

Why Learn Theory?

Q. What can a computer do with limited resources?

e.g., Intel Core 2 Duo running Linux kernel 2.6

- Pioneering work in the 1930s.

- 
- A black and white portrait of a man with a white beard and mustache, wearing a dark suit, a white shirt, and a dark tie. He is also wearing a light-colored hat with a dark band. The portrait is set against a plain, light background.

David Hilbert



Kurt Gödel



Alan Turing



Alonzo Church



John von Neumann

4

Regular Expressions


Describing a Pattern

PROSITE. Huge database of protein families and domains.

Q. How to describe a protein motif?

Ex. [signature of the C_2H_2 -type zinc finger domain]

- **C**
- Between 2 and 4 amino acids.
- **C**
- 3 more amino acids.
- One of the following amino acids: **LIVMFYWCX**.
- 8 more amino acids.
- **H**
- Between 3 and 5 more amino acids.
- **H**



CAASCGGPY**ACGGWAGYHAGWH**

6

Pattern Matching Applications

Test if a string matches some pattern.

- Process natural language.
- Scan for virus signatures.
- Access information in digital libraries.
- Search-and-replace in a word processors.
- Filter text (spam, NetNanny, ads, Carnivore, malware).
- Validate data-entry fields (dates, email, URL, credit card).
- Search for markers in human genome using PROSITE patterns.

Parse text files.

- Compile a Java program.
- Crawl and index the Web.
- Read in data stored in TOY input file format.
- Automatically create Java documentation from Javadoc comments.

Regular Expressions: Basic Operations

Regular expression. Notation to specify a set of strings.

operation	regular expression	matches	does not match
concatenation	aabaab	aabaab	<i>every other string</i>
wildcard	.u.u.u.	cumulus jugulum	succubus tumultuous
union	aa baab	aa baab	<i>every other string</i>
closure	ab*a	aa abbba	ab ababa
parentheses	a(a b)aab	aaaab abaab	<i>every other string</i>
	(ab)*a	a ababababa	aa abbba

Regular Expressions: Examples

Regular expression. Notation is surprisingly expressive.

regular expression	matches	does not match
<code>.*spb.*</code> <i>contains the trigraph spb</i>	raspberry crispbread	subspace subspecies
<code>a* (a*ba*ba*ba*)*</code> <i>multiple of three b's</i>	bbb aaa bbbaababbaa	b bb baabbbbaa
<code>.*0....</code> <i>fifth to last digit is 0</i>	1000234 98701234	11111111 403982772
<code>gcg(cgg agg)*ctg</code> <i>fragile X syndrome indicator</i>	gcgctg gcgcggctg gcgcggaggctg	gcgcgg cgccggcgctg gcgcaggctg

9

Generalized Regular Expressions

Regular expressions are a standard programmer's tool.

- Built in to Java, Perl, Unix, Python,
- Additional operations typically added for convenience.
 - Ex 1: `[a-e]+` is shorthand for `(a|b|c|d|e)(a|b|c|d|e)*`.
 - Ex 2: `\s` is shorthand for "any whitespace character" (space, tab, ...).

operation	regular expression	matches	does not match
one or more	<code>a(bc)+de</code>	abcde abcbcede	ade bcde
character class	<code>[A-Za-z][a-z]*</code>	lowercase Capitalized	camelCase 4illegal
exactly k	<code>[0-9]{5}-[0-9]{4}</code>	08540-1321 19072-5541	11111111 166-54-1111
negation	<code>[^aeiou]{6}</code>	rhythm	decade

10

Regular Expressions in Java

Validity checking. Is input in the set described by the re?

```
public class Validate
{
    public static void main(String[] args) {
        String re = args[0];
        String input = args[1];
        StdOut.println(input.matches(re));
    }
}
```

powerful string library method

```
% java Validate "C.{2,4}C...[LIVMFYWC].{8}H.{3,5}H" CAASC GGPyACGGAAGYHAGAH
true
% java Validate "[$_A-Za-z][$_A-Za-z0-9]*" ident123
true
% java Validate "[a-z]+@[a-z]+\.(edu|com)" wayne@cs.princeton.edu
true
```

Annotations:
 - `C2H2` type zinc finger domain
 - legal Java identifier
 - valid email address (simplified)
 - need quotes to "escape" the shell

11

String Searching Methods

```
public class String (Java's String library)
```

```
boolean matches(String re) does this string match the given regular expression

String replaceAll(String re, String str) replace all occurrences of regular expression with the replacement string

int indexOf(String r, int from) return the index of the first occurrence of the string r after the index from

String[] split(String re) split the string around matches of the given regular expression
```

```
String s = StdIn.readAll();
s = s.replaceAll("\\s+", " ");
```

replace each sequence of at least one whitespace character with a single space

RE that matches any sequence of whitespace characters (at least 1).
 Extra \ distinguishes from the string \s+

12

String Searching Methods

`public class String` *(Java's String library)*

<code>boolean matches(String re)</code>	<i>does this string match the given regular expression</i>
<code>String replaceAll(String re, String str)</code>	<i>replace all occurrences of regular expression with the replacement string</i>
<code>int indexOf(String r, int from)</code>	<i>return the index of the first occurrence of the string r after the index from</i>
<code>String[] split(String re)</code>	<i>split the string around matches of the given regular expression</i>

```
String s = StdIn.readAll();
String[] words = s.split("\\s+");
```

create an array of the words in StdIn

13

DFAs


Solving the Pattern Match Problem

Regular expressions are a concise way to describe patterns.

- How would you implement the method `matches()` ?
- Hardware: build a deterministic finite state automaton (DFA).
- Software: simulate a DFA.

DFA: simple machine that solves a pattern match problem.

- Different machine for each pattern.
- Accepts or rejects string specified on input tape.
- Focus on `true` or `false` questions for simplicity.




15

Deterministic Finite State Automaton (DFA)

Simple machine with N states.

- Begin in start state.
- Read first input symbol.
- Move to new state, depending on current state and input symbol.
- Repeat until last input symbol read.
- Accept input string if last state is labeled Y.



16


DFA and RE Duality

RE. Concise way to **describe** a set of strings.

DFA. Machine to **recognize** whether a given string is in a given set.

Duality.

- For any DFA, there exists a RE that describes the same set of strings.
- For any RE, there exists a DFA that recognizes the same set.



Practical consequence of duality proof: to match RE

- build DFA
- simulate DFA on input string.

17

Implementing a Pattern Matcher

Problem. Given a RE, create program that tests whether given input is in set of strings described.

Step 1. Build the DFA.

- A compiler!
- See COS 226 or COS 320.

Step 2. Simulate it with given input.

```
State state = start;
while (!StdIn.isEmpty())
{
    char c = StdIn.readChar();
    state = state.next(c);
}
StdOut.println(state.accept());
```

18

Application: Harvester

Harvest information from input stream.

- Harvest patterns from DNA.

```
% java Harvester "gcgcggcggcggcggcggcgtg" chromosomeX.txt
gcgcggcggcggcggcggcggcggcgtg
gcgcgtg
gcgcgtg
gcgcggcggcggcggcggcggcggcgtg
```

- Harvest email addresses from web for spam campaign.

```
% java Harvester "[a-z]+@[a-z]+\.(edu|com)" http://www.princeton.edu/~cos126
rs@cs.princeton.edu
maia@cs.princeton.edu
doug@cs.princeton.edu
wayne@cs.princeton.edu
```

19

Application: Harvester

Harvest information from input stream.

- Use `Pattern` data type to compile regular expression to NFA.
- Use `Matcher` data type to simulate NFA.

equivalent, but more efficient representation of a DFA

```
import java.util.regex.Pattern;
import java.util.regex.Matcher;

public class Harvester
{
    public static void main(String[] args)
    {
        String re      = args[0];
        In in          = new In(args[1]);
        String input    = in.readAll();
        Pattern pattern = Pattern.compile(re);
        Matcher matcher = pattern.matcher(input);

        while (matcher.find())
            StdOut.println(matcher.group());
    }
}
```

Annotations in the original image:

- create NFA from RE (pointing to `Pattern.compile(re)`)
- create NFA simulator (pointing to `Matcher matcher = pattern.matcher(input);`)
- look for next match (pointing to `matcher.find()`)
- the match most recently found (pointing to `matcher.group()`)

20

Application: Parsing a Data File

Ex: parsing an NCBI genome data file.

header info


```

LOCUS AC146846 128142 bp DNA linear HTG 13-NOV-2003
DEFINITION Ornithorhynchus anatinus clone CLM1-393H9,
ACCESSION AC146846
VERSION AC146846.2 GI:38304214
KEYWORDS HTG; HTGS_PHASE2; HTGS_DRAFT.
SOURCE Ornithorhynchus anatinus (platypus)
ORIGIN
  1 tgtatttcat ttgacogtgc tgttttttcc oggtttttca gtaacgtgtt agggagccac
  61 gtgattctgt ttgttttatg ctgcgaata gctgctgat gaattctgc atagacagct
  121 gcgcagggga gaaatgacca gtttgtgatg acaaatgta ggaaagctgt ttcttcataa
  ...
128101 ggaaatgca ccccaagct aatgtacagc ttcttagat tg
//
    
```

line numbers

spaces

comments



Goal. Extract the data as a single **actg** string.

21

Application: Parsing a Data File

```


import java.util.regex.Pattern;
import java.util.regex.Matcher;

public class ParseNCBI
{
    public static void main(String[] args)
    {
        String re = "[ ]*[0-9]+([actg ]*).*";
        Pattern pattern = Pattern.compile(re);
        In in = new In(args[0]);
        String data = "";
        while (!in.isEmpty())
        {
            String line = in.readLine();
            Matcher matcher = pattern.matcher(line);
            if (matcher.find())
            {
                data += matcher.group(1).replaceAll(" ", "");
            }
        }
        System.out.println(data);
    }
}
    
```

extract the part of match in ()

```

LOCUS AC146846 128142 bp DNA linear HTG 13-NOV-2003
DEFINITION Ornithorhynchus anatinus clone CLM1-393H9,
ACCESSION AC146846
VERSION AC146846.2 GI:38304214
KEYWORDS HTG; HTGS_PHASE2; HTGS_DRAFT.
SOURCE Ornithorhynchus anatinus (platypus)
ORIGIN
  1 tgtatttcat ttgacogtgc tgttttttcc oggtttttca gtaacgtgtt agggagccac
  61 gtgattctgt ttgttttatg ctgcgaata gctgctgat gaattctgc atagacagct // a comment
  121 gcgcagggga gaaatgacca gtttgtgatg acaaatgta ggaaagctgt ttcttcataa
  ...
128101 ggaaatgca ccccaagct aatgtacagc ttcttagat tg
//
    
```



22

Regular Expressions



<http://xkcd.com/208/>

23

Summary

Programmer.

- Regular expressions are a powerful pattern matching tool.
- Implement regular expressions with finite state machines.

Theoretician.

- RE is a compact description of a set of strings.
- DFA is an abstract machine that solves RE pattern match problem.

You. Practical application of core CS principles.

24

Fundamental Questions

Q. Are there patterns that **cannot** be described by any RE?

A. Yes.

- Bit strings with equal number of 0s and 1s.
- Strings that represent legal REs.
- Decimal strings that represent prime numbers.
- DNA strings that are Watson-Crick complemented palindromes.

25

Fundamental Questions

Q. Are there languages that **cannot** be recognized by any DFA?

A. Yes.

- Bit strings with equal number of 0s and 1s.
- Strings that represent legal REs.
- Decimal strings that represent prime numbers.
- DNA strings that are Watson-Crick complemented palindromes.

26

Fundamental Questions

Q. Are there languages that **cannot** be recognized by any DFA?

A. Yes: Bit strings with equal number of 0s and 1s.

Proof sketch.

- Suppose that you have such a DFA, with N states.
- Give it N+1 0s followed by N+1 1s.
- Some state is revisited.
- Delete substring between visits.
- DFA recognizes that string, too.
- It does not have equal number of 0s and 1s.
- Contradiction.
- No such DFA exists.

0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

0	1	3	5	6	8	7	3	5	.	.	.
---	---	---	---	---	---	---	---	---	---	---	---

0	0	0	0	1	1	1	1	1	1	1	1	1	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---

0	1	3	5	.	.	.
---	---	---	---	---	---	---

27

Fundamental Questions

Q. Are there languages that **cannot** be recognized by any DFA?

A. Yes.

- Bit strings with equal number of 0s and 1s.
- Strings that represent legal REs.
- Decimal strings that represent prime numbers.
- DNA strings that are Watson-Crick complemented palindromes.

Fundamental problem: DFA lacks **memory**.

28

Fundamental Questions

Q. Are there machines that are **more powerful** than a DFA?

A. Yes.

A **1-stack DFA** can recognize

- Bit strings with equal number of 0s and 1s.
- Legal REs.
- Watson-Crick complemented palindromes.

29

Fundamental Questions

Q. Are there machines that are **more powerful** than a 1-stack DFA?

A. Yes.

A **2-stack DFA** can recognize


- Prime numbers.
- Legal Java Programs.

30

Fundamental Questions

Q. Are there machines that are more powerful than a 2-stack DFA?

A. No! Not even a supercomputer!



2-stack DFAs are equivalent to **Turing machines** [stay tuned].

31