

Finite Automata

Part One

Computability Theory

What problems can we solve with a computer?

What kind of
computer?



Two Challenges

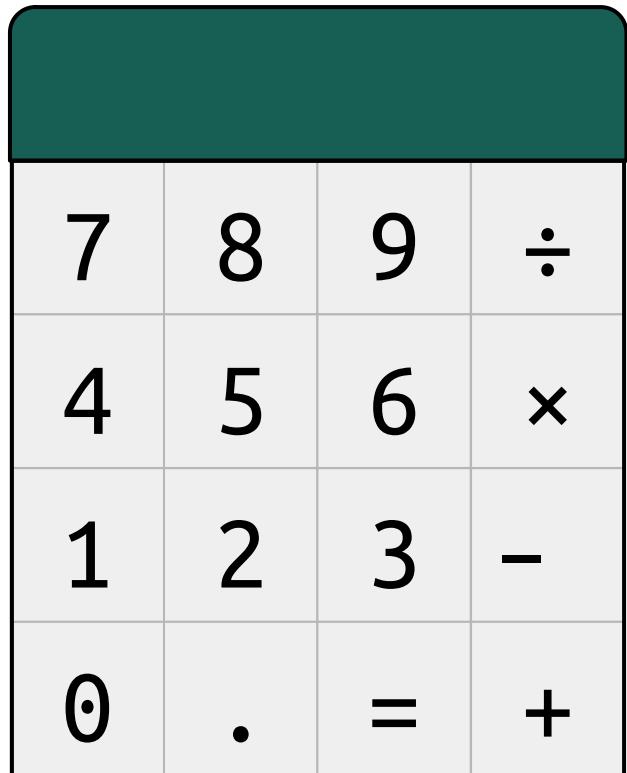
- Computers are dramatically better now than they've ever been, and that trend continues.
- Writing proofs on formal definitions is hard, and computers are *way* more complicated than sets, graphs, or functions.
- ***Key Question:*** How can we prove what computers can and can't do...
 - ... so that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

Enter Automata

- An **automaton** 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.

What do these automata look like?

A Tale of Two Computers



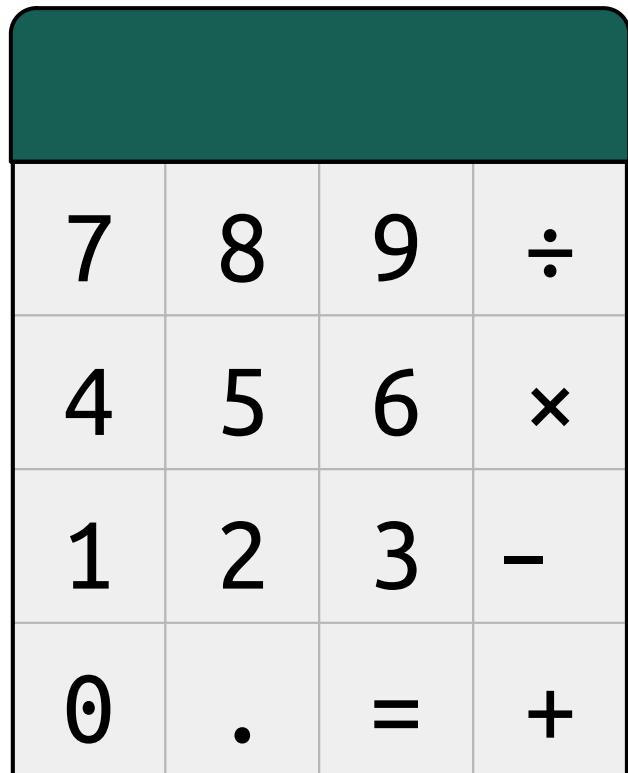
Why does this
computer...

...“feel” less
powerful than
this one?

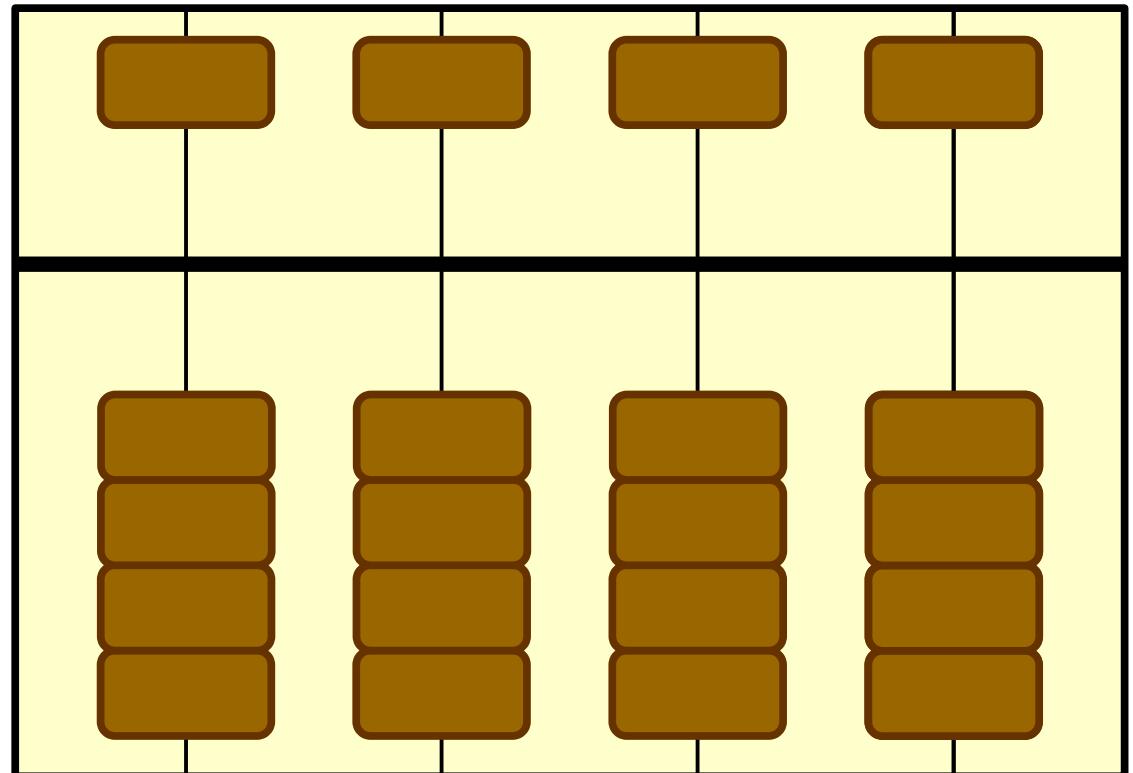
Calculators vs. Desktops

- 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.

Computing with Finite Memory



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



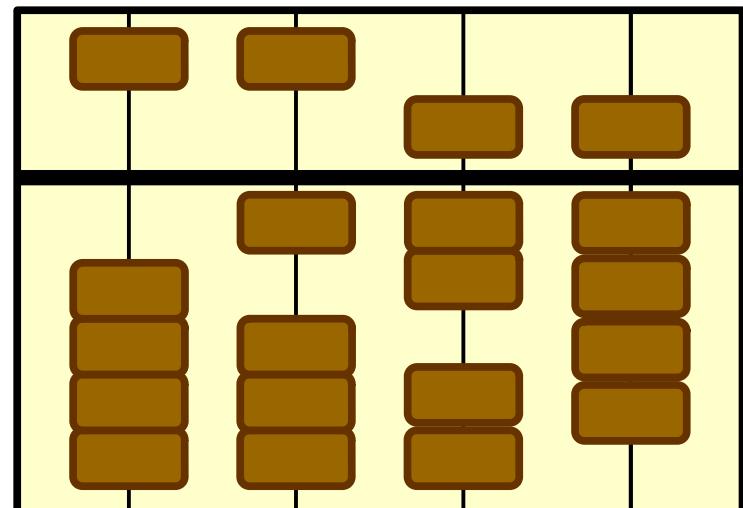
Data stored in wood.
Algorithm is in brain.
Memory limited by beads.

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

What's in Common?

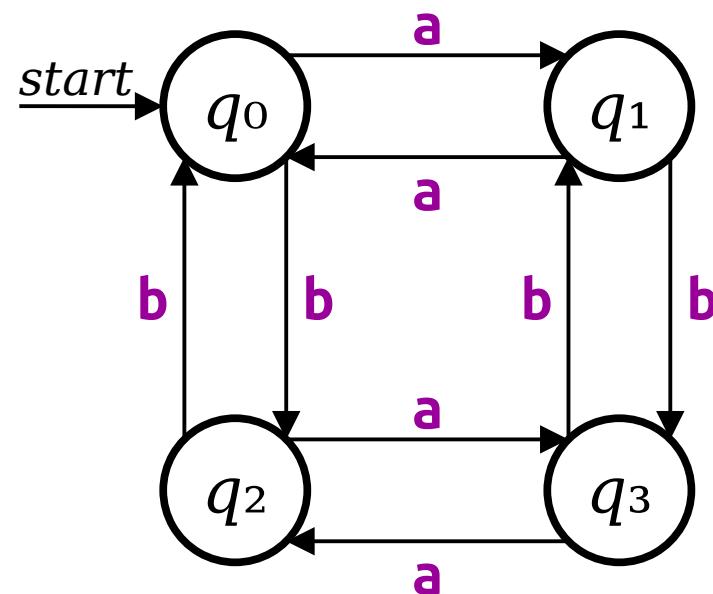
- These machines ***receive input*** from an external source.
- That input is provided ***sequentially***, one discrete unit at a time.
- Each input causes the device to ***change configuration***. This change, big or small, is where the computation happens.
- Once all input is provided, we can ***read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

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



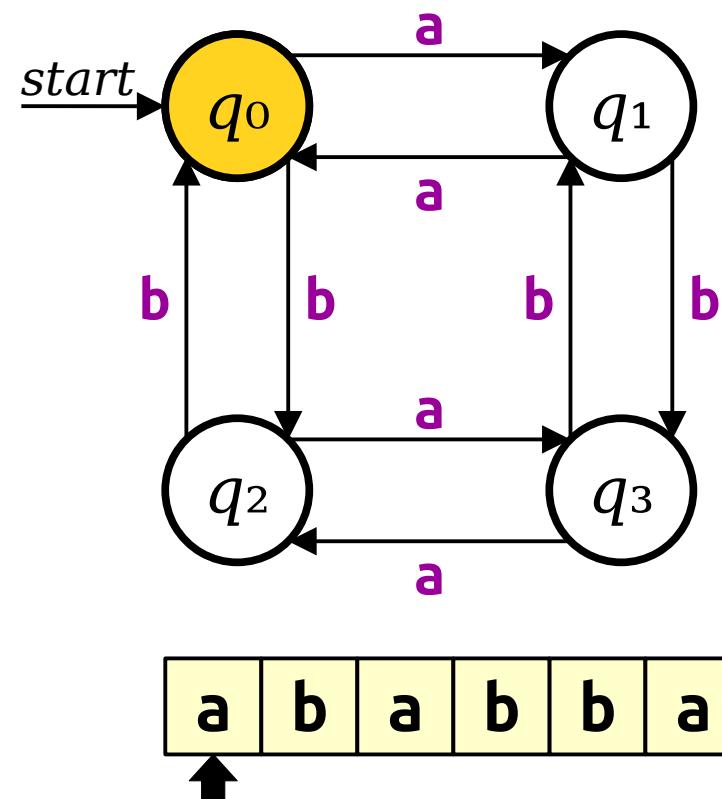
Modeling Finite Computation

- 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.
- Each transition indicates how memory changes in response to inputs.
- Some state is designated as the **start state**. The computation begins in that state.



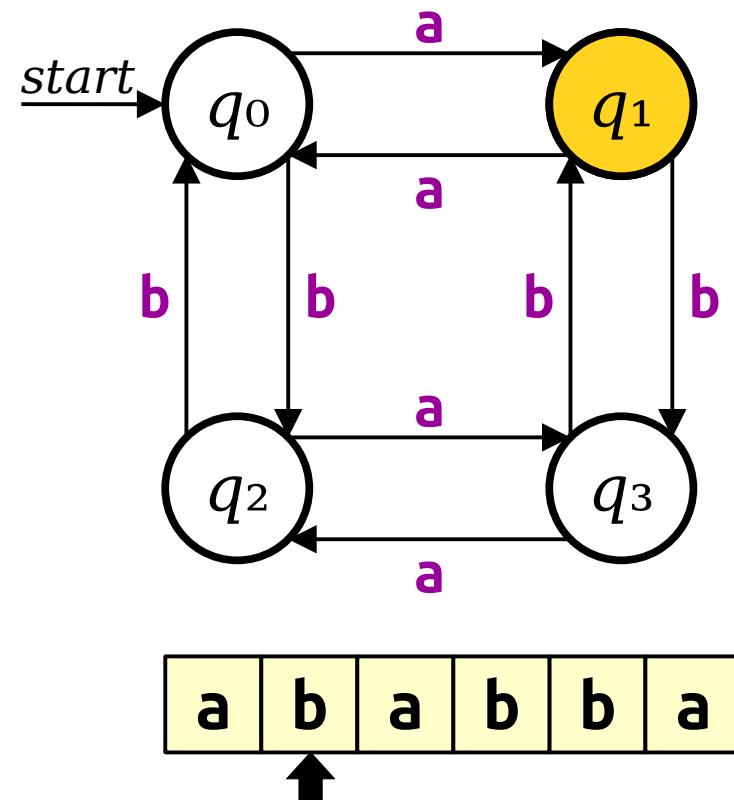
Modeling Finite Computation

- 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.



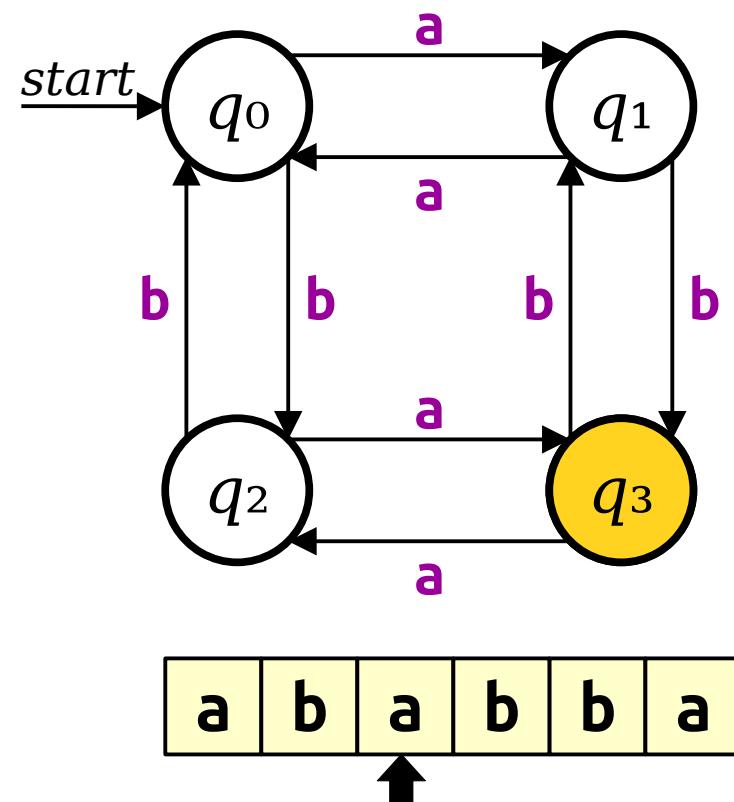
Modeling Finite Computation

- 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.



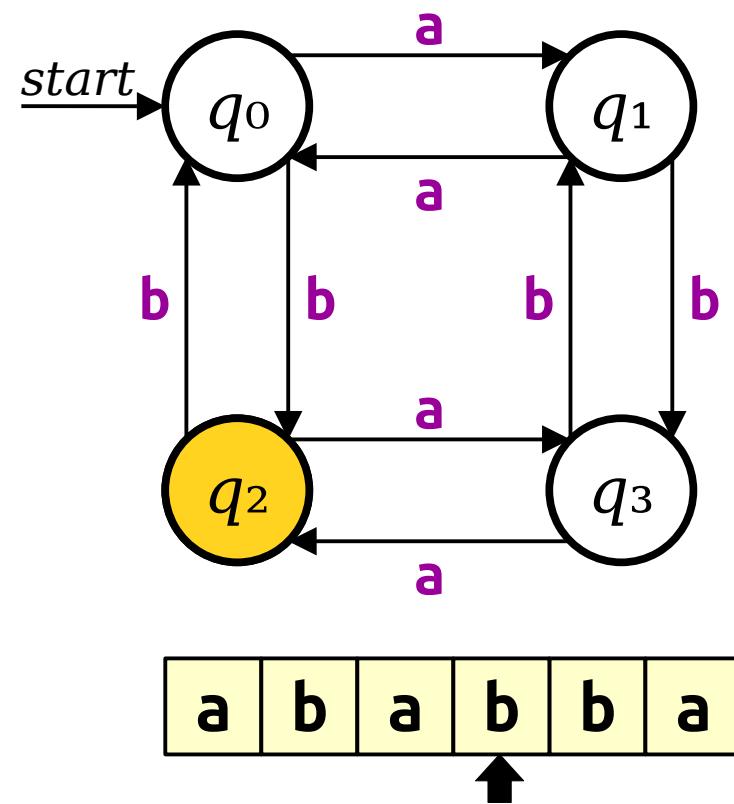
Modeling Finite Computation

- 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.



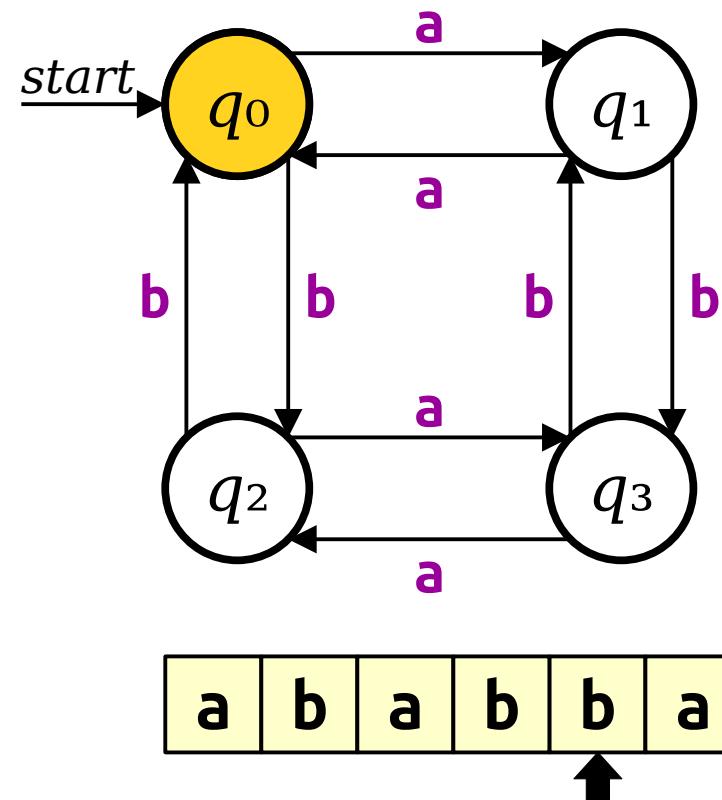
Modeling Finite Computation

- 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.



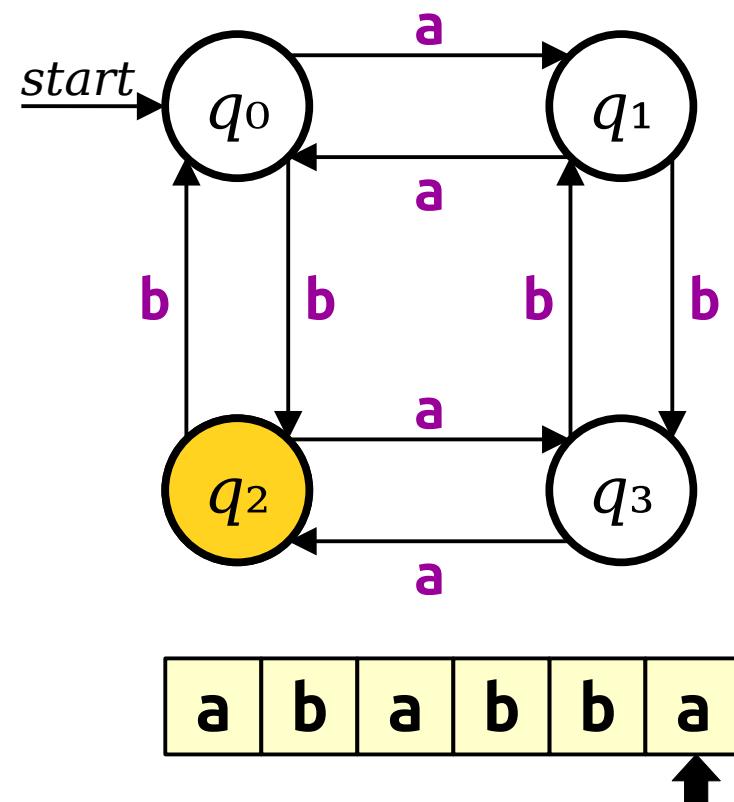
Modeling Finite Computation

- 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.



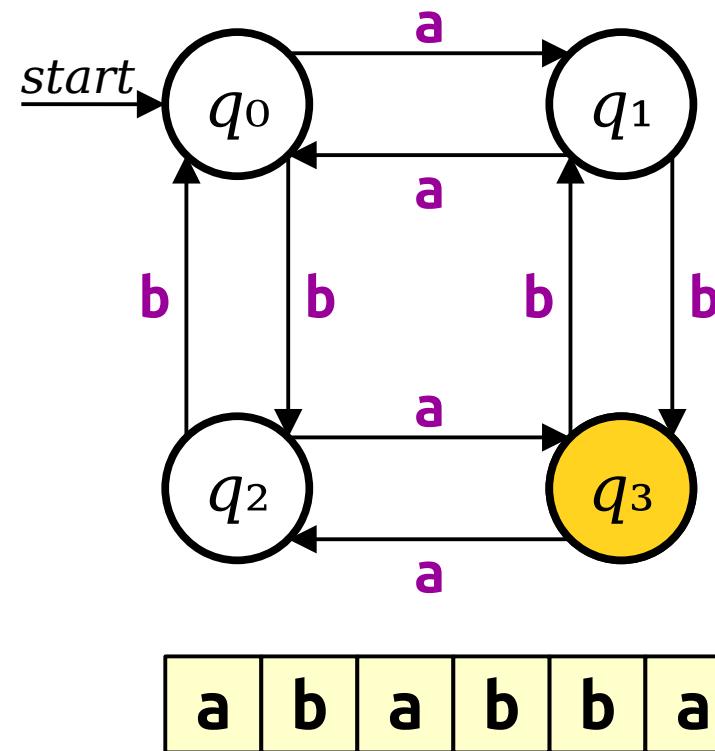
Modeling Finite Computation

- 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.



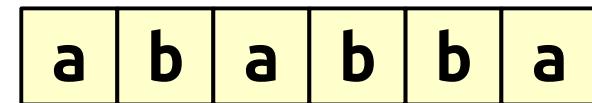
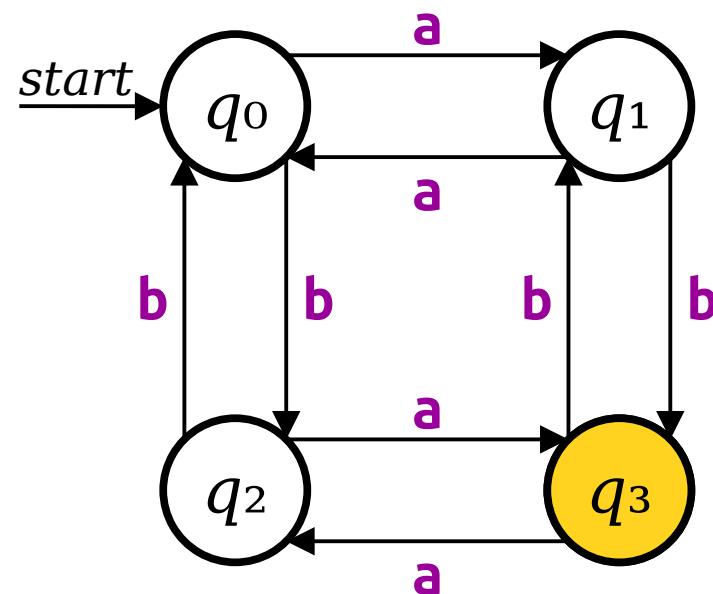
Modeling Finite Computation

- 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.



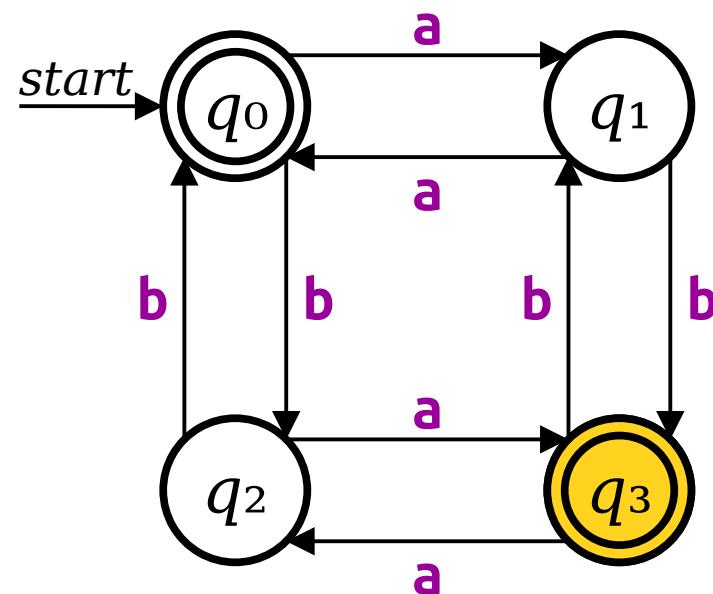
Modeling Finite Computation

- 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 me 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).



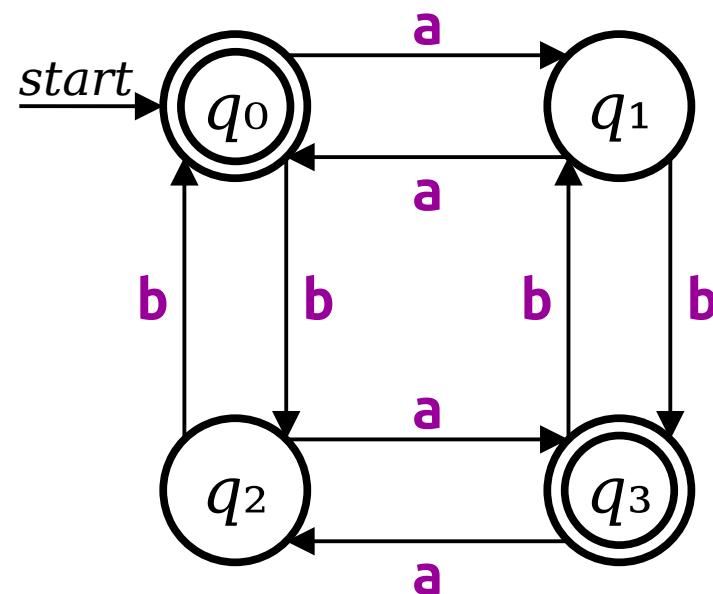
Modeling Finite Computation

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

aab

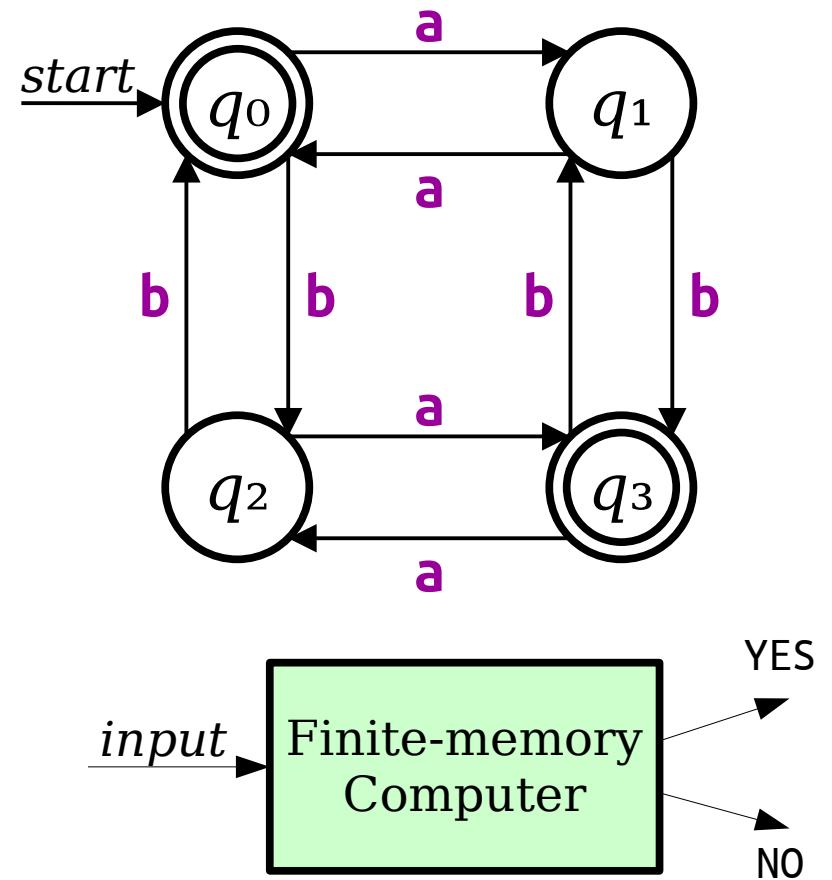
aabb

abbababba



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 aabaaaabbabaaaabaaaabbb abbababba
- The **empty string** has no characters and is denoted ϵ .
- Calling attention to an earlier point: since all strings are finite sequences of characters from Σ , you cannot have a string of infinite length.

Languages

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

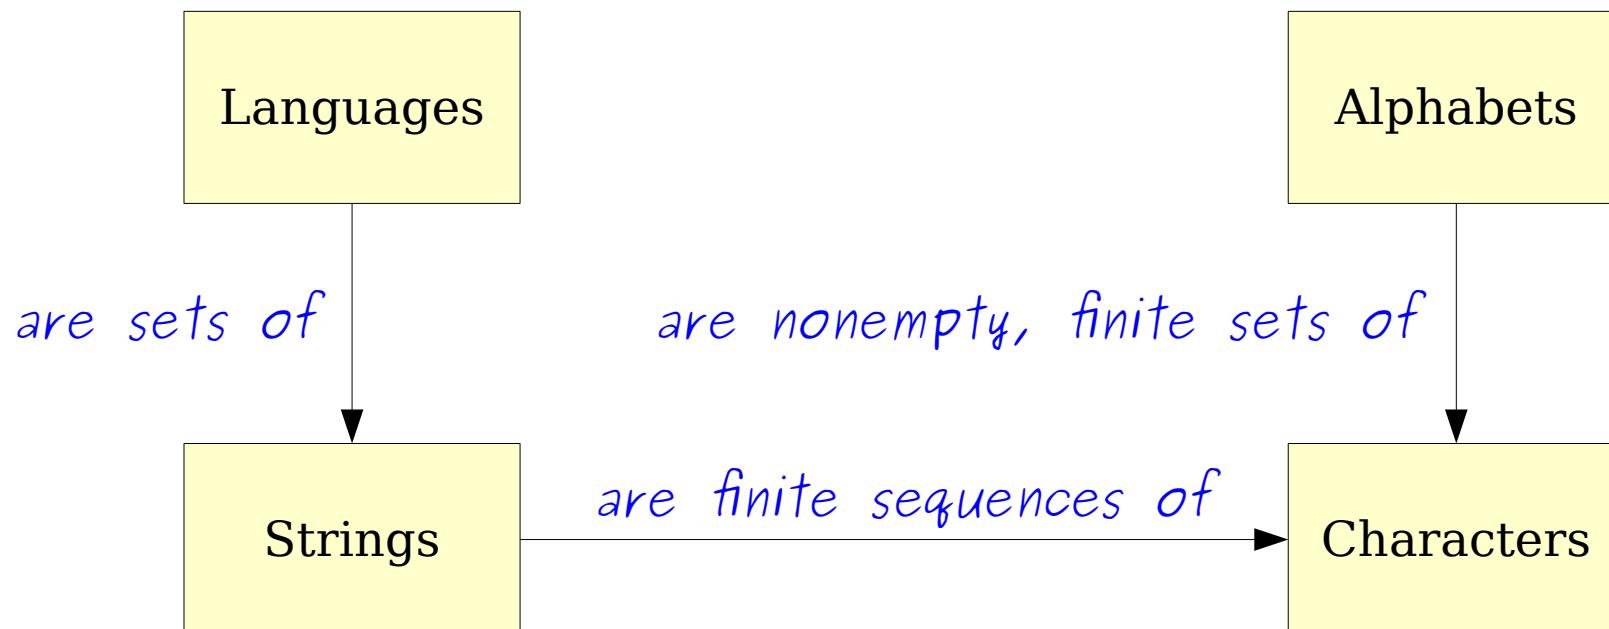
Mathematical Lookalikes

- 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!***

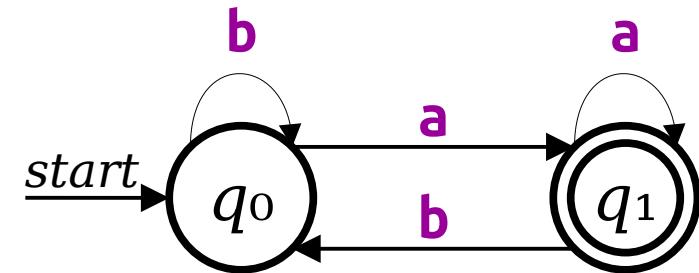
The Cast of Characters

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



Finite Automata and Languages

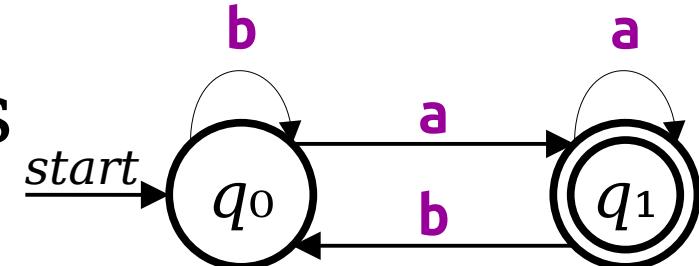
- Let A be an automaton that processes strings drawn from an alphabet Σ .
- The ***language of A*** , denoted $\mathcal{L}(A)$, is the set of strings over Σ that A accepts:



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

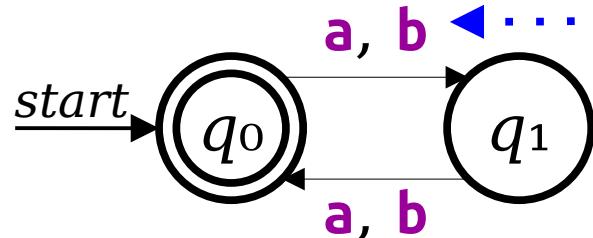
Finite Automata and Languages

- 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 $\mathcal{L}(D) = \{ w \in \{\mathbf{a}, \mathbf{b}\}^* \mid w \text{ ends in } \mathbf{a} \}$.

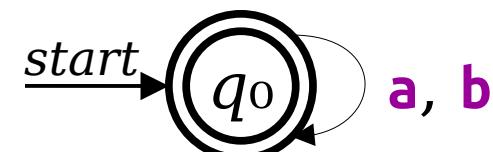
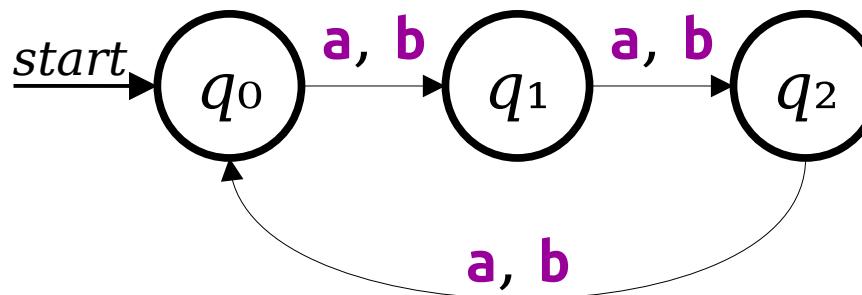
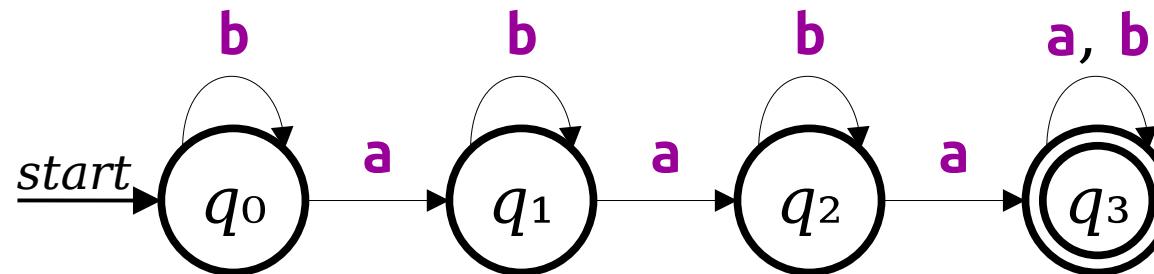


$$\mathcal{L}(A) = \{ w \in \Sigma^* \mid A \text{ accepts } w \}$$

Finite Automata and Languages



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



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

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.

Time-Out For Announcements!

Midterm Graded

- We've finished grading the midterm exam.
- Graded exams are available online through Gradescope.
- Solutions are up on the course website, along with statistics and common mistakes.
- ***Please read the solutions.*** They contain statistics, common mistakes, and advice for going forward.

Your Questions

“I did horribly on Midterm 1-- this isn't exactly a question, but any advice?”

For starters – I'm sorry to hear that!

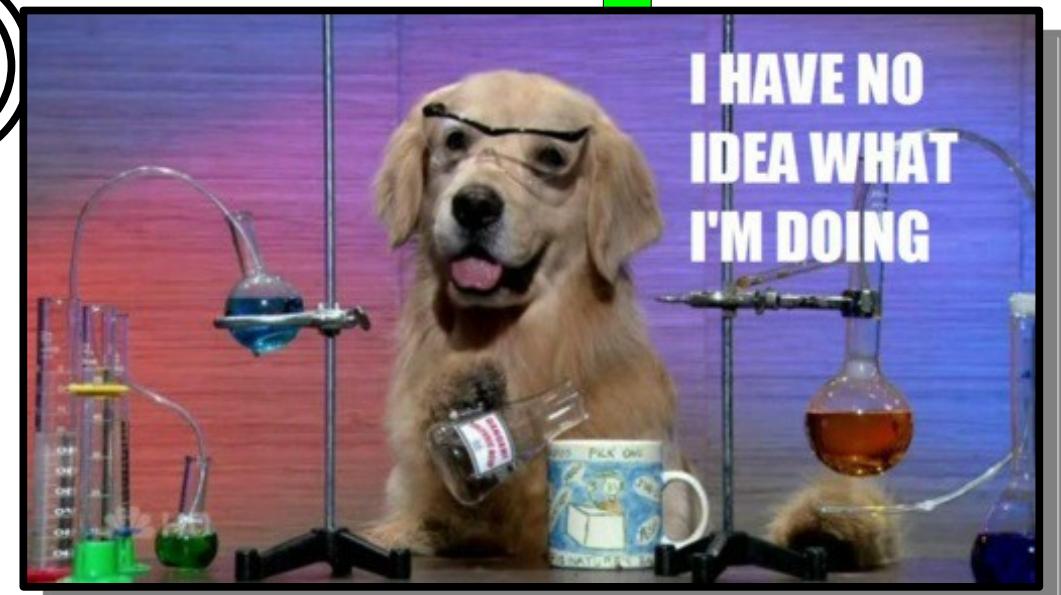
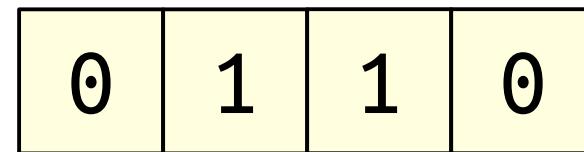
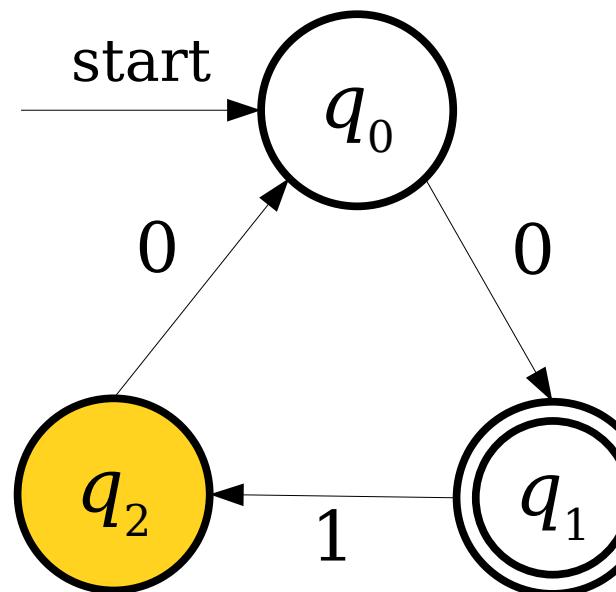
So let's talk next steps. First, if you'd like to meet with me or one of the TAs, email the staff list or post on Edstem. We'd be happy to chat one-on-one with you to figure out what's working and what could use some tweaking.

Next, review how you're engaging with the course and see what you can do differently. We have a “How to Improve in CS103” document up on the course website that offers advice about how to do this. Some of our suggestions are about ways to be proactive with your learning rather than reactive. Others are about specific things to do to make sure you're properly interpreting problem set feedback. And some has to do with how to study for exams.

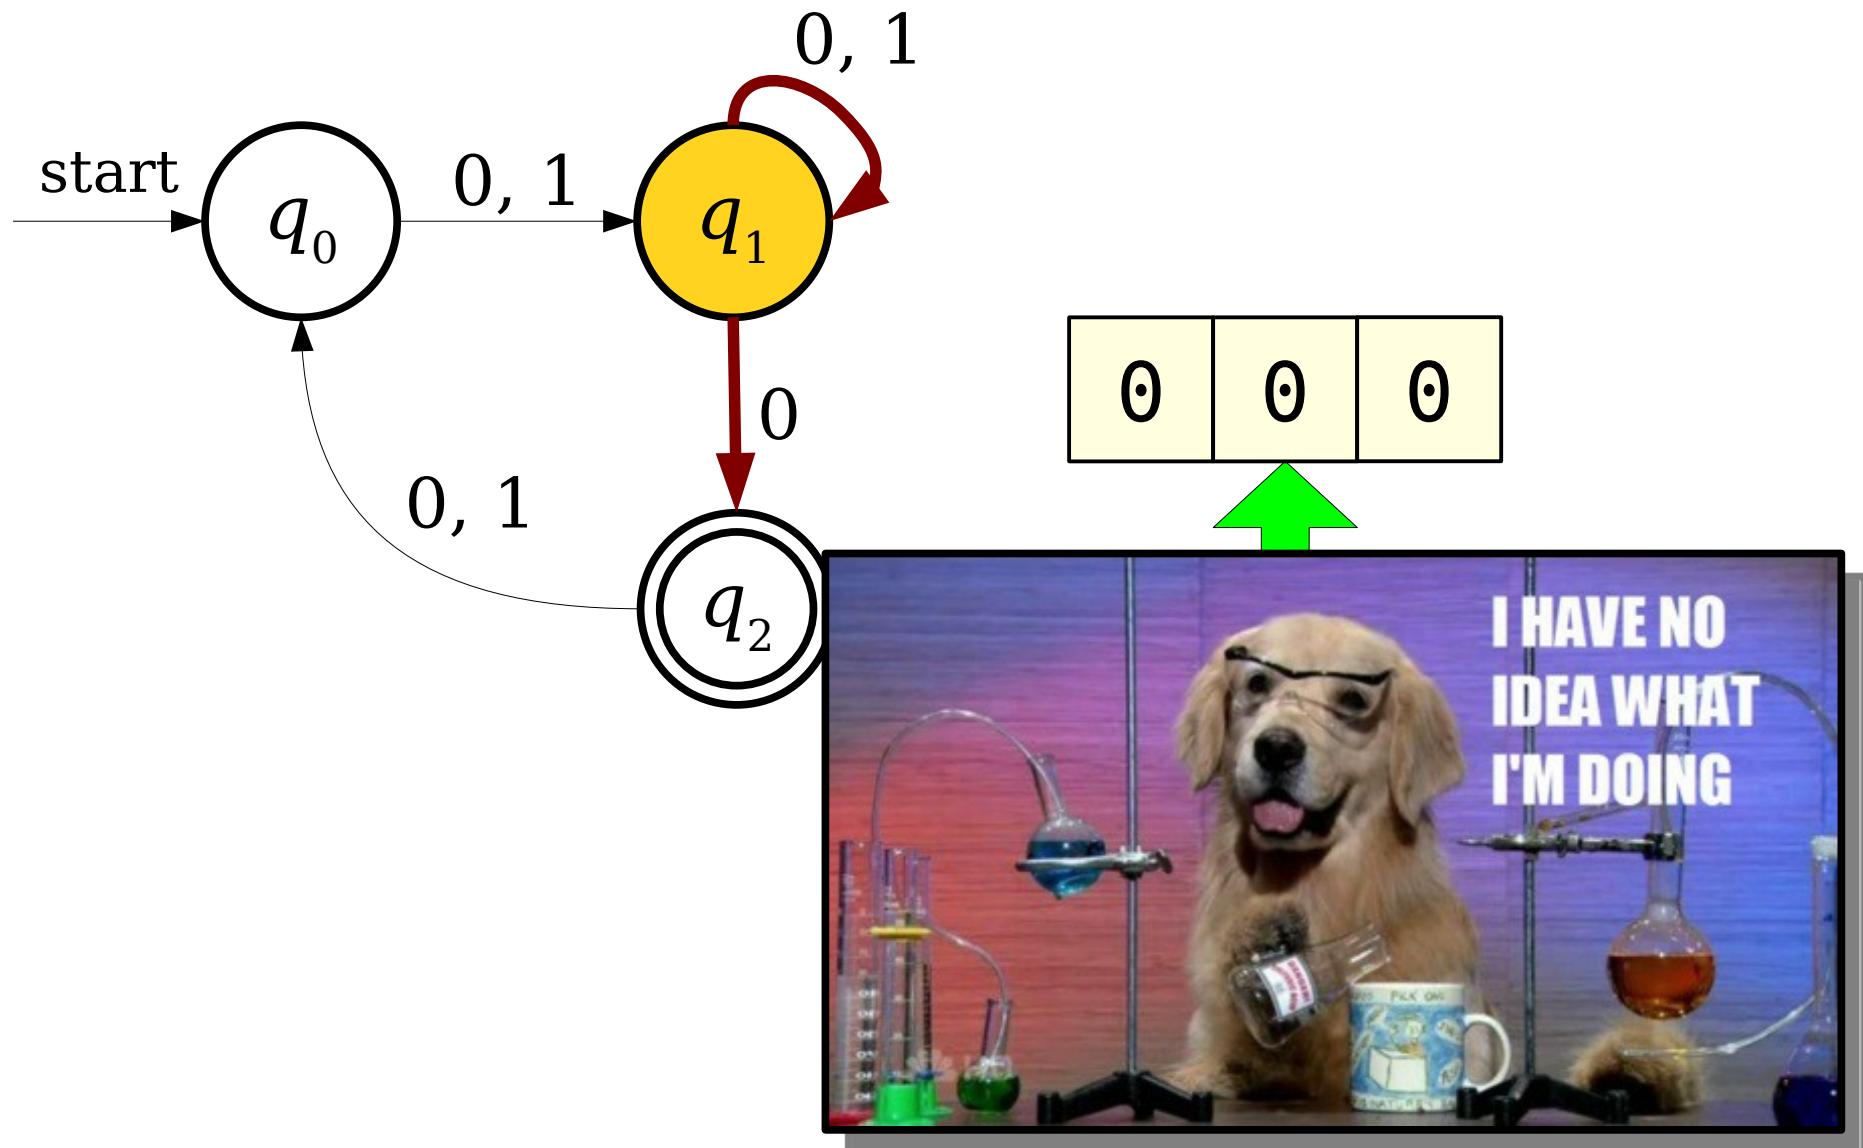
Hang in there. We're at the halfway point in the quarter and there's still plenty of time to master this material. I've seen people make incredible progress since the first midterm, and we're confident you can do this as well!

Back to CS103!

A Small Problem



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?

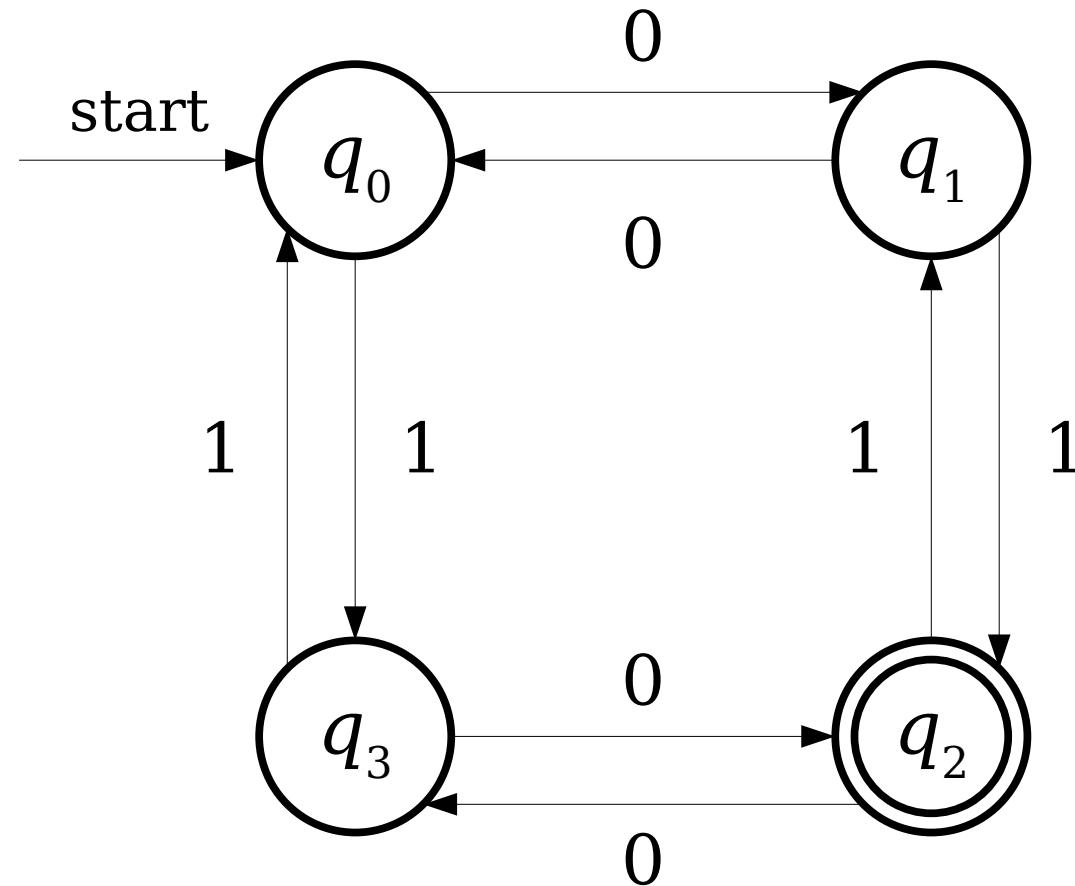
DFAs

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

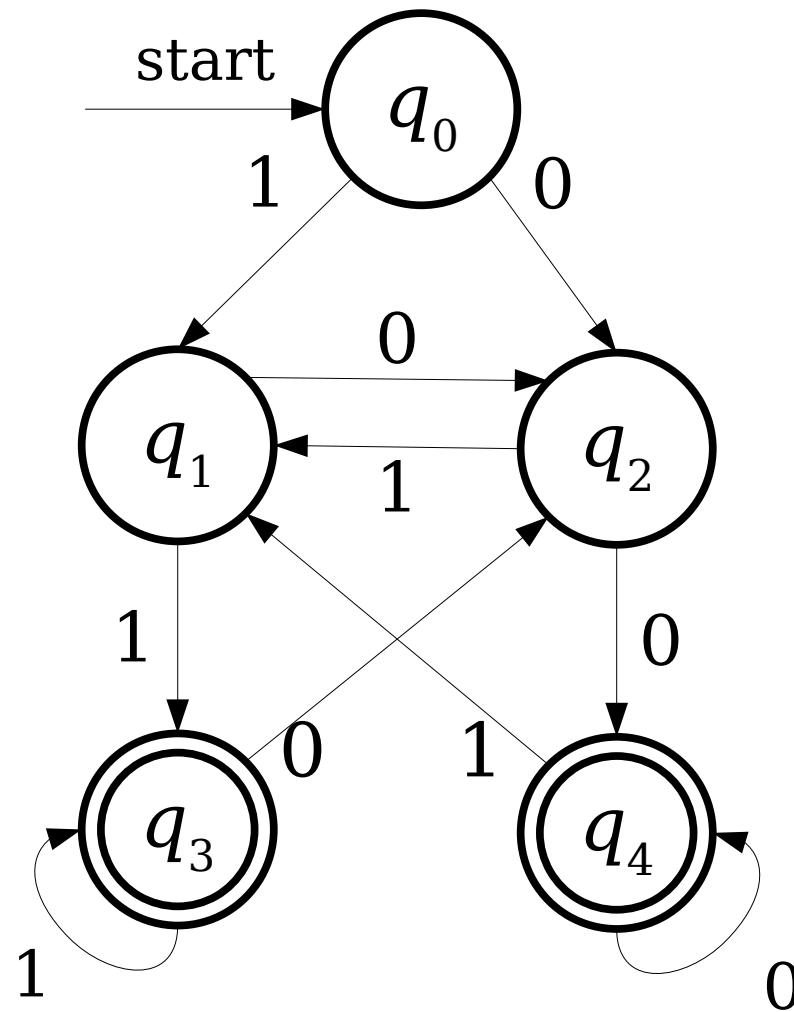
DFAs

- 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.

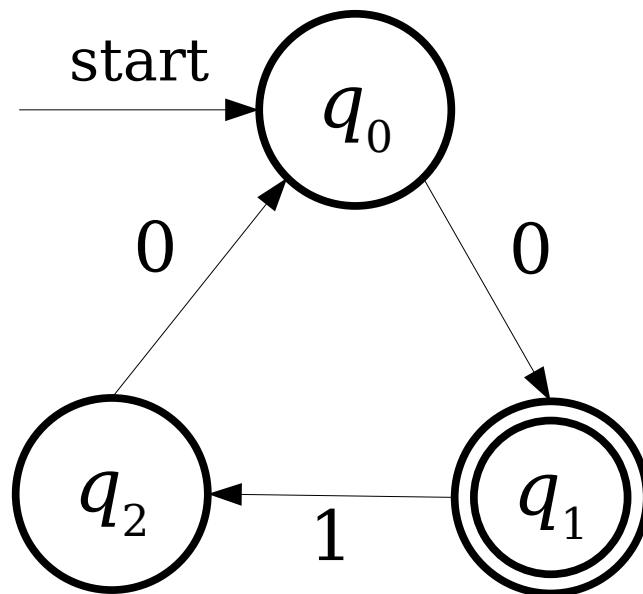
Is this a DFA over $\{0, 1\}$?



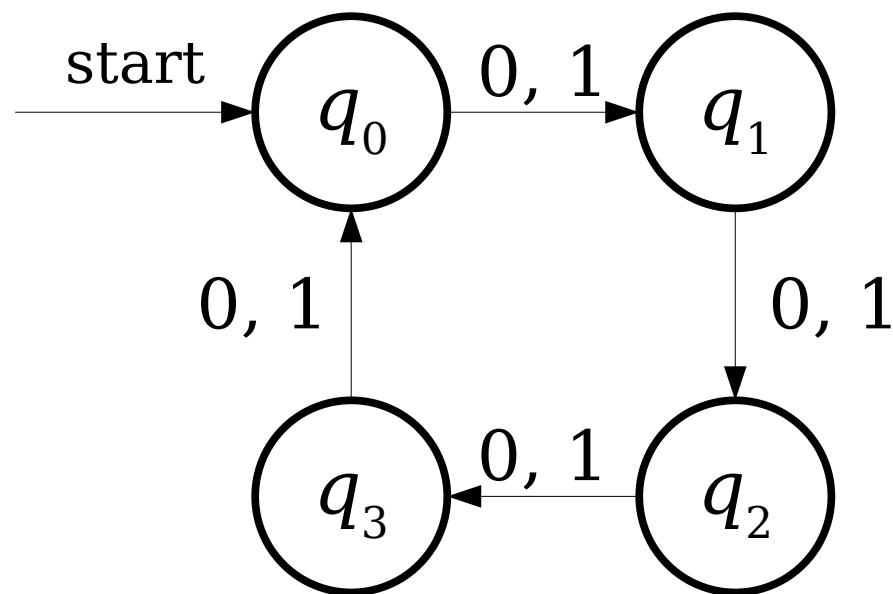
Is this a DFA over $\{0, 1\}$?



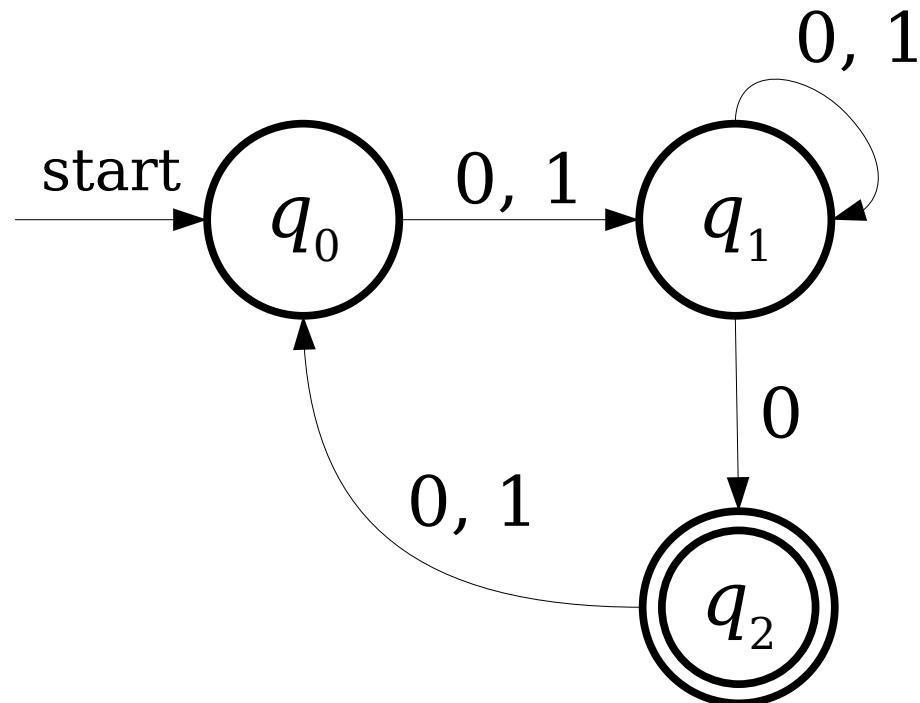
Is this a DFA over $\{0, 1\}$?



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



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



Is this a DFA?



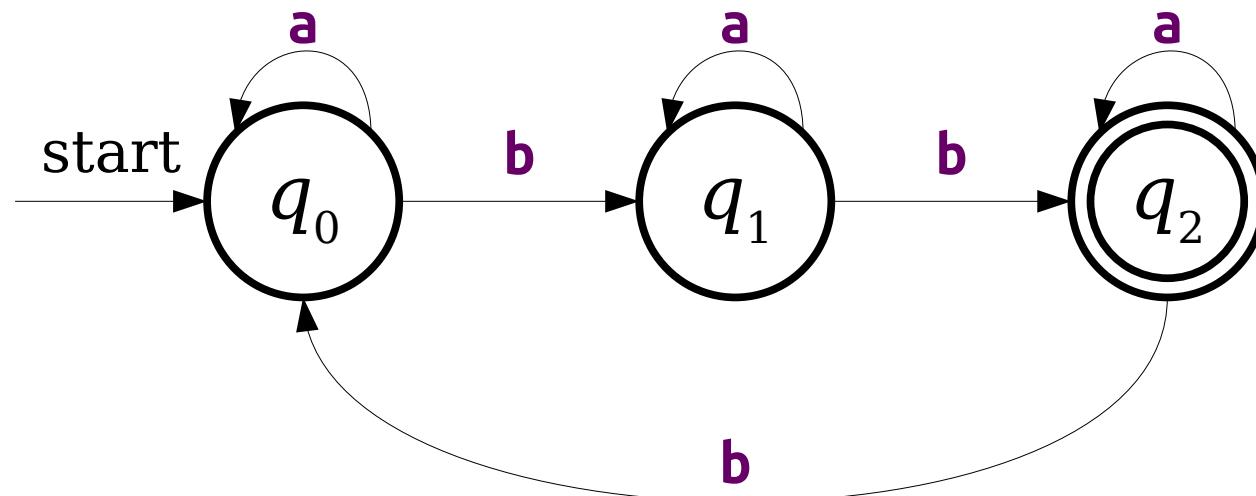
Drinking **F**amily of **A**ardvarks

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

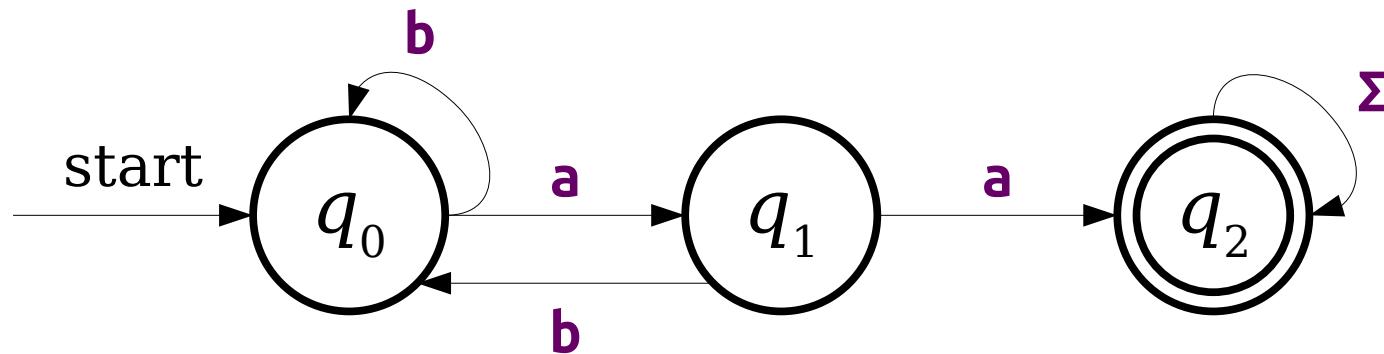
$L = \{ w \in \{a, b\}^* \mid \text{the number of } b\text{'s in } w \text{ is congruent to two modulo three} \}$



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

Recognizing Languages with DFAs

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



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:

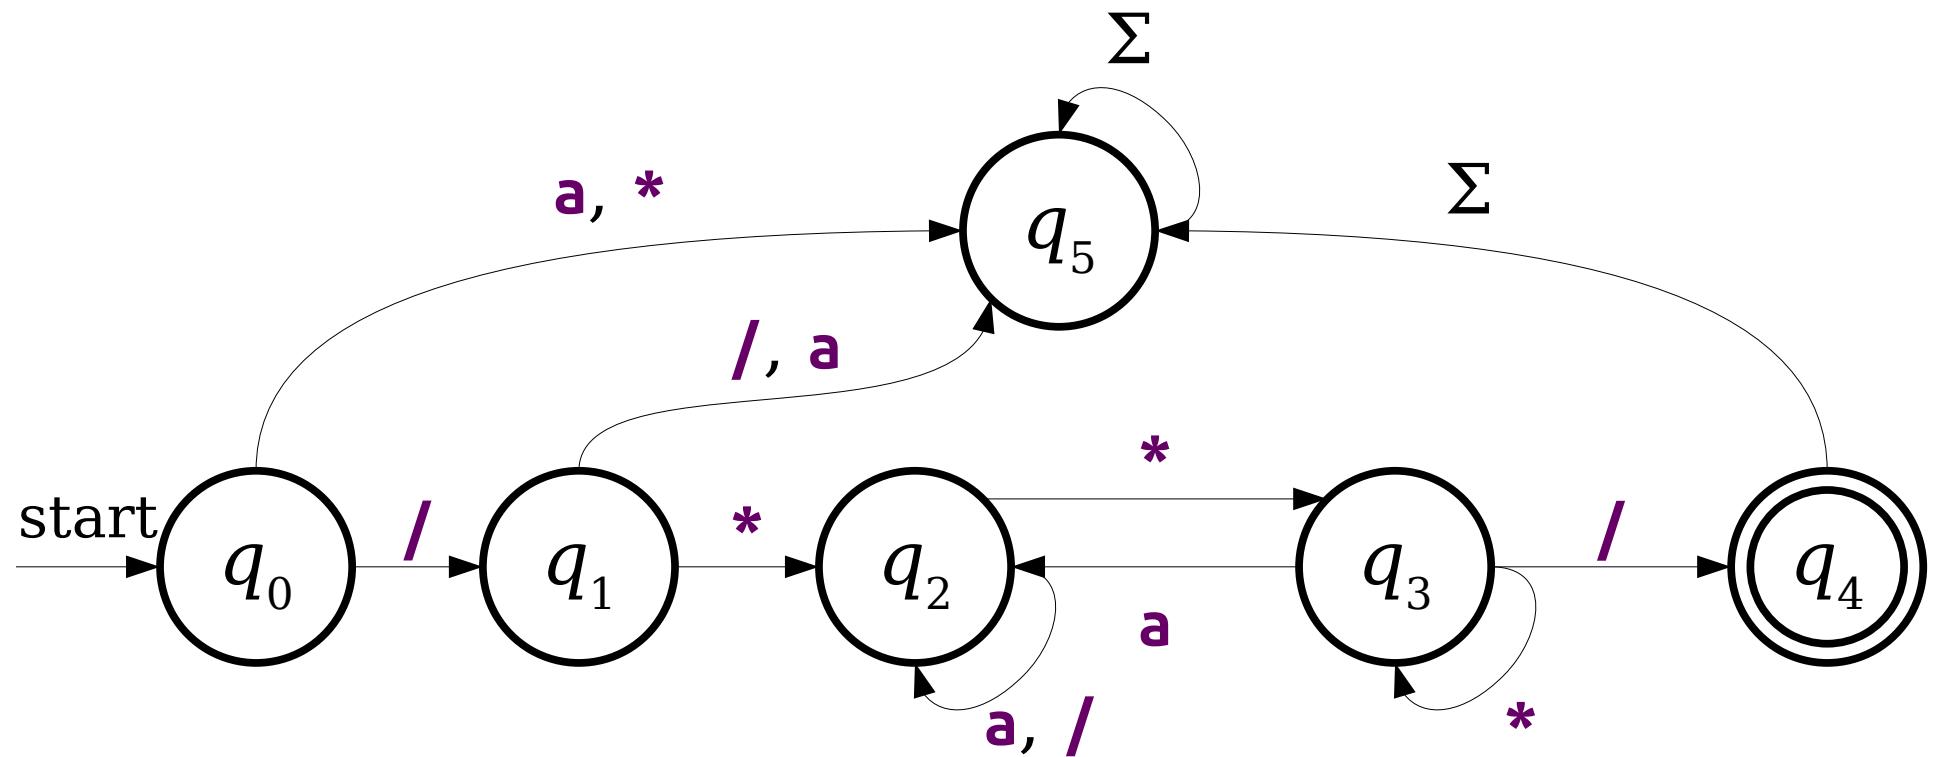
/*a*/
/**/
/***/
/*aaa*aaa*/
/*a/a*/

Rejected:

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

More Elaborate DFAs

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



Next Time

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