

Basic Text Processing

- **Regular Expressions**
- **Text Normalization**
 - **Word Tokenization**
- **Subword Tokenization**

Basic Text Processing: **Regular Expressions**

Regular Expressions

- **Regular expressions** are the most important tool to describe text patterns and we can use them to specify the strings to be extracted from the text.
 - Regular expressions are widely used in many text preprocessing tasks.
 - A set of text preprocessing tasks is called as **text normalization**.
 - Normalizing text means converting it to a more convenient, standard form.

Tokenization: Separating out or **tokenizing** words from text.

- English words are often separated from each other by whitespace (not enough).
 - For processing tweets we'll need to tokenize emoticons like :) or hashtags like #nlproc.

Lemmatization: Task of determining that words have the same root.

- Words *sings*, *singing*, *sang*, *sung* have the same root word (**lemma**) *sing*.
- Words *kitabım*, *kitaplar*, ... have the same root word *kitap*.
- **Stemming:** a simpler version of lemmatization in which we mainly just strip suffixes from the end of the word.

Sentence Segmentation: breaking up a text into individual sentences.

Regular Expressions

- Each **Regular Expression (RE)** represents a set of strings having certain pattern.
 - In NLP, we can use REs to find strings having certain patterns in a given text.
- Regular Expressions are an algebraic way to describe formal languages.
 - Regular Expressions describe exactly the regular languages.
- A regular expression is built up of simpler regular expressions (using defining rules).
- Simple Definition for Regular Expressions over alphabet Σ
 - ϵ is a regular expression
 - If $a \in \Sigma$, a is a regular expression
 - **or** : If E_1 and E_2 are regular expressions, then $E_1 \mid E_2$ is a regular expression
 - **concatenation** : If E_1 and E_2 are regular expressions, then E_1E_2 is a regular expression
 - **Kleene Closure**: If E is a regular expression, then E^* is a regular expression
 - **Positive Closure**: If E is a regular expression, then E^+ is a regular expression

Searching Strings with Regular Expressions

(using Python style REs)

- How can we search for any of following strings?
 - woodchuck
 - woodchucks
 - Woodchuck
 - Woodchucks
- Python style regular expressions:
 - We write regular expressions in Python → `r"..."`
`r"regularexpression"`
- The **simplest kind of regular expression** is a sequence of simple characters.
 - The regular expression `r"b"` will match with the string "b".
 - The regular expression `r"bc"` will match with the string "bc".
 - The regular expression `r"woodchuck"` will match with the string "woodchuck".
 - The regular expression `r"woodchucks"` will match with the string "woodchucks".
 - The regular expression `r"woodchuck"` will NOT match with the string "Woodchuck".

Regular Expressions: Disjunctions

disjunction of characters []

- **Disjunction of Characters:** The string of characters inside the braces [] specifies a **disjunction** of characters to match.
- The regular expression `r"[wW]"` matches patterns containing either w or W.

Regular Expression	Matches
<code>r"[wW]oodchuck"</code>	Woodchuck or woodchuck
<code>r"[1234567890]"</code>	Any digit

- **Ranges in []:** If there is a well-defined sequence associated with a set of characters, **dash (-)** in brackets can specify any one character in a **range**.

Regular Expression	Matches
<code>r"[A-Z]"</code>	An upper case letter
<code>r"[a-z]"</code>	A lower case letter
<code>r"[0-9]"</code>	A single digit

Regular Expressions: Disjunctions

Negations in []

- **Negations in []:**

- The square braces can also be used to specify what a single character cannot be, by use of the caret ^.
- If the caret ^ is the first symbol after the open square brace [, the resulting pattern is negated.

Regular Expression	Matches
<code>r"[^A-Z]"</code>	Not an upper case letter
<code>r"[^a-z]"</code>	Not a lower case letter
<code>r"[^Ss]"</code>	Neither 'S' nor 's'
<code>r"[^e^]"</code>	Neither e nor ^
<code>r"a^b"</code>	The pattern a^b

Regular Expressions: Disjunctions

or (disjunction) operator | (pipe symbol)

- If E_1 and E_2 are regular expressions, then $E_1 | E_2$ is a regular expression

Regular Expression	Matches
<code>r"woodchuck groundhog"</code>	woodchuck or groundhog
<code>r"a b c"</code>	a , b or c
<code>r"[gG]roundhog [Ww]oodchuck"</code>	woodchuck , Woodchuck , groundhog or Groundhog

Regular Expressions: Closure Operators

Kleene * and Kleene +

- **Kleene * (closure) operator:** The Kleene star means “zero or more occurrences of the immediately previous regular expression.
- **Kleene + (positive closure) operator:** The Kleene plus means “one or more occurrences of the immediately preceding regular expression.

Regular Expression	Matches
<code>r"ba*"</code>	b, ba, baa, baaa, ...
<code>r"ba+"</code>	ba, baa, baaa, ...

Regular Expressions: {} . ?

- **{m,n}** causes the resulting RE to match from m to n repetitions of the preceding RE.
- **{m}** specifies that exactly m copies of the previous RE should be matched
- The question mark **?** marks **optionality of the previous expression**.

Regular Expression	Matches
<code>r"woodchucks?"</code>	woodchuck or woodchucks
<code>r"colou?r"</code>	color or colour
<code>r"b{2,3}"</code>	bb, bbb

- A wildcard expression **dot .** matches any single character (except a carriage return).

Regular Expression	Matches
<code>r"beg.n"</code>	begin, begun, begxn, ...
<code>r"a.*b"</code>	any string starts with a and ends with b

Regular Expressions: Anchors ^ \$

- **Anchors** are special characters that anchor regular expressions to particular places in a string.
- The **caret** ^ matches the start of a string.
 - The regular expression **^The** matches the word **The** only at the start of a string.
- The **dollar sign** \$ matches the end of a line.

Regular Expression	Matches
<code>r"\$. \$"</code>	any character at the end of a string
<code>r"\. \$"</code>	dot character at the end of a string
<code>r"^[A-Z] "</code>	any uppercase character at the beginning of a string
<code>r"^The dog\. \$"</code>	a string that contains only the phrase The dog.

Regular Expressions: Precedence of Operators

- The order precedence of RE operator precedence, from highest precedence to lowest precedence is as follows
 - Parenthesis `()`
 - Counters `* + ? {}`
 - Sequences and anchors `^ $`
 - Disjunction `|`
- The regular expression **the*** matches **theeee** but not **thethe**

Regular Expressions: backslashed characters

- Aliases for common sets of characters

RE	Expansion	Match
<code>\d</code>	<code>[0-9]</code>	any digit
<code>\D</code>	<code>[^0-9]</code>	any non-digit
<code>\w</code>	<code>[a-zA-Z0-9_]</code>	any word character
<code>\W</code>	<code>[^a-zA-Z0-9_]</code>	any non-word character
<code>\s</code>	<code>[\t\n\r\f]</code>	any whitespace character
<code>\S</code>	<code>[^ \t\n\r\f]</code>	any non-whitespace character

Regular Expressions: backslashed characters

- Special characters need to be **backslashed**.

RE	Match
<code>\b</code>	a word boundary: A word boundary is the position between a word character (<code>\w</code>) and a non-word character (<code>\W</code>).
<code>\B</code>	a non-word boundary: This is the opposite of <code>\b</code> , and it matches any position that is not a word boundary.
<code>\n</code>	a newline character
<code>\t</code>	a tab character
escaping other special characters: We can escape almost any special character (like <code>*</code> , <code>+</code> , <code>?</code> , <code>(</code> , <code>)</code> , <code>{</code> , <code>}</code> , <code>[</code> , <code>]</code> , <code> </code> , <code>^</code> , <code>\$</code> , <code>.</code> , and <code>\</code>) by preceding it with a backslash (<code>\</code>).	
<code>\. \+ * \\ \? \(\) \[\] \ \^ \\$</code>	

Regular Expressions: Example

- We want to write a RE to find cases of the English article **the**
 - We can use **findall** method in **re** library for tokenization.

```
import re
sentence = "The book and the other book are not theology books. We breathe the air."
re.findall(r"the", sentence)
→ ['the', 'the', 'the', 'the', 'the']
The book and the other book are not theology books. We breathe the air.

re.findall(r"[tT]he", sentence)
→ ['The', 'the', 'the', 'the', 'the', 'the']
The book and the other book are not theology books. We breathe the air.

re.findall(r"[tT]he\b", sentence)
→ ['The', 'the', 'the', 'the']
The book and the other book are not theology books. We breathe the air.


re.findall(r"\b[tT]he", sentence)
→ ['The', 'the', 'the', 'the']
The book and the other book are not theology books. We breathe the air.

re.findall(r"\b[tT]he\b", sentence)
→ ['The', 'the', 'the']
The book and the other book are not theology books. We breathe the air.
```

Regular Expressions & FSAs

- Any regular expression can be realized as a **finite state automaton (FSA)**
- There are two kinds of FSAs
 - **Deterministic Finite State Automatons (DFAs)**
 - Non-deterministic Finite State Automatons (NFAs)
 - Any NFA can be converted into a corresponding DFA.
- A DFA (and a regular expression) represents a **regular language**.

Regular Expressions




Finite Automata

Regular Languages

Regular Expressions: A DFA and A NFA

- **A regular language:** The strings whose second characters from the right end are 1.
- Regular Expression for this regular language: $(0|1)^*1(0|1)$



- A DFA for this language :

Formal Definition of Finite-State Automaton

- FSA is $Q \times \Sigma \times q_0 \times F \times \delta$
- Q : a finite set of N states $q_0, q_1, \dots, q_N$
- Σ : a finite input alphabet of symbols
- q_0 : the start state
- F : the set of final states
- $\delta(q,i)$: transition function
 - DFA : There is exactly one arc leaving a state q with a symbol a .
There is no arc with the empty string.

Basic Text Processing: **Text Normalization**

Text Normalization

- Almost every natural language processing task needs to do text normalization.
- Three tasks are commonly applied as part of any normalization process:
 1. Segmenting/tokenizing words from the text
 2. Normalizing word formats
 3. Segmenting sentences in the text.

Words

- Before processing words, we need to decide what counts as a **word**.
- How many words are in the following sentence?

He stepped out into the hall, was delighted to encounter a water brother.

- If we do NOT count punctuations as words → 13 words
 - If we count punctuations as words → 15 words
- Punctuations can be useful to identify boundaries of things and some aspects of meaning.
- Are capitalized tokens and uncapitalized tokens the same word?
 - The and the big possibly
 - US and us may be not (US: united states of America)

Words

- Are the inflected forms like **cat** and **cats** the same word?
- They have the same **lemma** **cat**, but they have different **wordforms**.
- A **lemma** is a set of lexical forms having the same stem, the same major part-of-speech, and the same word sense.
- The **wordform** is the full inflected or derived form of the word.
 - For morphologically complex languages, we often need to deal with lemmatization.
 - For many tasks in English, however, wordforms are sufficient.

Words: *How many words are there in English?*

- A **type** is a distinct Word in a corpus.
- **V: Vocabulary** is the set of types.
 - $|V|$ is the size of the vocabulary.
- Each **word** in a corpus is a **token**.
 - N is the number of tokens in the corpus.

Corpus	# of Tokens = N	# of Types = $ V $
Shakespeare	884,000	31 thousand
Switchboard phone conversations	2.4 million	20 thousand
Brown corpus	1 million	38 thousand
Google N-grams	1 trillion	13 million

Tokenization

- **Word Tokenization**

Word Tokenization and Normalization

- **Tokenization** is the task of segmenting the text into words.
- **Normalization** is the task of putting words in a standard format.
- We can use **regular expressions** to segment the text into words for **tokenization** task.
 - Since tokenization needs to be run before any other language processing, it is important for it to be very fast.
 - The method for tokenization/normalization is to use **deterministic algorithms based on regular expressions** compiled into very efficient **finite state automata**.

Tokenization

- Normally we want to break off **punctuations** as separate tokens, but sometimes we want to keep them in words internally.
- Punctuations as separate tokens: He ate apple, orange and banana.
- Punctuations kept internally:
 - m.p.h. Ph.D. AT&T Prices: \$43.55 Dates: 27/09/2019
 - URLs: <http://www.hacettepe.edu.tr> Twitter hashtags: #nlproc
- A tokenizer can also expand **clitic** contractions that are marked by apostrophes.
 - what're to two tokens what are
 - we're to two tokens we are
- Tokenization algorithms may also tokenize multiword expressions like New York or rock 'n' roll as a single token.

Tokenization in Python

- We can use methods in **re** or **nltk** library for tokenization.
 - We can use **findall** method in **re** library for tokenization.
 - We can use **word_tokenize** or **regex_tokenize** methods in **nltk** library for tokenization.
- Token patterns are described with regular expressions.

`re.findall(pattern, string)`

- Returns all non-overlapping matches of ***pattern*** in ***string***, as a list of strings or tuples.
- The ***string*** is scanned left-to-right, and matches are returned in the order found.
- The result depends on the number of capturing groups in the pattern.
 - If there are no groups, return a list of strings matching the whole pattern.
 - If there is exactly one group, return a list of strings matching that group.
 - If multiple groups are present, return a list of tuples of strings matching the groups.

`nltk.regex_tokenize(string, pattern)`

`nltk.word_tokenize(string)`

- `nltk.word_tokenize` may require to download some library.
 - `nltk.download('punkt_tab')`

Tokenization in Python

- In regular expressions, **parentheses** are used for two main purposes:
 1. **Grouping:** Parentheses (. . .) are used to group parts of the pattern together.
 2. **Non-Grouping:** Parentheses (?: . . .) are used to create non-capturing groups.
 - This group is still used to group parts of the pattern together (like regular parentheses), but it does not create a capturing group.

```
str = "the evening shows start at 7:00pm and 10:15pm. the morning show at 9:00."
```

```
# Parentheses identify a group within the pattern
```

```
matches = re.findall(r"(\d\d?:\d\d) (am|pm)?", str)
```

```
➔ [('7:00', 'pm'), ('10:15', 'pm'), ('9:00', '')]
```

```
matches = re.findall(r"(?:\d\d?:\d\d) (am|pm)?", str)
```

```
➔ ['pm', 'pm', '']
```

```
matches = re.findall(r"(\d\d?:\d\d) (?:am|pm)?", str)
```

```
➔ ['7:00', '10:15', '9:00']
```

```
matches = re.findall(r"(?:\d\d?:\d\d) (?:am|pm)?", str)
```

```
➔ ['7:00pm', '10:15pm', '9:00']
```

Tokenization in Python

```
import nltk
import re
sentence = 'That U.S.A. poster-print costs $12.40...'

nltk.word_tokenize(sentence)
➔ ['That', 'U.S.A.', 'poster-print', 'costs', '$', '12.40', '...']

pattern = r"\w+" # tokens are any non-empty sequence of word characters
nltk.regexp_tokenize(sentence, pattern) # or findall of re
re.findall(pattern, sentence)
➔ ['That', 'U', 'S', 'A', 'poster', 'print', 'costs', '12', '40']

pattern = r"\S+" # tokens are white space separated
nltk.regexp_tokenize(sentence, pattern)
➔ ['That', 'U.S.A.', 'poster-print', 'costs', '$12.40...']

pattern = r"""(?x)          # multiline and set verbose for regular expressions
(?:[A-Z]\.)+              # abbreviations
| \w+(?:-\w+)*             # words with optional internal hyphens
| \$?\d+(?:\.\d+)?%?       # currency, percentages, e.g. $12.40, 45%
| \.\.\.                  # ellipsis (three dots)
| [][.,;"'()? :_-]        # other single character tokens
"""

nltk.regexp_tokenize(sentence, pattern)
➔ ['That', 'U.S.A.', 'poster-print', 'costs', '$', '12.40', '...']
```

Tokenization: Language Issues

- French:
 - *L'ensemble* to two words *un ensemble*
- German noun compounds are not segmented:
 - *Lebensversicherungsgesellschaftsangestellter*
 - German tokenizer needs **compound splitter**.
- Chinese and Japanese no spaces between words:
 - 莎拉波娃现在居住在美国东南部的佛罗里达。
 - 莎拉波娃 现在 居住 在 美国 东南部 的 佛罗里达
 - Sharapova now lives in US southeastern Florida

Word Tokenization in Chinese

- Word tokenization is also called **Word Segmentation**
- Chinese words are composed of characters
 - Characters are generally 1 syllable and 1 morpheme.
 - Average word is 2.4 characters long.
- Standard baseline segmentation algorithm: **Maximum Matching**

Given a wordlist of Chinese, and a string.

1. Start a pointer at the beginning of the string
2. Find the longest word in dictionary that matches the string starting at pointer
3. Move the pointer over the word in string
4. Go to 2

Max-match segmentation

- Thecatinthehat the cat in the hat
- Thetabledownthere the table down there
 theta bled own there
 - Doesn't generally work in English!
- But works well in Chinese
 - 莎拉波娃现在居住在美国东南部的佛罗里达。
 - 莎拉波娃 现在 居住 在 美国 东南部 的 佛罗里达
- Modern probabilistic segmentation algorithms even better

Text Normalization

- **Tokens** can also be **normalized**, in which a single normalized form is chosen for words with multiple forms like **USA** and **US**.
 - This standardization may be valuable, despite the spelling information that is lost in the normalization process.
 - For information retrieval, we want a query for US to match a document that has USA.
- **Case folding** is another kind of normalization: Reduce all letters to lower case.
 - For most applications (information retrieval), case folding is helpful.
 - For some NLP applications (MT, information extraction) cases can be helpful.
 - US versus us are important

Lemmatization

- **Lemmatization** is the task of determining that two words have the **same root**, despite their surface differences.
 - *am, are, is* → *be*
 - *car, cars, car's, cars'* → *car*
- **Lemmatization**: have to find correct dictionary **headword form** of the Word.
- The most sophisticated methods for lemmatization involve complete **morphological parsing** of the word.
- **Morphology** is the study of the way words are built up from smaller meaning-bearing units called **morphemes**.
- Two broad classes of morphemes can be distinguished:
 - **Stems** : the central morpheme of the word, supplying the main meaning
 - **Affixes** : adding “additional” meanings of various kinds.

Lemmatization

- **Lemmatization** algorithms can be complex.
- For this reason we sometimes make use of a simpler but cruder method, which mainly consists of chopping off word-final affixes.
- This naive version of morphological analysis is called **stemming**.
- One of the most widely used stemming algorithms is **Porter Stemmer**.
 - The algorithm is based on series of **rewrite rules** run in series, in which the output of each pass is fed as input to the next pass.
 - Some rules are:
 - ATIONAL $\rightarrow$ ATE (e.g., relational $\rightarrow$ relate)
 - ING $\rightarrow$ ε if stem contains vowel (e.g., motoring $\rightarrow$ motor)
 - SSES $\rightarrow$ SS (e.g., grasses $\rightarrow$ grass)

Sentence Segmentation

- **Sentence segmentation** is another important step in text processing.
- The most useful cues for segmenting a text into sentences are **punctuation**, like **periods, question marks, exclamation points**.
- Question marks and exclamation points are relatively unambiguous markers of sentence boundaries.
- Periods, on the other hand, are more ambiguous.
 - Abbreviations like **Inc.** or **Dr.**
 - Numbers like **.02%** or **4.3**
- Build a binary classifier
 - Looks at a “.”
 - Decides EndOfSentence/NotEndOfSentence
 - Classifiers: hand-written rules, regular expressions, or machine-learning

Tokenization

- **Subword Tokenization**

Subword Tokenization

- Another option for text tokenization → **Subword Tokenization**
 - Instead of white-space word segmentation, use single-character segmentation
 - Use the data (**training corpus**) to tell us how to tokenize.
- In **Subword Tokenization**, tokens can be parts of words as well as whole words.
- To deal with this ***unknown word problem***, modern tokenizers (used by Large Language Models) automatically induce sets of tokens that include tokens smaller than words, called **subwords**.
 - Subwords can be arbitrary substrings, or they can be meaning-bearing units like the morphemes -est or -er.
 - In modern tokenization schemes, most tokens are words, but some tokens are frequently occurring morphemes or other subwords like **-er**.
 - Every unseen word like **lower** can thus be represented by some sequence of known subword units, such as **low** and **er**, or even as a sequence of individual letters if necessary.

Subword Tokenization

- Two common algorithms for subword tokenization:
 - **Byte-Pair Encoding (BPE)** tokenization (Sennrich et al., 2016)
 - **Unigram Language Modeling (ULM)** tokenization (Kudo, 2018)
 - sometimes ULM tokenization is confusingly called as "**SentencePiece**" since it is the library it's in SentencePiece library.
- All algorithms have 2 parts:
 - A **token learner** that takes a raw training corpus and induces a vocabulary (a set of tokens).
 - A **token segmenter** that takes a raw test sentence and tokenizes it according to that vocabulary

Byte-Pair Encoding (BPE) - Token Learner

- Let vocabulary be the set of all individual characters
 $= \{A, B, C, D, \dots, a, b, c, d, \dots\}$
- Repeat:
 - Choose the two symbols that are most frequently adjacent in the training corpus (say 'A', 'B')
 - Add a new merged symbol 'AB' to the vocabulary
 - Replace every adjacent 'A' 'B' in the corpus with 'AB'.
- Until k merges have been done.

Byte-Pair Encoding (BPE)

- Token Learner

function BYTE-PAIR ENCODING(strings C , number of merges k) **returns** vocab V

$V \leftarrow$ all unique characters in C # initial set of tokens is characters

for $i = 1$ **to** k **do** # merge tokens k times

$t_L, t_R \leftarrow$ Most frequent pair of adjacent tokens in C

$t_{NEW} \leftarrow t_L + t_R$ # make new token by concatenating

$V \leftarrow V + t_{NEW}$ # update the vocabulary

 Replace each occurrence of t_L, t_R in C with t_{NEW} # and update the corpus

return V

Byte-Pair Encoding (BPE)

- Token Learner: Example

- Most subword algorithms are run inside space-separated tokens.
- So, we commonly first explicitly mark the spaces between words with a special word-seperator-symbol '▯'.
- Next, separate into letters (characters).

- Tiny Corpus:

"set new new renew reset renew"

➔ **set▯new▯new▯renew▯reset▯renew**

Byte-Pair Encoding (BPE)

- Token Learner: Example

- First, we break up the corpus into words, with leading whitespace, together with their counts; no merges will be allowed to go beyond these word boundaries.

corpus

```
2  □ n e w
2  □ r e n e w
1  s e t
1  □ r e s e t
```

vocabulary

```
□ e n r s t w
```

- Merge **n e** to **ne** (most frequent pair which occurs 4 times)

corpus

```
2  □ ne w
2  □ r e ne w
1  s e t
1  □ r e s e t
```

vocabulary

```
□ e n r s t w ne
```

Byte-Pair Encoding (BPE)

- Token Learner: Example

- Merge **ne w** to **new** (most frequent pair which occurs 4 times)

corpus

```
2  □ new
2  □ r e new
1  s e t
1  □ r e s e t
```

vocabulary

```
□ e n r s t w ne new
```

- Merge **□ r** to **□r** (most frequent pair which occurs 3 times)

corpus

```
2  □ new
2  □r e new
1  s e t
1  □r e s e t
```

vocabulary

```
□ e n r s t w ne new □r
```

Byte-Pair Encoding (BPE)

- Token Learner: Example

- Merge **▮r e** to **▮re** (most frequent pair which occurs 3 times)

corpus

```
2  ▮ new
2  ▮re new
1  s e t
1  ▮re s e t
```

vocabulary

```
▮ e n r s t w ne new ▮r ▮re
```

- Next merges: (pairs occur 2 times)

merges

```
▮ new
▮re new
s e
se t
```

vocabulary

```
▮ e n r s t w ne new ▮r ▮re ▮new
▮ e n r s t w ne new ▮r ▮re ▮new ▮renew
▮ e n r s t w ne new ▮r ▮re ▮new ▮renew se
▮ e n r s t w ne new ▮r ▮re ▮new ▮renew se set
```

Byte-Pair Encoding (BPE)

- Token Learner: Example

- Merge **▮r e** to **▮re** (most frequent pair which occurs 3 times)

corpus

```
2  ▮ new
2  ▮re new
1  s e t
1  ▮re s e t
```

vocabulary

```
▮ e n r s t w ne new ▮r ▮re
```

- Next merges: (pairs occur 2 times)

corpus

```
2  ▮new
2  ▮renew
1  set
1  ▮re set
```

vocabulary

```
▮ e n r s t w ne new ▮r ▮re ▮new ▮renew se set
```

Byte-Pair Encoding (BPE)

- Token Segmenter

- On the test data, run each merge learned from the training data:
 - Greedily
 - In the order we learned them

- Learned **vocabulary**

 ▯ e n r s t w ne new ▯r ▯re ▯new ▯renew se set
Merges: 1 2 3 4 5 6 7 8

- So: merge every **n e** to **ne**, then merge **ne w** to **new** etc.
- Result:
 - Test set " newer" → "▯ n e w e r" would be tokenized as a full word "**▯new e r**"
 - Test set " replay" → "▯ r e p l a y" would be two tokens: "**▯re p l a y**"

Properties of Byte-Pair Encoding Tokens

- Usually include frequent words as tokens
- And frequent subwords
 - Which are often morphemes like *-est* or *-er*
- A **morpheme** is the smallest meaning-bearing unit of a language
 - *unlikeliest* has 3 morphemes *un-*, *likely*, and *-est*

Byte-Pair Encoding (BPE) in Practice

UTF-8 (Unicode Transformation Format 8)

- We assumed that BPE algorithm takes a list of characters as the initial vocabulary.
- In real applications, BPE runs on the individual bytes of UTF-8-encoded text.
 - ie. BPE starts with an initial vocabulary with 256 tokens (256 possible values of a byte)
 - No unknown tokens.
 - In real applications, BPE runs on huge Unicode corpora to learn a vocabulary with size 50000 to 20000.
- UTF-8 is a variable length encoding. In UTF-8, characters can be 1,2,3 or 4 bytes.
 - Python uses Unicode to represent strings. Each string is a sequence of Unicode code points.
 - Unicode code points can map UTF-8 encoding as follows.

Code Points		UTF-8 Encoding			
From - To	Bit Value	Byte 1	Byte 2	Byte 3	Byte 4
U+0000-U+007F	0xxxxxxx	xxxxxxx			
U+0080-U+07FF	00000yyy yyxxxxxx	110yyyyy	10xxxxxx		
U+0800-U+FFFF	zzzzyyyy yyxxxxxx	1110zzzz	10yyyyyy	10xxxxxx	
U+10000-U+10FFFF	000uuuuu zzzzyyyy yyxxxxxx	11110uuu	10uuzzzz	10yyyyyy	10xxxxxx

Byte-Pair Encoding (BPE) in Practice

Visualizing GPT4o tokens

- The BPE tokenizer used in OpenAI GPT4o has 200K tokens, which is a comparatively large number.
 - We can use **Tiktokenizer** visualizer (<https://tiktokenizer.vercel.app/>) to see the number of tokens in a given sentence.
 - Visualize the sentence: "We can use a visualizer to see the tokens in a given sentence."

We·can·use·a·visualizer·to·see·the·tokens·in·a·given·sentence.

Token count: 15

Tokens: 2167, 665, 1199, 261, 11256, 4492, 316, 1921, 290, 20290, 306, 261, 4335, 21872, 13

- The visualization shows colors to separate out words, but the true output of the tokenizer is simply a sequence of unique token ids.

Byte-Pair Encoding (BPE) in Practice

Visualizing GPT4o tokens

Emojis 🍷✅🤖🌟🥹🔥⭐🤨

One word and 8 emojis are represented by 13 tokens.

Tokens: 2585, 105199, 116548, 65941, 138021, 97375, 112808, 117, 96606, 62160, 4103, 104, 102

The·slowest·car·and·slowest·train

Tokens: 976, 9800, 376, 1669, 326, 9800, 376, 8513

a·slowest·car·a·slowest·train·a·fastest·car

Tokens: 64, 9800, 376, 1669, 261, 9800, 376, 8513, 261, 32840, 1669

Byte-Pair Encoding (BPE) in Practice

Tokenizing across languages

- Even though BPE tokenizers use multilingual LLM training data, the data is still vastly dominated by English
 - Most BPE tokens used for English, leaving less for other languages.
 - Words in other languages are often split up.

He·climbed·the·mountain·slowly·He·was·so·tired·

- Tokens: 2066, 85578, 290, 21689, 22936, 13, 1679, 673, 813, 25920, 13
- 11 tokens. All 9 English words are separate tokens.

O·dağı·yavaşca·tırmandı·O·çok·yorulmuştu·

- Tokens: 46, 194313, 342, 168200, 3743, 260, 6637, 3842, 612, 13, 532, 24931, 159564, 361, 49052, 10191, 13
- 17 tokens. 7 Turkish words are represented by 15 tokens.

Summary

Word Tokenization

- The **regular expression** language is a powerful tool for pattern-matching.
- **Word tokenization** and **normalization** are generally done by cascades of simple *regular expression substitutions* or finite automata.

Subword Tokenization

- **Words and morphemes** are useful units of representation, but difficult to define formally.
- **Unicode** is a system for representing characters in the many scripts used to write the languages of the world.
- Each character is represented internally with a unique id called a **code point**, and can be encoded in a file via encoding methods like **UTF-8**, which is a variable-length encoding.
- **Byte-Pair Encoding (BPE)** is the standard way to induce tokens in a data driven way. It is the first step in most large language models.
- **BPE tokens** are often roughly word or morpheme-sized, although they can be as small as single characters.