


CS103
WINTER 2025




Lecture 14:
Finite Automata

Part 1 of 3

Problem Set Three Graded



Midterm 1



- **80th** Percentile:
48 / 50 (**96%**)
- **60th** Percentile:
45 / 50 (**90%**)
- **40th** Percentile:
41 / 50 (**82%**)
- **20th** Percentile:
34 / 50 (**68%**)

Thoughts and Observations

- This is only 12.5% of your grade.
- We want everyone to be wildly successful!
 - 1-on-1s (contact Emily)
 - Review feedback
 - Assess (small scattered point losses? one large loss?)
- Assuming comfort going forward:
 - Contrapositive
 - Negations (implication, quant.)
 - Assume/Prove table
 - Proofwriting Checklist

CoDa Update

- The coffee shop is now open! (7 AM – 2 PM)
- Starting March 3 it will be open later. (7 AM – 5 PM)
- Feel free to make use of the beautiful study spaces throughout the building!

DESTINATIONS	
TOKYO 6.2 <i>Espresso, milk, cherry blossom water, vanilla, brown sugar</i>	MANILA 6.2 <i>Espresso, milk, ube, mango, coconut crème, vanilla, brown sugar</i>
VALENCIA 6.2 <i>Espresso, milk, orange blossom water, coconut, vanilla, brown sugar</i>	SANTIAGO 6.2 <i>Espresso, milk, cocoa, cayenne, ginger, cinnamon, brown sugar</i>
LEXINGTON 6.2 <i>Espresso, milk, Jim Beam bourbon reduction, vanilla, brown sugar</i>	BALI 6.2 <i>Espresso, milk, coconut crème, honey, cayenne</i>

TRADITIONAL	NOT COFFEE
ESPRESSO 3.75	LOOSE LEAF TEA 4 • Black: Earl Grey, China Breakfast • Green: Jasmine, Quince • Herbal: Chamomille, Peppermint, African Sunset Rooibos
MACCHIATO 4	LONDON FOG 5.3
CORTADO 4.5	SNOWBIRD 5.8
FLAT WHITE 4.8	CHAI LATTE 5.3
CAPPUCCINO 4.8	MATCHA LATTE 5.5 • Add Strawberry +0.60
LATTE 5.3	LEMONADE 4.5 • Add Strawberry
MOCHA 5.6	ORANGE JUICE 5
DRIP 4	
AMERICANO 4.3	
COLD BREW 5	

FLAVOR \$0.5 | ALMOND, OAT, SOY \$0.5 | EXTRA ESPRESSO \$1

BRUNCHIE BITES	
YOGURT & HOUSE-MADE GRANOLA • GF* 6.5 <i>Straus Greek yogurt, housemade lemon zest & honey butter granola. Topped with a drizzle of honey & fresh strawberries!</i>	
CHOCO GANACHE & ALMOND CHIA PUDDING • V, GF* 7 <i>Chia seeds, almond milk, brown sugar, salt, semi-sweet chocolate, coconut cream, almonds & cacao nibs. Topped with fresh strawberries!</i>	
PB & JAM OVERNIGHT OATS • V, GF* 7 <i>Rolled oats, chia seeds, oat milk, and maple syrup, with peanut butter & raspberry jam. Topped with cacao nibs!</i>	
BREAKFAST BURRITO 10 <i>Bombay spiced potatoes, tender scrambled eggs, & white cheddar folded into a burrito. House-made chipotle sauce on the side!</i>	
EGG FRITTATA • GF* 6.5 <i>Eggs & cream whipped with seasonal vegetables & cheese, baked as a soufflé!</i>	

V – VEGAN, GF* – GLUTEN FRIENDLY


Outline for Today

- ***Computability Theory***
 - What problems can we solve with a computer?
- ***Formal Language Theory***
 - Stringy thingies.
- ***Finite Automata***
 - A very simple model of a computing device.

Computability Theory

What problems can we solve with a computer?

What kind of
computer?




Two Challenges

- Computers are dramatically better now than they've ever been, and that trend continues.
- Writing proofs on formal definitions is hard, and computers are *way* more complicated than sets, graphs, or functions.
- ***Key Question:*** How can we prove what computers can and can't do...
 - ... so that our results are still true in 20 years?
 - ... without multi-hundred page proofs?


Enter Automata

- An **automaton** (plural: **automata**) is a mathematical model of a computing device.
- It's an **abstraction** of a real computer, the way that graphs are abstractions of social networks, transportation grids, etc.
- The automata we'll explore are
 - powerful enough to capture huge classes of computing devices, yet
 - simple enough that we can reason about them in a small space.
- They're also fascinating and useful in their own rights. More on that later.

Toward a Model of Computation...



Why does this computer "feel" less powerful...



...than this one?

Calculators vs. Desktops


- A calculator has a ***small amount of memory***. A desktop computer has a ***large amount of memory***.
- A calculator performs a ***fixed set of functions***. A desktop is ***reprogrammable*** and can run many different programs.
- These two distinctions account for much of the difference between “calculator-like” computers and “desktop-esque” computers.
- In CS103, we’ll first explore “small-memory” computers in detail, then discuss “large-memory” computers in depth.

Let's Focus on Computing with Finite
Memory

Our Goal: A Unifying Abstraction

7	8	9	÷
4	5	6	×
1	2	3	-
0	.	=	+

Data stored electronically.
Algorithm is in silicon.
Memory limited by display.




Data stored in wood.
Algorithm is in brain.
Memory limited by beads.

How do we model “memory” and
“an algorithm” when they can take
on so many forms?

What's in Common?


- These machines **receive input** from an external source.
- That input is provided **sequentially**, one discrete unit at a time.
- Each input causes the device to **change configuration**. This change, big or small, is where the computation happens.
- Once all input is provided, we can **read off an answer** based on the configuration of the device.

179			
7	8	9	÷
4	5	6	×
1	2	3	-
0	.	=	+




Modeling Finite Computation

- We will model a finite-memory computer as a collection of **states** linked by **transitions**.
- Each state corresponds to one possible configuration of the device's memory. This is super abstract!
- Each transition indicates how memory changes in response to inputs.
- Some state is designated as the **start state**. The computation begins in that state.




Modeling Finite Computation

- This device processes **strings** made of **characters**.
 - Each character represents some external input to the device.
 - The string represents the full sequence of inputs to the device.
- To run this device, we begin in our start state and scan the input from left to right.
- Each time the machine sees a character, it **changes state** by following the transition labeled with that character.




Modeling Finite Computation

- Once we've finished entering all the characters of our input, we need to obtain the result of the computation.
- In general, computers can produce all sorts of things as the result of a computation: a number, a piece of text, etc.
- As a simplifying assumption, we'll assume that we just need to get a single bit of output. That is, our machines will just say YES or NO.
- (This can be generalized – come talk to us after class if you're curious how!)



Modeling Finite Computation

- Some of the states in our computational device will be marked as **accepting states**. These are denoted with a double ring.
- If the device ends in an accepting state after seeing all the input, **accepts** the input (says YES).
- If the device does not end in an accepting state after seeing all the input, it **rejects** the input (says NO).




Modeling Finite Computation

- Try it yourself!
Which of these strings does this device accept?

aab

aabb

abbababba




Answer at

<https://cs103.stanford.edu/pollev>

Finite Automata

- This type of computational device is called a **finite automaton** (plural: **finite automata**).
- Finite automata model computers where (1) memory is finite and (2) the computation produces as YES/NO answer.
- In other words, finite automata model predicates, and do so with a fixed, finite amount of memory.



Formalizing Things

Strings

- An **alphabet** is a finite, nonempty set of symbols called **characters**.
 - Typically, we use the symbol Σ to refer to an alphabet.
- A **string over an alphabet Σ** is a finite sequence of characters drawn from Σ .
- Example: Let $\Sigma = \{a, b\}$. Here are some strings over Σ :
a aabaaabbabaaabaaaabbb abbababba
- But wait! There are no quotes here!
- The **empty string** has no characters and is denoted ϵ .

Languages


- A **language over Σ** is a set L consisting of strings over Σ .
- Example: The language of palindromes over $\Sigma = \{a, b, c\}$ is the set
 - $\{\varepsilon, a, b, c, aa, bb, cc, aaa, aba, aca, bab, \dots\}$
- The set of all strings composed from letters in Σ is denoted Σ^* .
 - Formally: $\Sigma^* = \{ w \mid w \text{ is a string over } \Sigma \}$.
- Formally, we say that L is a language over Σ when $L \subseteq \Sigma^*$.

Mathematical Lookalikes

- We now have \in , ε , Σ , and Σ^* . Yikes!
- The symbol \in is the ***element-of*** relation.
- The symbol ε is the ***empty string***.
- The symbol Σ denotes an ***alphabet***.
- The expression Σ^* means “all strings that can be made from characters in Σ .”
- That lets us write things like
 - We have $\varepsilon \in \Sigma^*$, but $\varepsilon \notin \Sigma$.
- Ever get confused? ***Just ask!***


The Cast of Characters

- **Languages** are sets of strings.
- **Strings** are finite sequences of characters.
- **Characters** are individual symbols.
- **Alphabets** are sets of characters.



Finite Automata and Languages


- Let A be an automaton that processes strings drawn from an alphabet Σ .
- The **language of A** , denoted $\mathcal{L}(A)$, is the set of strings over Σ that A accepts:



$$\mathcal{L}(A) = \{ w \in \Sigma^* \mid A \text{ accepts } w \}$$




Finite Automata and Languages

- Let D be the automaton shown to the right. It processes strings over $\{\mathbf{a}, \mathbf{b}\}$.
- Notice that D accepts all strings of \mathbf{a} 's and \mathbf{b} 's that end in \mathbf{a} and rejects everything else.
- So $\mathcal{L}(D) = \{ w \in \{\mathbf{a}, \mathbf{b}\}^* \mid w \text{ ends in } \mathbf{a} \}$.



$$\mathcal{L}(A) = \{ w \in \Sigma^* \mid A \text{ accepts } w \}$$

This means “take this transition if you see an **a** or a **b**.”




What are the languages of these automata? Answer at <https://cs103.stanford.edu/pollev>

$$\mathcal{L}(A) = \{ w \in \Sigma^* \mid A \text{ accepts } w \}$$



The Story So Far

- A ***finite automaton*** is a collection of ***states*** joined by ***transitions***.
- Some state is designated as the ***start state***.
- Some number of states are designated as ***accepting states***.
- The automaton processes a string by beginning in the start state and following the indicated transitions.
- If the automaton ends in an accepting state, it ***accepts*** the input.
- Otherwise, the automaton ***rejects*** the input.
- The ***language*** of an automaton is the set of strings it accepts.


A Small Problem



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



Another Small Problem



The Need for Formalism

- In order to reason about the limits of what finite automata can and cannot do, we need to formally specify their behavior in *all* cases.
- All of the following need to be defined or disallowed:
 - What happens if there is no transition out of a state on some input?
 - What happens if there are *multiple* transitions out of a state on some input?


DFAs

- A ***DFA*** is a
 - ***D***eterministic
 - ***F***inite
 - ***A***utomaton
- DFAs are the simplest type of automaton that we will see in this course.


DFA's

- A DFA is defined relative to some alphabet Σ .
- For each state in the DFA, there must be *exactly one* transition defined for each symbol in Σ .
 - This is the “deterministic” part of DFA.
- There is a unique start state.
- There are zero or more accepting states.


Is this a DFA over $\{0, 1\}$?




Is this a DFA over $\{0, 1\}$?




Is this a DFA over $\{0, 1\}$?



Is this a DFA over $\{0, 1\}$?



Is this a DFA over $\{0, 1\}$?




Designing DFAs

- At each point in its execution, the DFA can only remember what state it is in.
- ***DFA Design Tip:*** Build each state to correspond to some piece of information you need to remember.
 - Each state acts as a “memento” of what you're supposed to do next.
 - Only finitely many different states means only finitely many different things the machine can remember.

Recognizing Languages with DFAs


$L = \{ w \in \{a, b\}^* \mid \text{the number of } b\text{'s in } w \text{ is congruent to two modulo three} \}$



Each state remembers the remainder of the number of **b**s seen so far modulo three.

Recognizing Languages with DFAs

$L = \{ w \in \{a, b\}^* \mid w \text{ contains } aa \text{ as a substring} \}$



More Elaborate DFAs

$L = \{ w \in \{a, *, /\}^* \mid w \text{ represents a C-style comment} \}$

Let's have the **a** symbol be a placeholder for "some character that isn't a star or slash."

Try designing a DFA for comments! Here's some test cases to help you check your work:

Accepted:


```
/*a*/  
/**/  
/***/  
/*aaa*aaa*/  
/*a/a*/
```

Rejected:

```
/**  
/**/a/*aa*/  
aaa/**/aa  
/*/  
/**a/  
//aaaa
```

More Elaborate DFAs

$L = \{ w \in \{a, *, /\}^* \mid w \text{ represents a C-style comment} \}$



Next Time

- ***Regular Languages***
 - An important class of languages.
- ***Nondeterministic Computation***
 - Why must computation be linear?
- ***NFAs***
 - Automata with Magic Superpowers.