

Turing Machines

Part Two

Outline for Today

- ***The Church-Turing Thesis***
 - Just how powerful are TMs?
- ***What Does it Mean to Solve a Problem?***
 - Rethinking what “solving” a problem means, and two possible answers to that question.

Recap from Last Time

Turing Machines

- A ***Turing machine*** is a program that controls a tape head as it moves around an infinite tape.
- There are six commands:
 - **Move** *direction*
 - **Write** *symbol*
 - **Goto** *label*
 - **Return** *boolean*
 - **If** *symbol command*
 - **If Not** *symbol command*
- Despite their limited vocabulary, TMs are surprisingly powerful.

A Sample Turing Machine

- Here's a sample TM.
- It receives inputs over the alphabet $\Sigma = \{\mathbf{a}, \mathbf{b}\}$.
- What strings does this TM accept?
- Can you write a regex that matches precisely the strings this TM accepts?

```
Start:  
    If Not 'a' Return False  
  
Loop:  
    Move Right  
    If Not Blank Goto Loop  
    Move Left  
    Move Left  
    If Not 'b' Return False  
    Return True
```

New Stuff!

Main Questions for Today:

Just how powerful are Turing machines?

What problems can you solve with a computer?

Real and “Ideal” Computers

- A real computer has memory limitations: you have a finite amount of RAM, a finite amount of disk space, etc.
- However, as computers get more and more powerful, the amount of memory available keeps increasing.
- An *idealized computer* is like a regular computer, but with unlimited RAM and disk space. It functions just like a regular computer, but never runs out of memory.

Theorem: Turing machines are equal in power to idealized computers. That is, any computation that can be done on a TM can be done on an idealized computer and vice-versa.

Key Idea: Two models of computation are equally powerful if they can simulate each other.

Simulating a TM

- The individual commands in a TM are simple and perform only basic operations:
Move Write Goto Return If
- The memory for a TM can be thought of as a string with some number keeping track of the current index.
- To simulate a TM, we need to
 - see which line of the program we're on,
 - determine what command it is, and
 - simulate that single command.
- **Claim:** This is reasonably straightforward to do on an idealized computer.
 - My “core” logic for the TM simulator is under fifty lines of code, including comments.

Simulating a Computer

- Programming languages provide a set of simple constructs.
 - Think things like variables, arrays, loops, functions, classes, etc.
- ***Key Idea:*** If a TM is powerful enough to simulate each of these individual pieces, it's powerful enough to simulate anything a real computer can do.

Can TMs do: Loops?

- We've seen TMs use loops to solve problems.
 - Our $\{ a^n b^n \mid n \in \mathbb{N} \}$ TM repeatedly pulls off the first and last character from the string.
 - Our sorting TM repeatedly finds **ba** and replaces it with **ab**.
- In some sense, the existence of Goto and labels means that TMs have loops.
- Hopefully, it's not too much of a stretch to think that TMs can do `while` loops, `for` loops, etc.

Can TMs do: arithmetic?

- We've seen TMs that perform basic arithmetic.
 - We can check if two numbers are equal.
 - We can check if a number is a Fibonacci number.
- Hopefully, it's not too much of a stretch to believe we could also do addition and subtraction, compute powers of numbers, do ceilings and floors, etc.

Can TMs do: variables?

- We've seen TMs that maintain variables.
 - You can think of our TM for $\{ \mathbf{a}^n \mathbf{b}^n \mid n \in \mathbb{N} \}$ as storing two variables – one that counts a number of **a**'s, and one that counts a number of **b**'s.
 - Our TM for Fibonacci numbers kinda sorta ish tracks the last two Fibonacci numbers, plus the length of the input string.
- It's a bit larger of a jump to make, but hopefully you're comfortable with the idea that TMs, in principle, can maintain variables.

Can TMs do: helper functions?

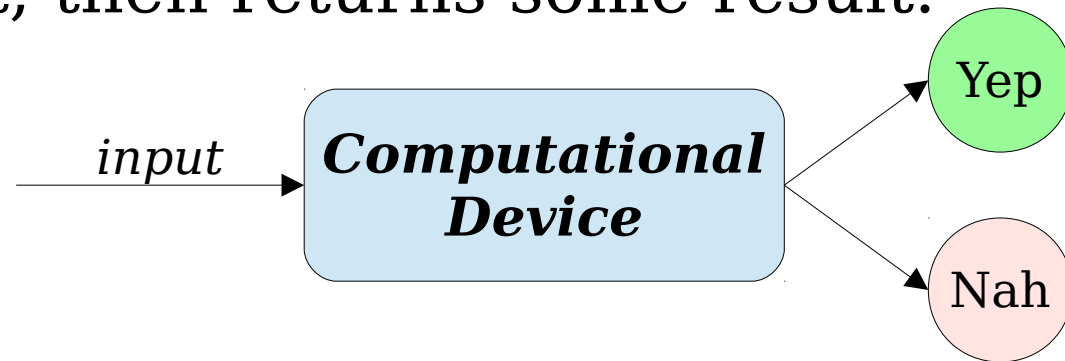
- We've seen TMs with helper functions.
 - We saw how to check for equal numbers of **a**'s and **b**'s by first sorting the string, then checking if the string has the form **aⁿbⁿ**.
 - We can check if a decimal number is a Fibonacci number by converting it to unary, then running our unary Fibonacci checker.
- Hopefully you're comfortable with the idea that a TM could have multiple "helper functions" that work together to solve some larger problem.

What Else Can TMs Do?

- Maintain strings and arrays.
 - Store their elements separated with some special separator character.
- Support pointers.
 - Maintain an array of what's in memory, where each item is tagged with its "memory address."
- Support function call and return.
 - It's hard, but you can do this if you can do helper functions and variables.

A Leap of Faith

- **Claim:** A TM is powerful enough to simulate any computer program that gets an input, processes that input, then returns some result.



- The resulting TM might be colossal, or really slow, or both, but it would still faithfully simulate the computer.
- We're going to take this as an article of faith in CS103. If you curious for more details, come talk to me after class.

Can a TM Work With...

“cat pictures?”

Sure! A picture is just a 2D array of colors, and a color can be represented as a series of numbers.



Can a TM Work With...

~~“cat pictures?”~~
“cat videos?”

If you think about it, a
video is just a series of
pictures!



Can a TM Work With...

“music?”

Sure! Music is encoded as a compressed waveform. That's just a list of numbers.

“deep learning?”

Sure! That's just applying a bunch of matrices and nonlinear functions to some input.

Just how powerful *are* Turing machines?

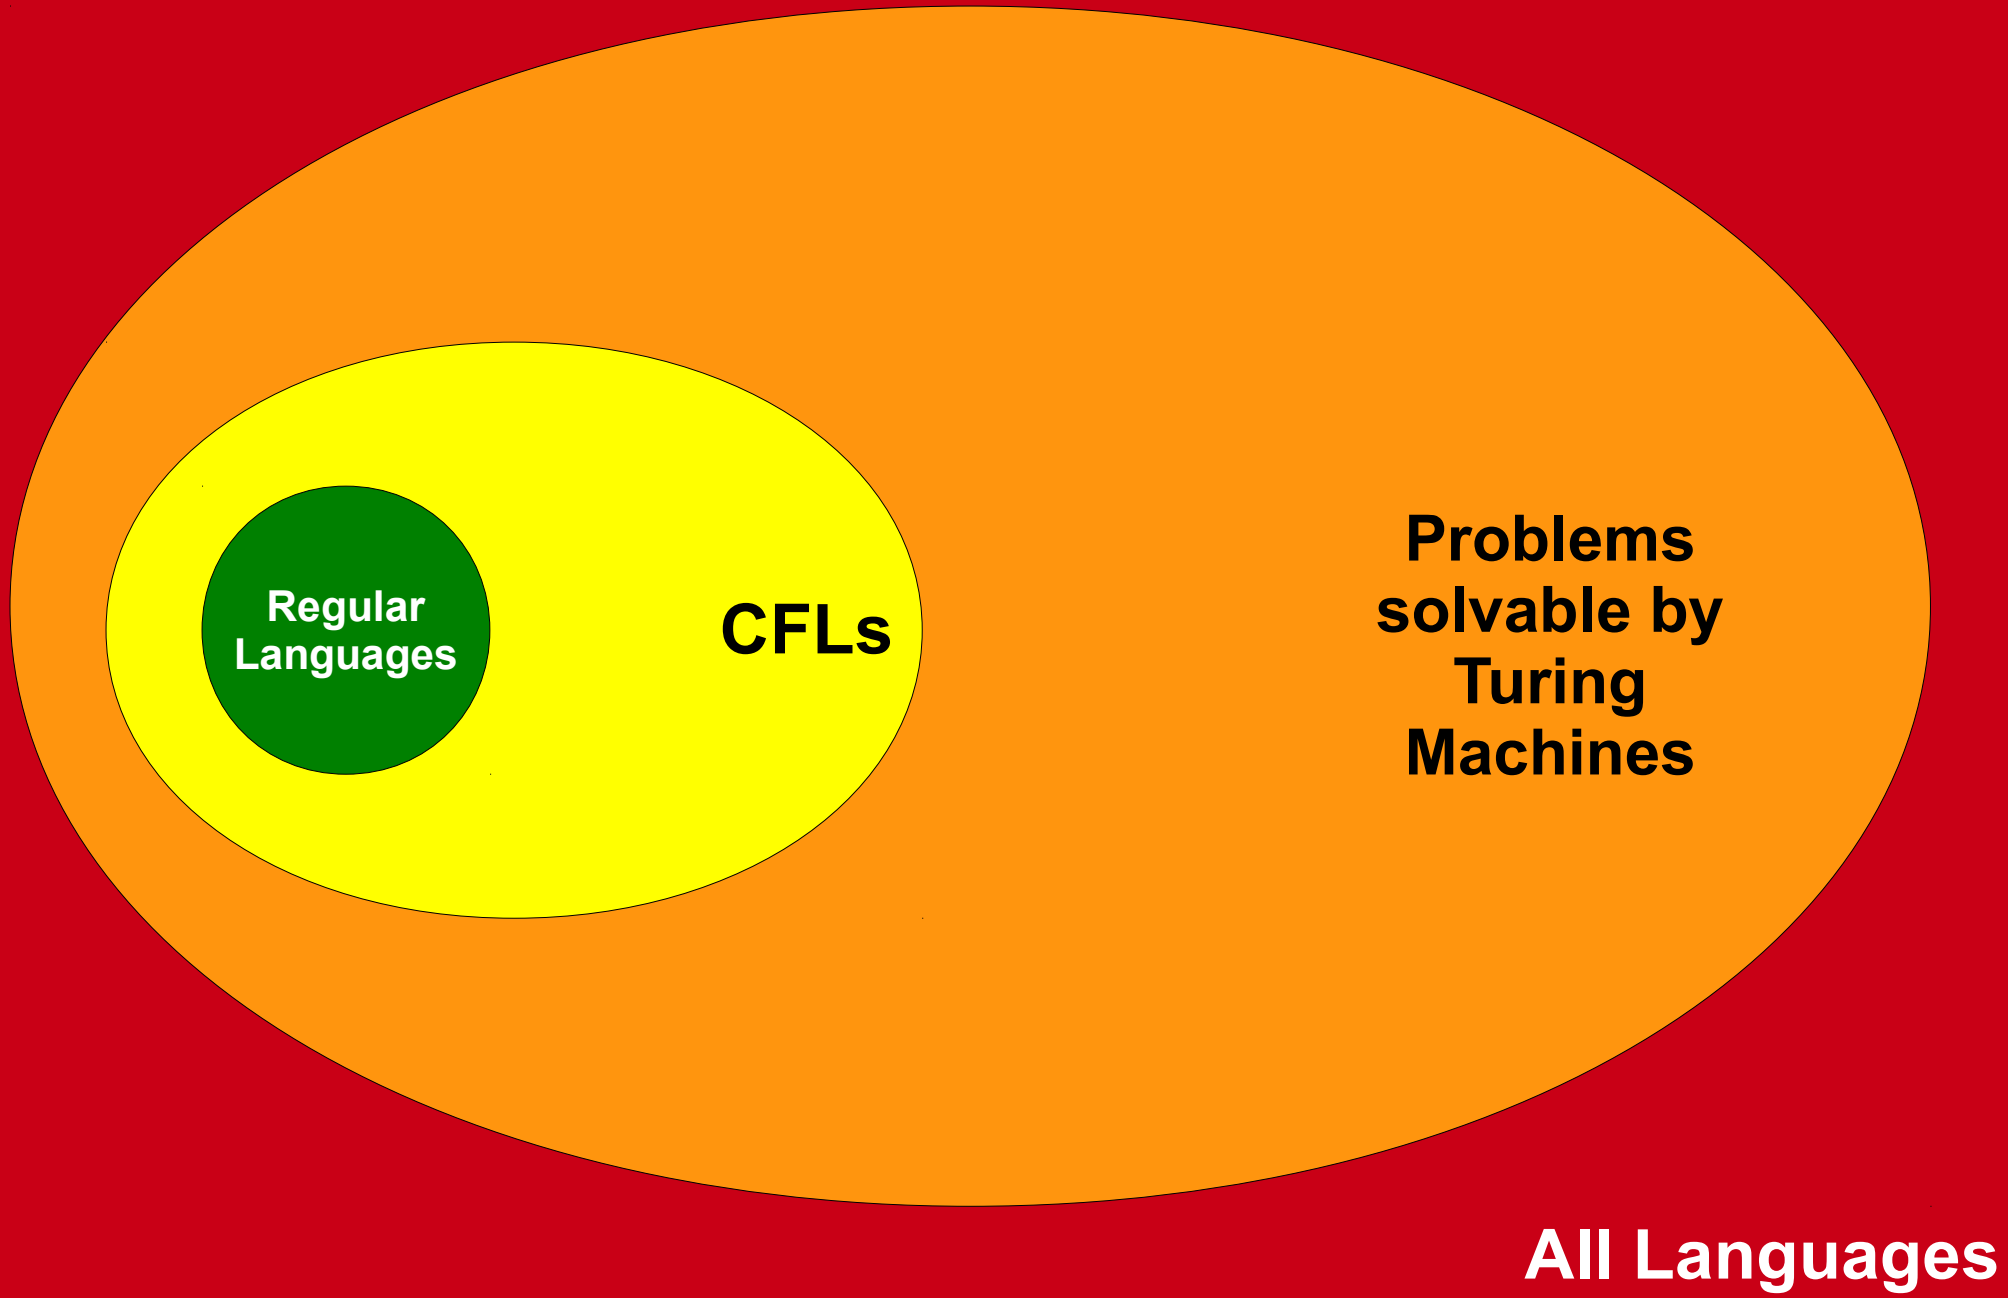
Effective Computation

- An ***effective method of computation*** is a form of computation with the following properties:
 - The computation consists of a set of steps.
 - There are fixed rules governing how one step leads to the next.
 - Any computation that yields an answer does so in finitely many steps.
 - Any computation that yields an answer always yields the correct answer.
- This is not a formal definition. Rather, it's a set of properties we expect out of a computational system.

The *Church-Turing Thesis* claims that
*every effective method of computation
is either equivalent to or weaker than
a Turing machine.*

“This is not a theorem – it is a
falsifiable scientific hypothesis.
And it has been thoroughly
tested!”

- Ryan Williams



Regular Languages

CFLs

Problems solvable by Turing Machines

All Languages

TMs and Computation

- Because Turing machines have the same computational powers as regular computers, we can (essentially) reason about Turing machines by reasoning about actual computer programs.
- **Going forward, we're going to switch back and forth between TMs and more C++ or python-like computer programs based on whatever is most appropriate.**
- In fact, our eventual proofs about the existence of impossible problems will involve a good amount of pseudocode. Stay tuned for details!

Decidability and Recognizability

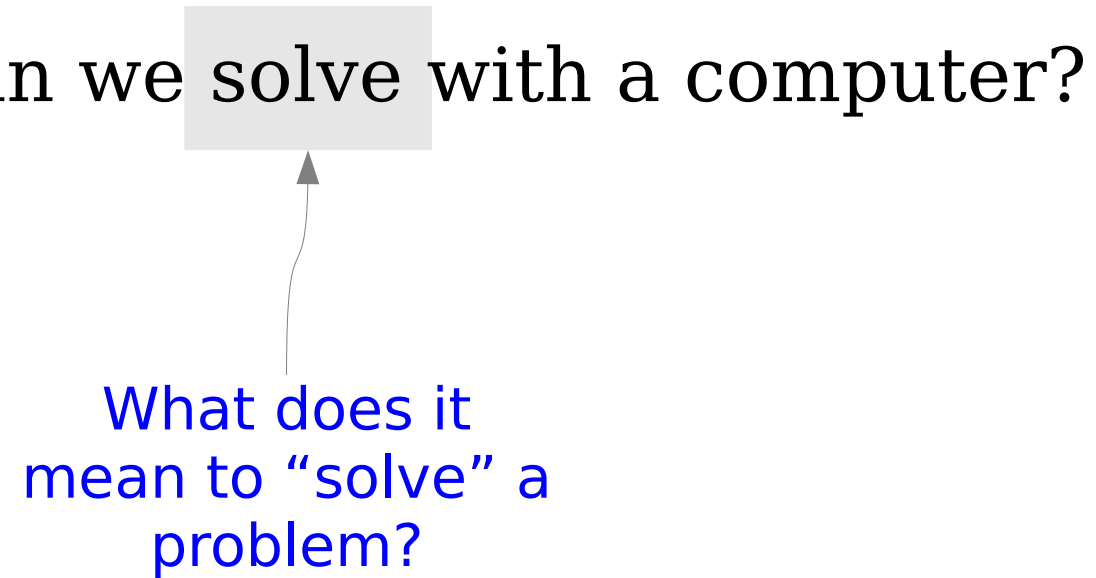
What problems can we solve with a computer?

What kind of
computer?

A diagram consisting of a blue question 'What kind of computer?' at the bottom. A thin grey line extends upwards from the question, curving slightly to the right, and ending in a small grey arrowhead pointing to the word 'computer' in the main question above. The word 'computer' in the main question is highlighted with a light grey rectangular background.

What problems can we solve with a computer?

What does it
mean to “solve” a
problem?

A diagram consisting of a main question at the top and a sub-question below it. The main question is "What problems can we solve with a computer?". The word "solve" in this question is highlighted with a light gray rectangular box. A thin, light gray arrow points from the bottom center of this box down to the sub-question "What does it mean to 'solve' a problem?". The sub-question is written in blue text.

A Simple Turing Machine

- Here's a TM.
- It receives inputs over the alphabet $\Sigma = \{\mathbf{a}, \mathbf{b}\}$.
- What strings does this TM accept?
- What happens when you give it these strings: \mathbf{a} , \mathbf{aaa} , \mathbf{bb} , \mathbf{baaa}


```
Start:  
    If 'a' Return True  
Loop:  
    Move Right  
    If Not 'b' Goto Loop
```

An Important Observation

- Unlike finite automata, which automatically halt after all the input is read, TMs keep running until they explicitly return true or return false.
- As a result, it's possible for a TM to run forever without accepting or rejecting.
- This leads to several important questions:
 - How do we formally define what it means to build a TM for a language?
 - What implications does this have about problem-solving?

Very Important Terminology


- Let M be a Turing machine.
- M **accepts** a string w if it returns true on w .
- M **rejects** a string w if it returns false on w .
- M **halts** on a string w if it returns on w (i.e., we don't care if it returns true or false, just that it returns at all).
- M **loops infinitely** (or just **loops**) on a string w if when run on w it neither returns true nor returns false (i.e., it doesn't halt).



Very Important Terminology

- Let M be a Turing machine.
- M **accepts** a string w if it returns true on w .
- M **rejects** a string w if it returns false on w .
- M **halts** on a string w if it returns on w (i.e., we don't care if it returns true or false, just that it returns at all).
- M **loops infinitely** (or just **loops**) on a string w if it neither returns true nor returns false on w .

To be in the language of M , a string must be accepted by M . No answer/looping is the same as rejection, in terms of meaning that the string is not in the language of M .



Very Important Terminology

- Let M be a Turing machine.
- M **accepts** a string w if it returns true on w .
- M **rejects** a string w if it returns false on w .
- M **halts** on a string w if it returns on w (i.e., we don't care if it returns true or false, just that it returns at all).
- M **loops infinitely** (or just **loops**) on a string w if it neither returns true nor returns false on w .

To be in the language of M , a string must be accepted by M . No answer/looping is the same as rejection, in terms of meaning that the string is not in the language of M .

Turing Machines practice affirmative consent! Nice!

does not reject

does not accept

Accept

Loop

Reject

halts

Recognizers and Recognizability

- A TM M is called a **recognizer** for a language L over Σ if the following statement is true:

$$\forall w \in \Sigma^*. (w \in L \leftrightarrow M \text{ accepts } w)$$

- If $w \in L$, then, eventually, M will accept w .
- For all strings not in L , M may reject or may infinite loop.
- Usefulness of this as a computer:
 - If you don't already know whether $w \in L$, running M on w may never tell you anything.
 - M might loop on w - but you can't differentiate between "it'll never give an answer" and "just wait a bit more."
 - This is a very weak definition of "solving a problem," but we are after all exploring the outer extremes of what computers can do

```
bool pizkwat(string input) {  
    return false;  
}
```

```
bool squigglebah(string input) {  
    while (true) {  
        // do nothing  
    }  
}
```

```
bool moozle(string input) {  
    int oot = 1;  
    while (input.size() != oot) {  
        oot += oot;  
    }  
    return true;  
}
```


```
bool humblegwah(string input) {  
    if (input.size() % 2 != 0) return false;  
  
    for (int i = 0; i < input.size() / 2; i++) {  
        if (input[2 * i] != input[2 * i + 1]) {  
            return false;  
        }  
    }  
  
    return true;  
}
```

$\forall w \in \Sigma^*. (w \in L \leftrightarrow M \text{ accepts } w)$

Each of these pieces of code is a recognizer for some language.
What language does each recognizer recognize?

Deciders and Decidability

- Some, but not all, TMs have the following property: the TM halts on all inputs.
- Such a TM is called a “Decider.”
 - All deciders are recognizers.
 - Not all recognizers are deciders.



Deciders and Decidability

- A TM M is called a **decider** for a language L over Σ if the following statements are true:

$\forall w \in \Sigma^*. M \text{ halts on } w.$

$\forall w \in \Sigma^*. (w \in L \leftrightarrow M \text{ accepts } w)$

- In other words, M accepts all strings in L and rejects all strings not in L .
- In other words, M is a recognizer for L , and M halts on all inputs.
- No matter what input string w you give a decider TM, you will always get a clear yes or no answer.

```
bool pizkwat(string input) {  
    return false;  
}
```

```
bool squigglebah(string input) {  
    while (true) {  
        // do nothing  
    }  
}
```

```
bool moozle(string input) {  
    int oot = 1;  
    while (input.size() != oot) {  
        oot += oot;  
    }  
    return true;  
}
```

```
bool humblegwah(string input) {  
    if (input.size() % 2 != 0) return false;  
  
    for (int i = 0; i < input.size() / 2; i++) {  
        if (input[2 * i] != input[2 * i + 1]) {  
            return false;  
        }  
    }  
  
    return true;  
}
```

$\forall w \in \Sigma^*. M$ halts on w

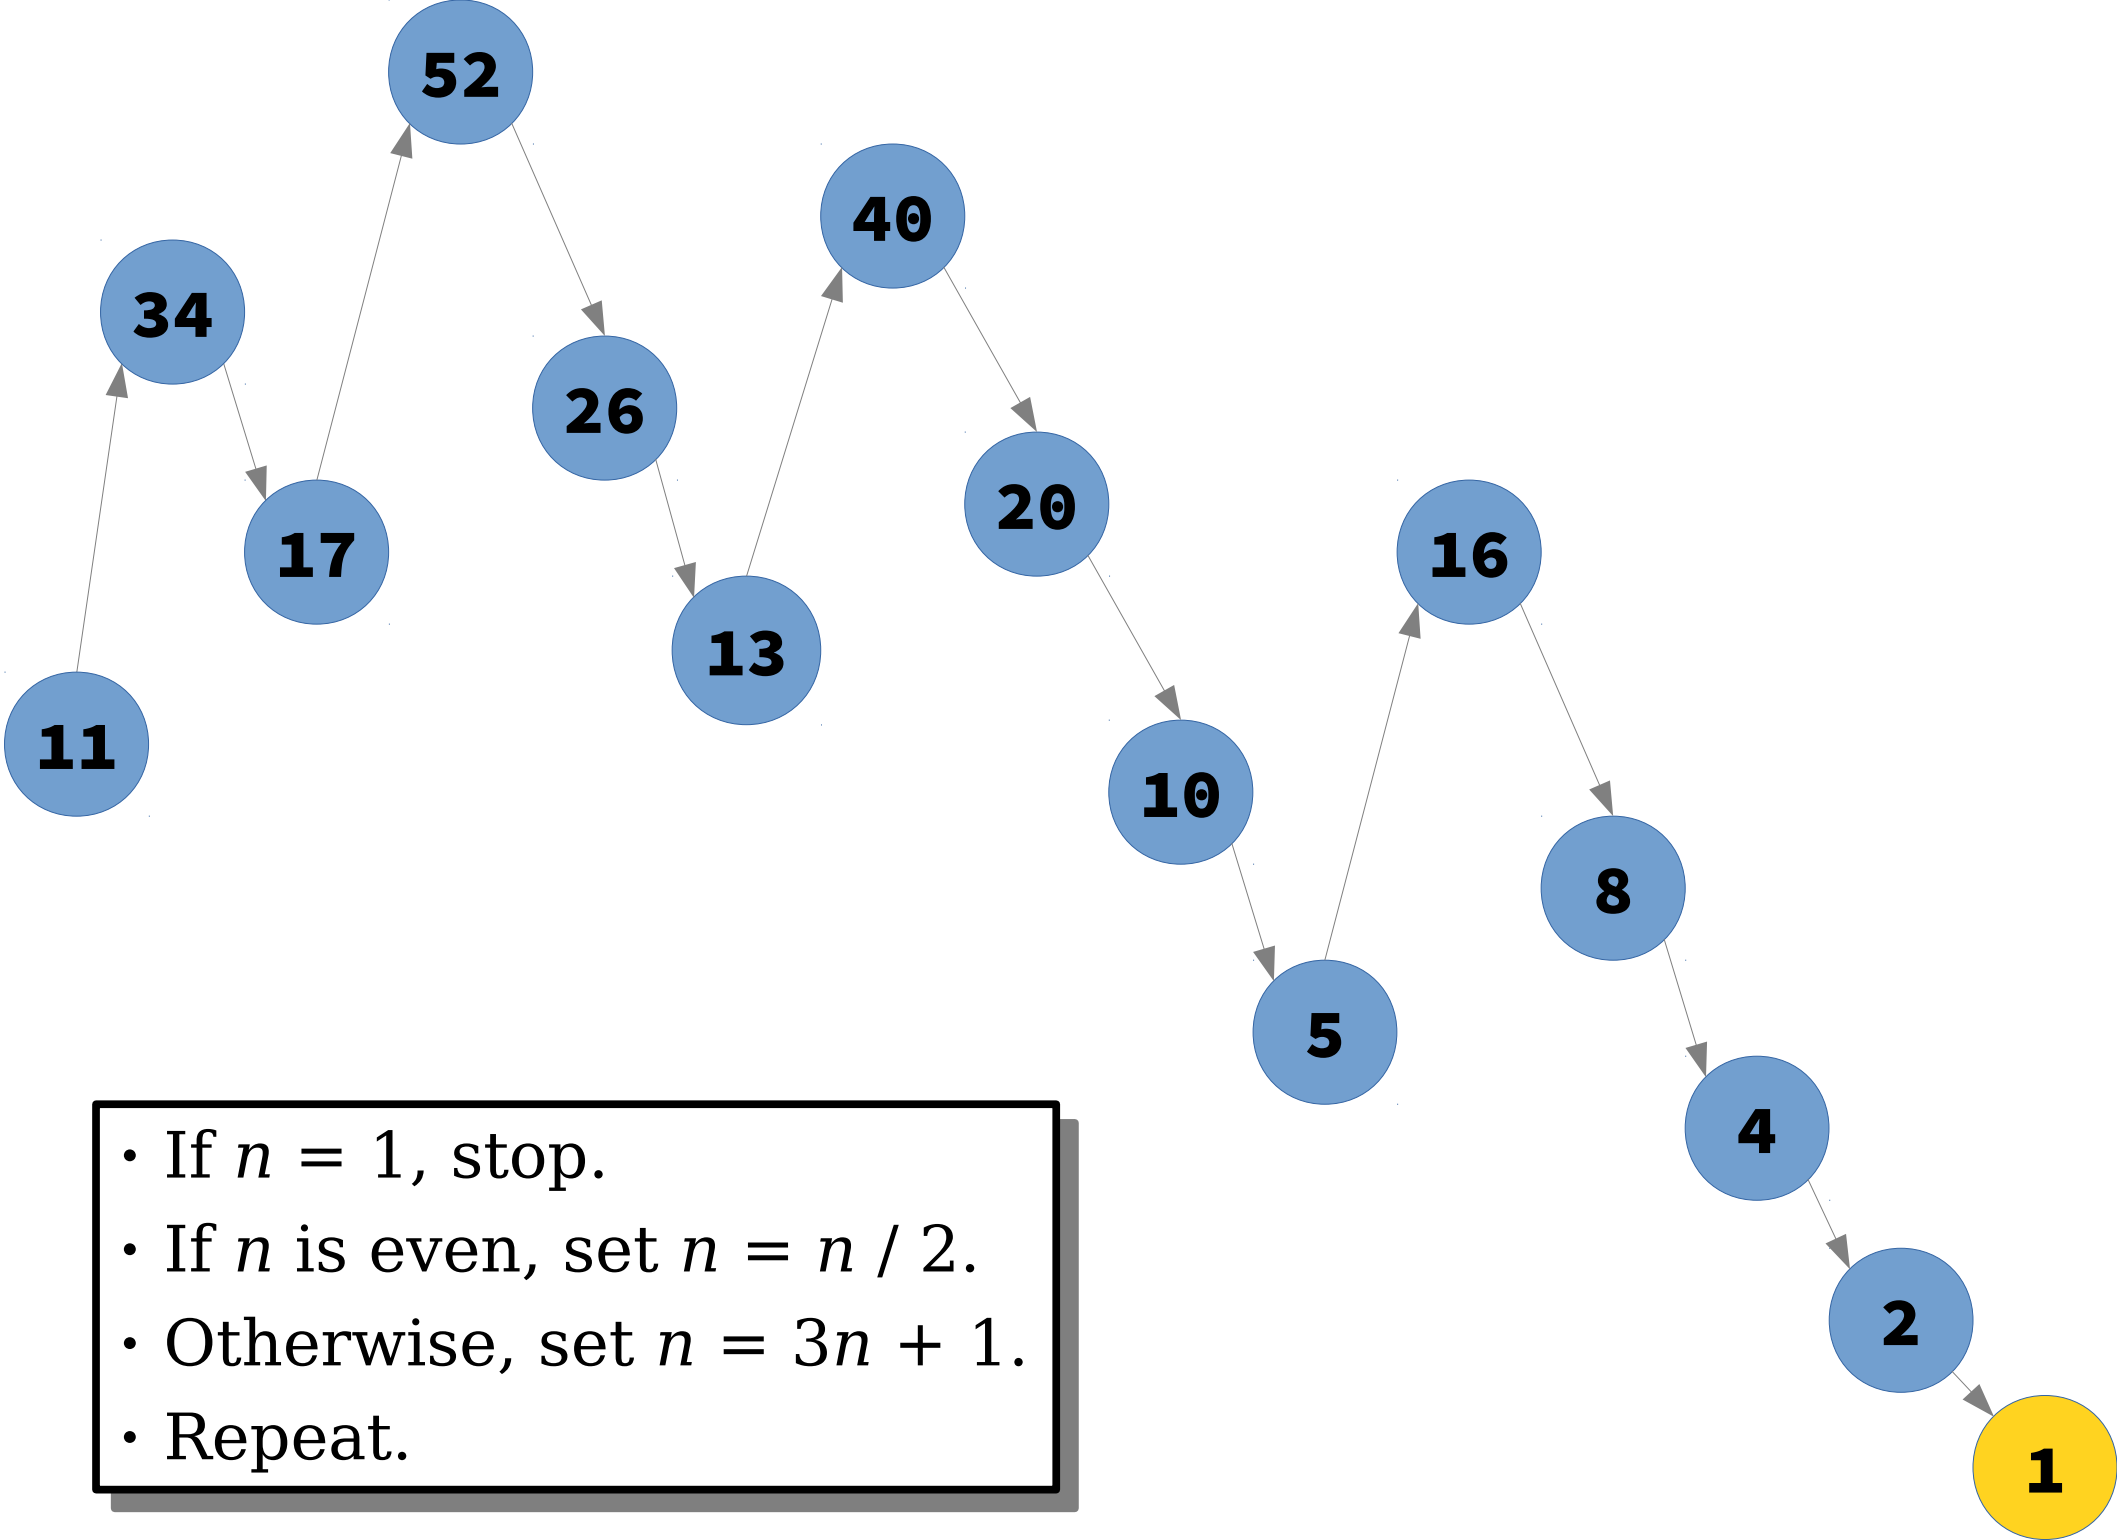
$\forall w \in \Sigma^*. (w \in L \leftrightarrow M \text{ accepts } w)$

Each piece of code is a recognizer for a language.
Which are deciders?

A Tricky TM

The Hailstone Sequence

- Consider the following procedure, starting with some $n \in \mathbb{N}$, where $n > 0$:
 - If $n = 1$, you are done.
 - If n is even, set $n = n / 2$.
 - Otherwise, set $n = 3n + 1$.
 - Repeat.
- **Question:** Given a natural number $n > 0$, does this process terminate?



- If $n = 1$, stop.
- If n is even, set $n = n / 2$.
- Otherwise, set $n = 3n + 1$.
- Repeat.

The Hailstone Sequence

- Consider the following procedure, starting with some $n \in \mathbb{N}$, where $n > 0$:
 - If $n = 1$, you are done.
 - If n is even, set $n = n / 2$.
 - Otherwise, set $n = 3n + 1$.
 - Repeat.
- Does the Hailstone Sequence terminate for...
 - $n = 5$?
 - $n = 20$?
 - $n = 7$?
 - $n = 27$?

The Hailstone Sequence

- Let $\Sigma = \{\mathbf{a}\}$ and consider the language
 $L = \{ \mathbf{a}^n \mid n > 0 \text{ and the hailstone sequence terminates for } n \}$.
- Could we build a TM for L ?

The Hailstone Turing Machine

- We can build a TM that works as follows:
 - If the input is ε , reject.
 - While the string is not **a**:
 - If the input has even length, halve the length of the string.
 - If the input has odd length, triple the length of the string and append a **a**.
 - Accept.

The Collatz Conjecture

- It is *unknown* whether this process will terminate for all natural numbers.
- In other words, no one knows whether the TM described in the previous slides will always stop running!
- The conjecture (unproven claim) that the hailstone sequence always terminates is called the ***Collatz Conjecture***.
- This problem has eluded a solution for a long time. The influential mathematician Paul Erdős is reported to have said “Mathematics may not be ready for such problems.”

Hailstone Decider?

- The hailstone TM M we saw earlier is a **recognizer** for the language

$$L = \{ \mathbf{a}^n \mid n > 0 \text{ and the hailstone sequence terminates for } n \}.$$

- If the hailstone sequence terminates for n , then M accepts \mathbf{a}^n . If it doesn't, then M does not accept \mathbf{a}^n .
- *Is it also a decider?*

Hailstone Decider?

- The hailstone TM M we saw earlier is a *recognizer* for the language

$$L = \{ \mathbf{a}^n \mid n > 0 \text{ and the hailstone sequence terminates for } n \}.$$

- If the hailstone sequence terminates for n , then M accepts \mathbf{a}^n . If it doesn't, then M does not accept \mathbf{a}^n .
- We honestly don't know if M is a decider for this language.
 - If the hailstone sequence always terminates, then M always halts and is a decider for L , and L turns out to be just all strings $\mathbf{a}^n \mid n > 0$ (a Regular language!).
 - If the hailstone sequence doesn't always terminate, then M will loop on some inputs and isn't a decider for L , and L is some strict subset of $\mathbf{a}^n \mid n > 0$.

Two new language classes

Recognizers and Recognizability

- The class **RE** consists of all recognizable languages.
- Formally speaking:

$$\mathbf{RE} = \{ L \mid L \text{ is a language and there's a recognizer for } L \}$$

- You can think of **RE** as “all problems with yes/no answers where “yes” answers can be confirmed by a computer.”

Deciders and Decidability

- The class **R** consists of all decidable languages.
- Formally speaking:
$$\mathbf{R} = \{ L \mid L \text{ is a language and there exists a decider for } L \}$$
- You can think of **R** as “all problems with yes/no answers that can be fully solved by computers.”
–
- The class **R** contains all the regular languages, all the context-free languages, most of CS161, etc.
- This is a “strong” notion of solving a problem.

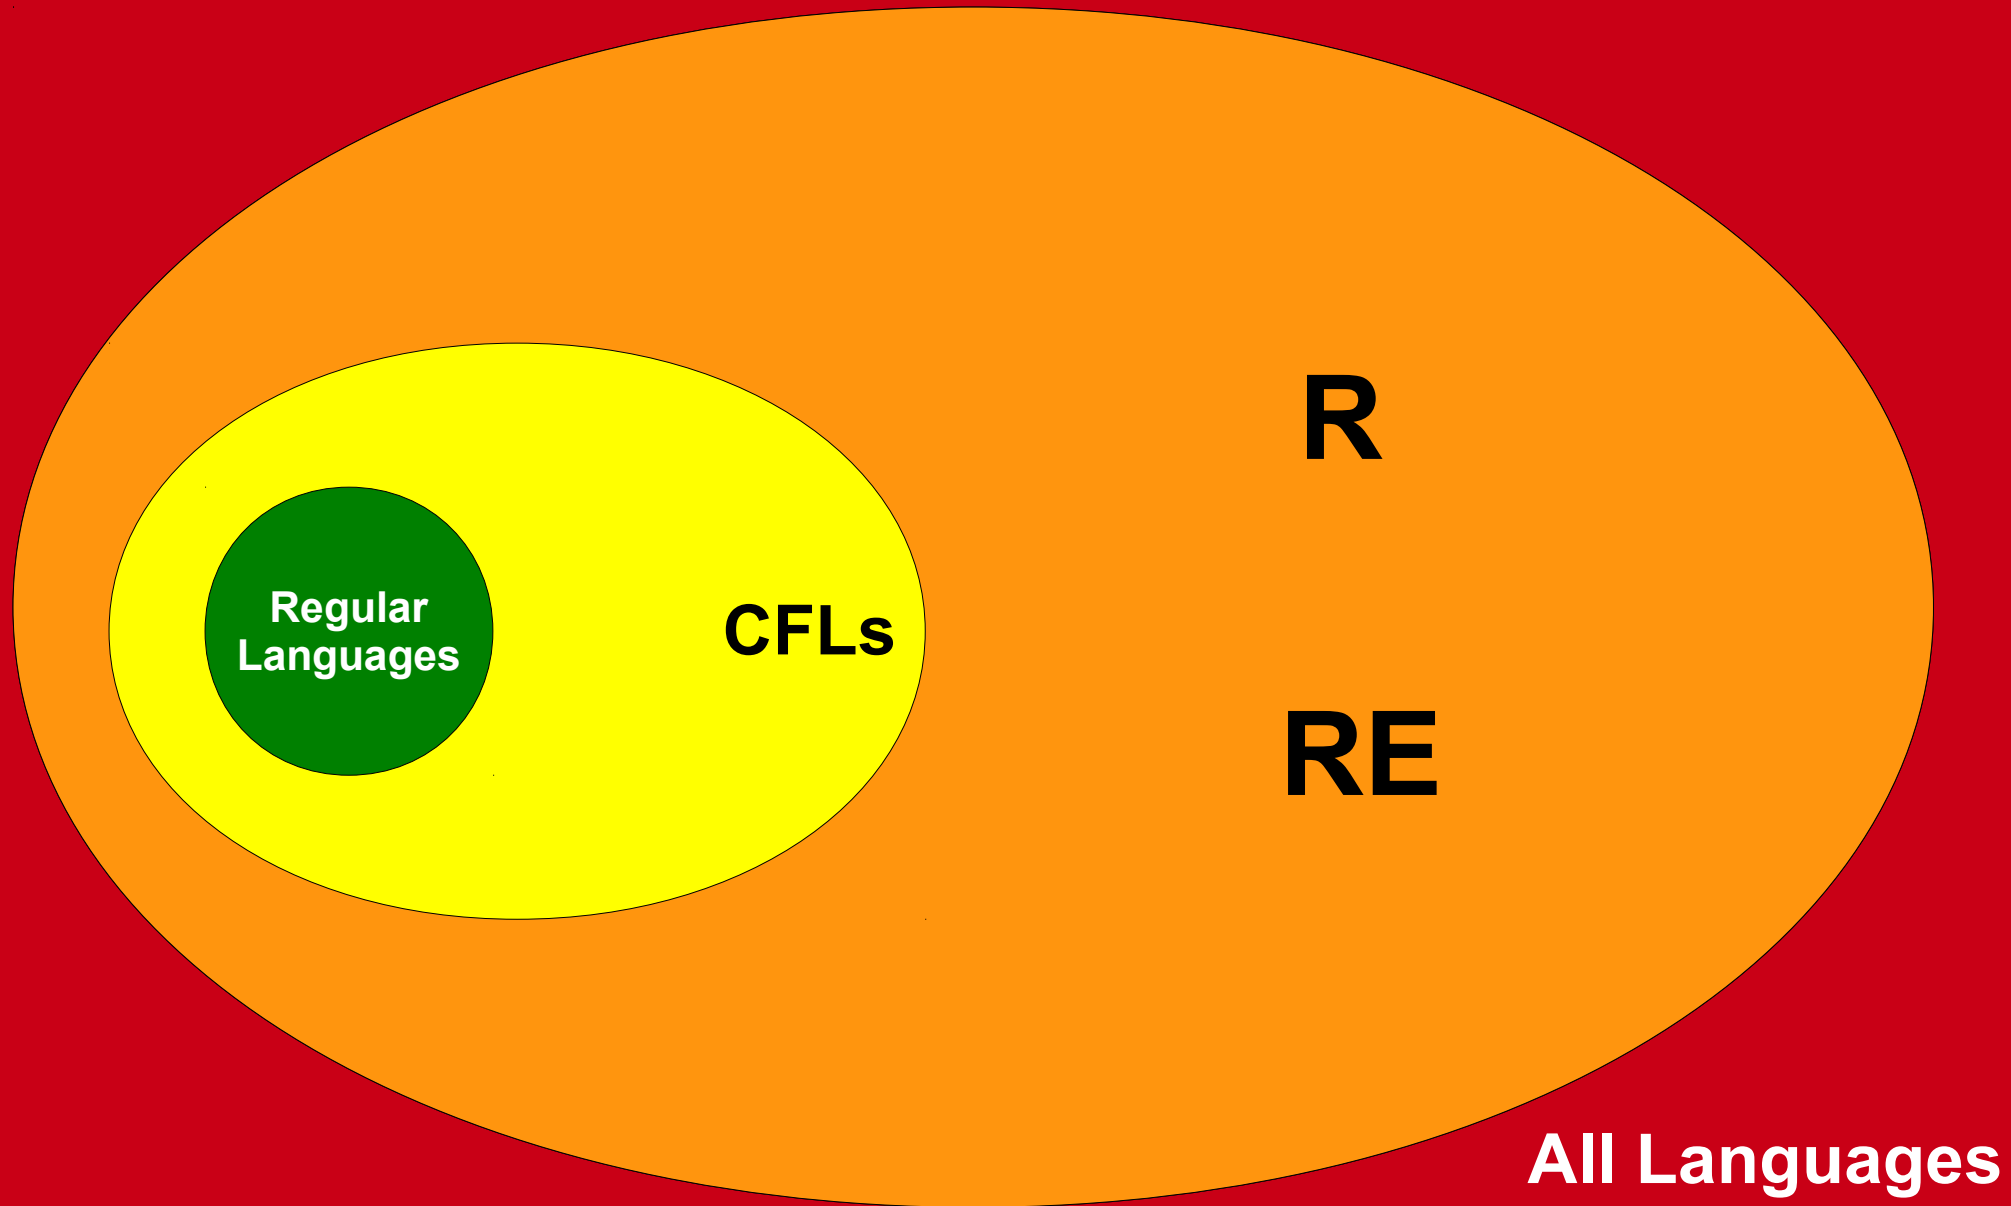
R and **RE** Languages

- Every decider for L is also a recognizer for L .
- This means that $\mathbf{R} \subseteq \mathbf{RE}$.
- Hugely important theoretical question:

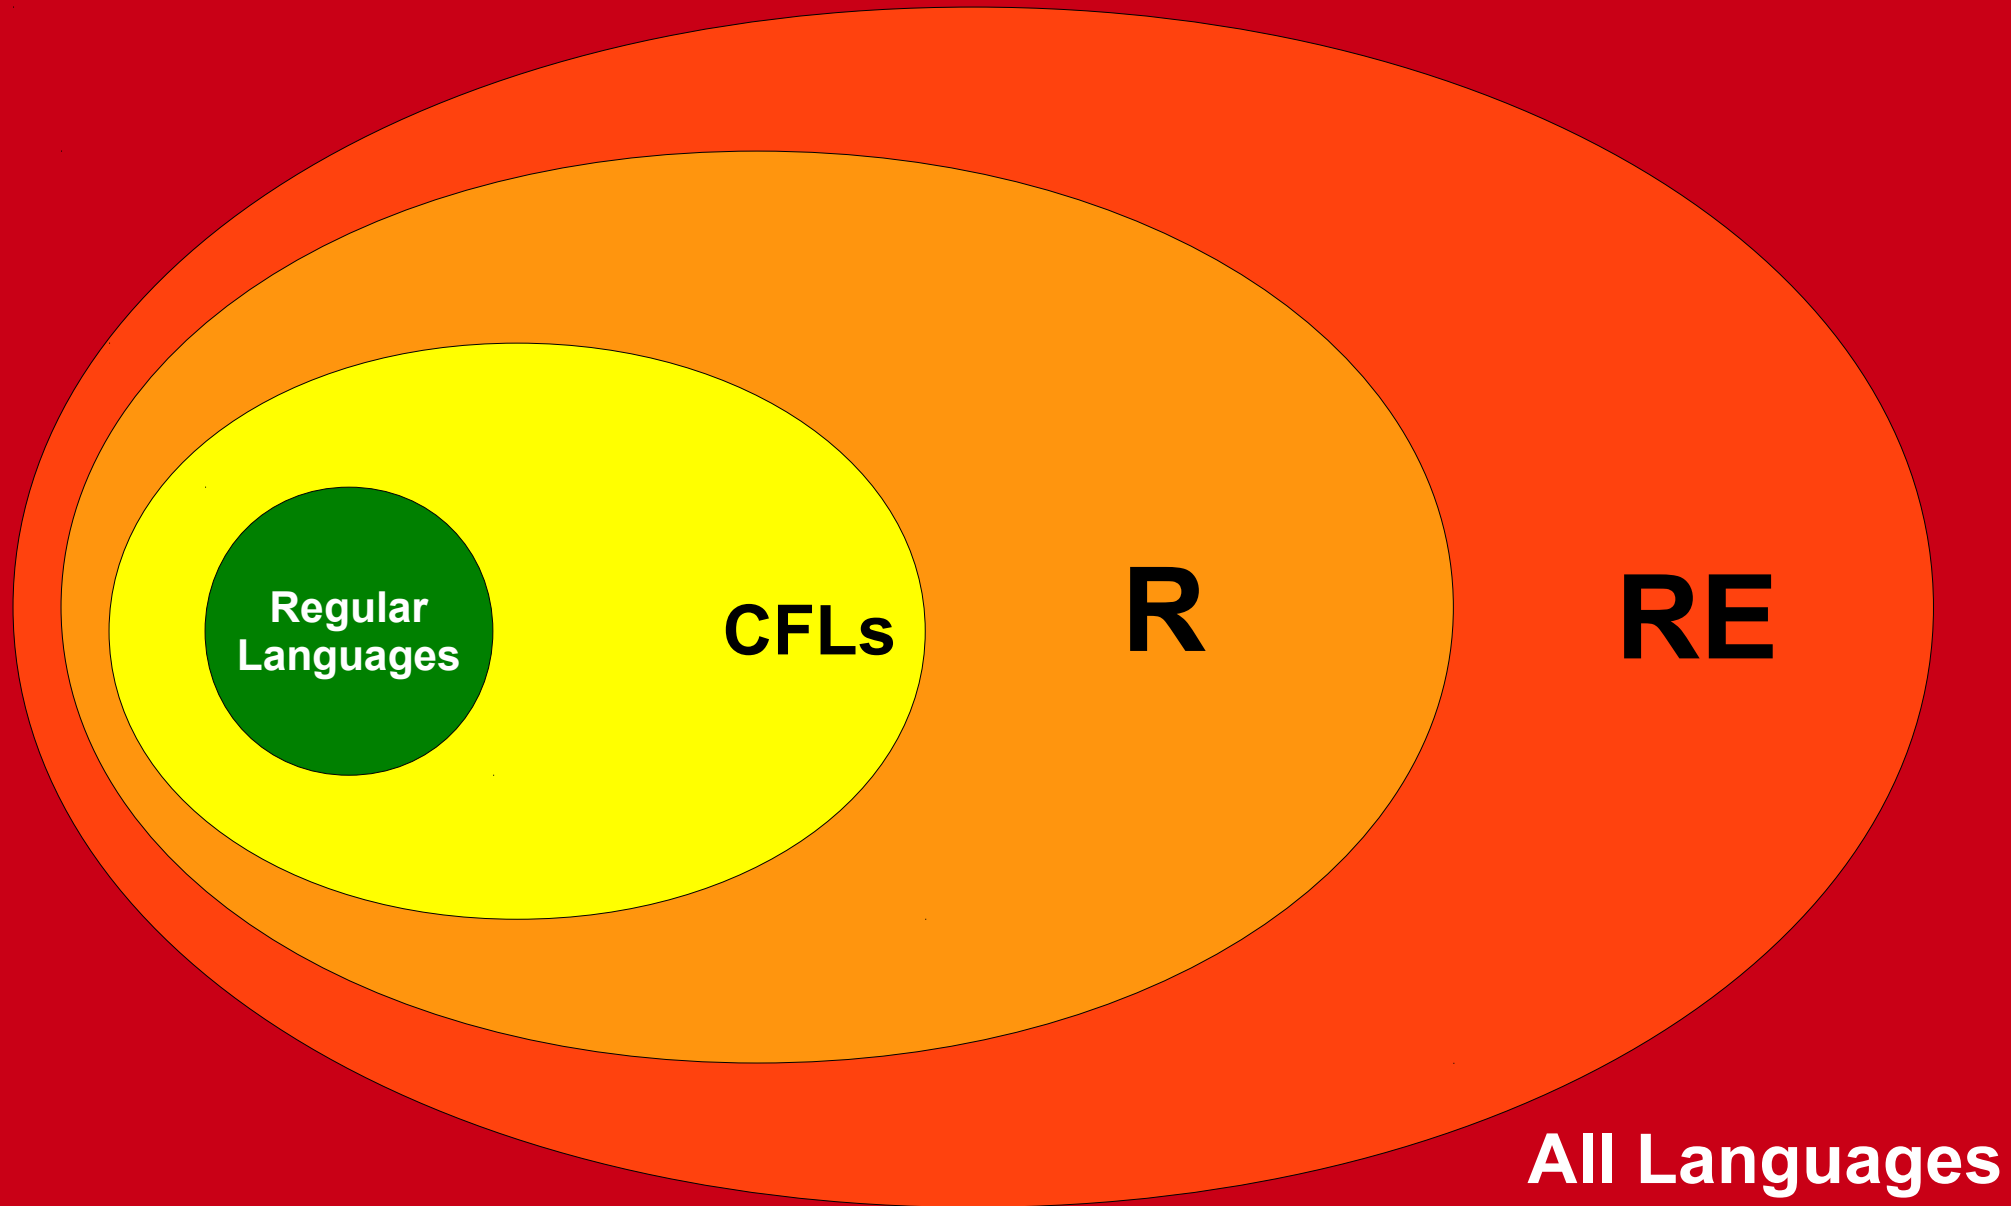
$$\mathbf{R} \stackrel{?}{=} \mathbf{RE}$$

- That is, if you can just confirm “yes” answers to a problem, can you necessarily *solve* that problem?

Which Picture is Correct?



Which Picture is Correct?



Unanswered Questions

- Why exactly is **RE** an interesting class of problems?
- What does the **R** $\stackrel{?}{=} \mathbf{RE}$ question mean?
- Is **R** = **RE**?
- What lies beyond **R** and **RE**?
- We'll see the answers to each of these in due time.

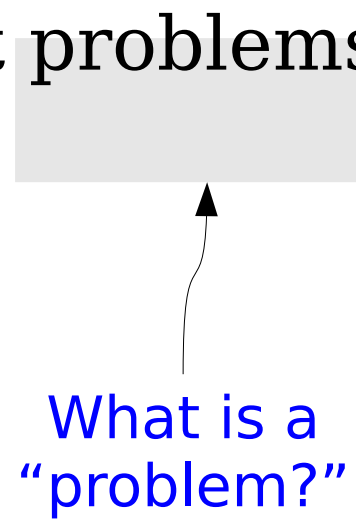
Next Time

- ***Emergent Properties***
 - Larger phenomena made of smaller parts.
- ***Universal Machines***
 - A single, “most powerful” computer.
- ***Self-Reference***
 - Programs that ask questions about themselves.

Strings, Languages, and Encodings

What problems can we solve with a computer?

What is a
“problem?”

A diagram consisting of a light gray rectangular box highlighting the word 'problems' in the main title. A thin black arrow points from the bottom center of this box down to the word 'problem' in the question 'What is a "problem?"' located below it.

Decision Problems

- A ***decision problem*** is a type of problem where the goal is to provide a yes or no answer.


- Example: Bin Packing

You're given a list of patients who need to be seen and how much time each one needs to be seen for. You're given a list of doctors and how much free time they have. Is there a way to schedule the patients so that they can all be seen?


- Example: Dominating Set Problem

You're given a transportation grid and a number k . Is there a way to place emergency supplies in at most k cities so that every city either has emergency supplies or is adjacent to a city that has emergency supplies?


A Model for Solving Problems




A Model for Solving Problems




A Model for Solving Problems




A Model for Solving Problems




A Model for Solving Problems



A Model for Solving Problems




A Model for Solving Problems




```
bool someFunctionName(string input) {  
    // ... do something ...  
}
```

A Model for Solving Problems




```
bool isAnBn(string input) {  
    // ... do something ...  
}
```

A Model for Solving Problems



```
bool isPalindrome(string input) {  
    // ... do something ...  
}
```


A Model for Solving Problems



```
bool isLinkageGraph(Graph G) {  
    // ... do something ...  
}
```

How does this
match our model?

A Model for Solving Problems



```
bool containsCat(Picture P) {  
    // ... do something ...  
}
```

How does this
match our model?

Humbling Thought:

*Everything on your computer is a
string over {0, 1}.*

Strings and Objects

- Think about how my computer encodes the image on the right.
- Internally, it's just a series of zeros and ones sitting on my hard drive.




Strings and Objects

- A different sequence of 0s and 1s gives rise to the image on the right.
- Every image can be encoded as a sequence of 0s and 1s, though not all sequences of 0s and 1s correspond to images.



Object Encodings

- If Obj is some mathematical object that is *discrete* and *finite*, then we'll use the notation $\langle Obj \rangle$ to refer to some way of encoding that object as a string.
- Think of $\langle Obj \rangle$ like a file on disk – it encodes some high-level object as a series of characters.



Object Encodings

- If Obj is some mathematical object that is *discrete* and *finite*, then we'll use the notation $\langle Obj \rangle$ to refer to some way of encoding that object as a string.
- Think of $\langle Obj \rangle$ like a file on disk – it encodes some high-level object as a series of characters.



$\langle Obj \rangle = 001101010001... \\ 001$


Object Encodings

- For the purposes of what we're going to be doing, we aren't going to worry about exactly *how* objects are encoded.
- For example, we can say $\langle G \rangle$ to mean “some encoding of a graph G ” without worrying about how it's encoded.
 - Analogy: do you need to know how numbers are represented in Python to be a Python programmer? That's more of a CS107 question.
- We'll assume, whenever we're dealing with encodings, that some person has figured out an encoding system for us and that we're using that encoding system.

Object Encodings

- **Great intuition:** If you can store an object as a file on disk, then you can encode it as a string.
- Here are a bunch of different types of objects. Which of these objects can *always* be encoded as a string?
 - A DFA over the alphabet $\{\mathbf{a}, \mathbf{b}\}$.
 - A regular expression.
 - A subset of $\{\mathbf{a}, \mathbf{b}\}^*$.
 - A function f from $\{k \in \mathbb{N} \mid k < n\}$ to itself, for some $n \in \mathbb{N}$.
 - A graph whose nodes are the set $\{k \in \mathbb{N} \mid k < n\}$, for some $n \in \mathbb{N}$.


A Model for Solving Problems



```
bool containsCat(Picture P) {  
    // ... do something ...  
}
```

Internally, this is
a sequence of
0s and **1**s.


A Model for Solving Problems



```
bool containsCat(Picture P) {  
    // ... do something ...  
}
```


Internally, this is
a sequence of
0s and **1**s.

A Model for Solving Problems




```
bool containsCat(Picture P) {  
    // ... do something ...  
}
```

A Model for Solving Problems



```
bool isLinkageGraph(Graph G) {  
    // ... do something ...  
}
```


A Model for Solving Problems



```
bool isDominatingSet(Graph G, Set D) {  
    // ... do something ...  
}
```

How does this
match our model?

A Model for Solving Problems




```
bool matchesRegex(string w, Regex R) {  
    // ... do something ...  
}
```

How does this
match our model?

Encoding Groups of Objects

- Given a group of objects $Obj_1, Obj_2, \dots, Obj_n$, we can create a single string encoding all these objects.
 - **Intuition 1:** Think of it like a .zip file, but without the compression.
 - **Intuition 2:** Think of it like a tuple or struct.
- We'll denote the encoding of all of these objects as a single string by **$\langle Obj_1, \dots, Obj_n \rangle$** .


A Model for Solving Problems



```
bool matchesRegex(string w, Regex R) {  
    // ... do something ...  
}
```

These form one large bitstring.

A Model for Solving Problems



```
bool matchesRegex(string w, Regex R) {  
    // ... do something ...  
}
```

These form one large bitstring.