coreference resolution
 - Word sense disambiguation
 - Semantic role labeling
 - ...

- We measure progress in NLP by evaluating systems designed to perform various **tasks**.
 - benchmarking
- Different tasks target different languages, parts of language, and applications.
- By way of example, we'll focus on the task of...

Sentiment Analysis



Goal: Predict the **opinion** expressed in a piece of text.

E.g., + or -. (Or a rating on a scale.)

Sentiment Analysis

Filled with horrific dialogue, laughable characters, a laughable plot, and really no interesting stakes during this film, "Star Wars Episode I: The Phantom Menace" is not at all what I wanted from a film that is supposed to be the huge opening to the segue into the fantastic Original Trilogy. The positives include the score, the sound effects, and most of the

KJ Proulx (/user/id/896976177/)

★ Super Reviewer

Extraordinarily faithful to the tone and style of the originals, The Force Awakens brings back the Old Trilogy's heart, humor, mystery, and fun. Since it is only the first piece in a new three-part journey it can't help but feel incomplete. But everything that's already there, from the stunning visuals, to the thrilling action sequences, to the charismatic new characters,

Matthew Samuel Mirliani (/user/id/896467979/)

★ Super Reviewer

Sentiment Analysis



Filled with horrific dialogue, laughable characters, a laughable plot, and really no interesting stakes during this film, "Star Wars Episode I: The Phantom Menace" is not at all what I wanted from a film that is supposed to be the huge opening to the segue into the fantastic Original Trilogy. The positives include the score, the sound effects, and most of the

KJ Proulx (/user/id/896976177/)

★ Super Reviewer



Extraordinarily faithful to the tone and style of the originals, The Force Awakens brings back the Old Trilogy's heart, humor, mystery, and fun. Since it is only the first piece in a new three-part journey it can't help but feel incomplete. But everything that's already there, from the stunning visuals, to the thrilling action sequences, to the charismatic new characters,

Matthew Samuel Mirliani (/user/id/896467979/)

★ Super Reviewer

Sentiment Analysis

★★

Filled with horrific dialogue, laughable characters, a laughable plot, and really no interesting stakes during this film, "Star Wars Episode I: The Phantom Menace" is not at all what I wanted from a film that is supposed to be the huge opening to the segue into the fantastic Original Trilogy. The positives include the score, the sound effects, and most of the

KJ Proulx (/user/id/896976177/)

★ Super Reviewer

★★★★★

Extraordinarily faithful to the tone and style of the originals, The Force Awakens brings back the Old Trilogy's heart, humor, mystery, and fun. Since it is only the first piece in a new three-part journey it can't help but feel incomplete. But everything that's already there, from the stunning visuals, to the thrilling action sequences, to the charismatic new characters,

Matthew Samuel Mirliani (/user/id/896467979/)

★ Super Reviewer

[RottenTomatoes.com](https://www.rottentomatoes.com) + intuitions about positive/negative cue words

So, you want to build a sentiment analyzer

Questions to ask yourself:

1. What is the input for each prediction? (sentence? full review text? text+metadata?)
2. What are the possible outputs? (+ or -)
3. How will it decide?
4. How will you measure its effectiveness?

The last one, at least, requires data!

BEFORE you build a system, choose a dataset for evaluation!

Why is data-driven evaluation important?

- Good science requires controlled experimentation.
- Good engineering requires benchmarks.
- Your intuitions about typical inputs are probably wrong.

Sometimes you want multiple evaluation datasets: e.g., one for **development** as you hack on your system, and one reserved for final **testing**.

Corpora

- A dataset with many instances of language use is called a **corpus** (pl. *corpora*).
 - ▶ < Latin 'body'
 - ▶ It can be a carefully curated sample for studying how language works.
 - ▶ Or, for building and evaluating NLP systems.
- Where can you get one?
 - ▶ A precompiled dataset, e.g. for NLP/machine learning contests
 - ▶ Create your own, e.g. by scraping websites

ChatGPT Activity




ChatGPT

- ChatGPT is a well-known system that responds to text prompts with generated text (or code).
 - It is a testament to the power of corpus data.
- Who has used it? What for?


Strengths/Weaknesses

1. Ask some factual questions with specific correct answers (e.g. birthday of a celebrity). Does the system provide a correct answer?
2. Engage in a dialogue about a favorite book, film, or research paper of your choice.
 - Can the system provide an accurate summary?
 - Are the answers on topic? Consistent?
 - If you ask about details, are the answers correct?
 - What happens if you ask a trick question?
3. Ask it to provide information in a creative form, e.g. lyrics in the style of your favorite singer.

Nathan vs. ChatGPT



Caution

- **It is hard to avoid anthropomorphizing systems like ChatGPT.**
 - They produce fluent text, which is something we tend to associate with human intelligence.
 - The system was trained on massive amounts of text on the web produced by humans, and remixes them to generate new strings. “Stochastic ”
 - The system is impressive and (in some respects) useful!
- **But would you trust ChatGPT to make an important decision/document?**
 - Much of the generated text contains incorrect details, and it is easy to trick the system. Fluency ≠ a mind at work.
 - Beware industry hype! (“Generative AI” is the new “blockchain”)

“Intelligent” image generation?



YOU

Make a wheres Waldo picture



CHATGPT



Here's your "Where's Waldo" style picture. Have fun trying to spot the character!

ALT