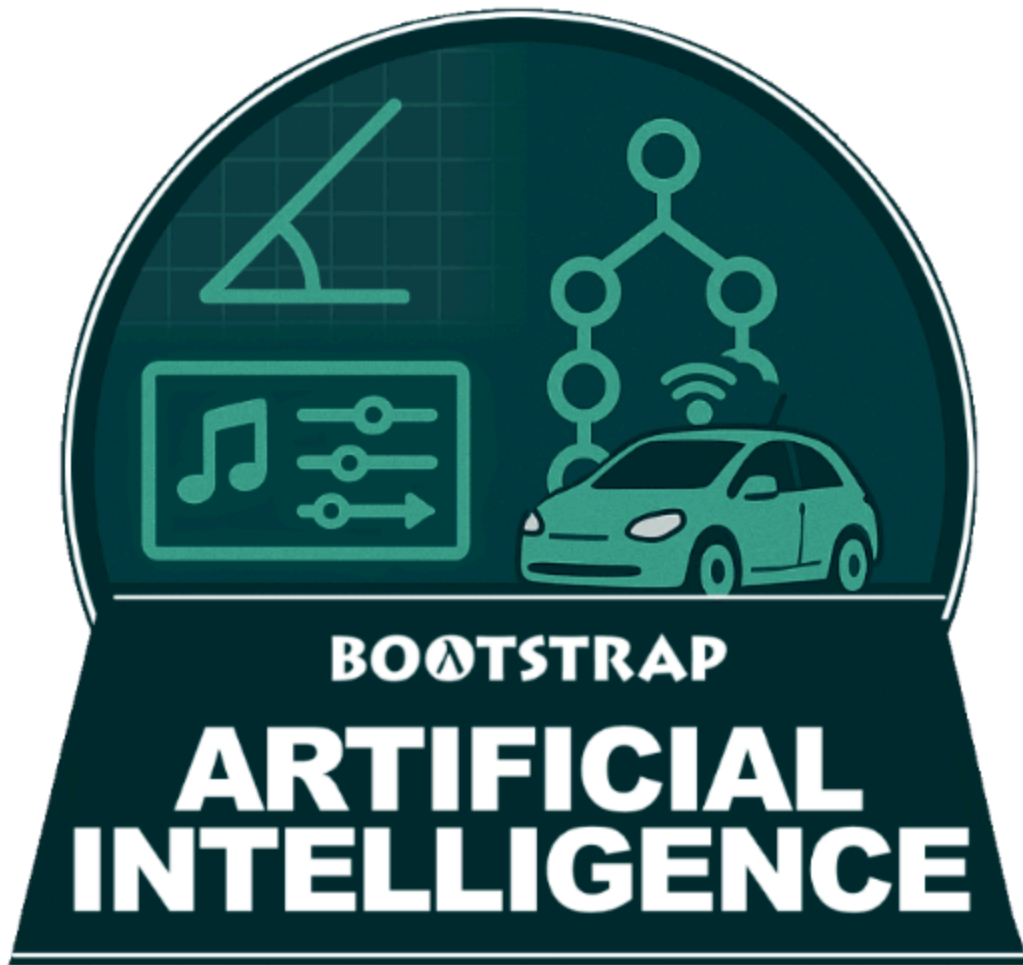


Name: _____



Student Workbook

Fall, 2025 — Pyret Edition



BOOTSTRAP

Equity • Scale • Rigor

Workbook v3.1

Brought to you by the Bootstrap team:

- Emmanuel Schanzer
- Kathi Fisler
- Shriram Krishnamurthi
- Dorai Sitaram
- Joe Politz
- Ben Lerner
- Nancy Pfenning
- Flannery Denny
- Rachel Tabak

Bootstrap is licensed under a Creative Commons 4.0 Unported License. Based on a work from www.BootstrapWorld.org.
Permissions beyond the scope of this license may be available at contact@BootstrapWorld.org.

Measuring Similarity

Previewing the Photos

1) What about the nine photos makes them all *similar* to one another?

2) What about the nine photos makes them all *different* from one another?

Sorting the Photos

3) "Alt text," short for alternative text, is a short description added to images on a webpage to provide context for users who cannot see the image. Write a brief description of Photo 1 that could serve as alt text, describing the image for someone who cannot see it.

4) Which is *more similar* to Photo 1: Photo 2 or Photo 3? Explain how you decided which photo is *more similar*.

5) Write a brief description of Photo 4 that could serve as alt text, providing as much detail as you can.

6) Which is *more similar* to Photo 4: Photo 5 or Photo 6? Explain how you decided which photo is *more similar*.

Using the Coordinate Plane

7) We just thought about where these images would land on a coordinate grid with its axes labeled "*body of water pixels*" and "*mountain pixels*". Think of other possible labels for the x and y axes. How many can you come up with?

Table of Contents

Introduction to Artificial Intelligence	1
Simple Data Types	3
Contracts for Strings and Images	6
Data-Driven Algorithms: Spell Checkers	15
Supervised Learning: Self-Driving Cars	19
Training Artificial Intelligence: Bags of Words	20
Training Artificial Intelligence: Language in Practice	26
Supervised Learning: Decision Trees	29
Statistical Language Modeling: Generating Text	34
Statistical Language Modeling: Chatbots	36

Table of Contents

Introduction to Artificial Intelligence	1
Simple Data Types	3
Contracts for Strings and Images	6
Data-Driven Algorithms: Spell Checkers	15
Supervised Learning: Self-Driving Cars	19
Training Artificial Intelligence: Bags of Words	20
Training Artificial Intelligence: Language in Practice	26
Supervised Learning: Decision Trees	29
Statistical Language Modeling: Generating Text	34
Statistical Language Modeling: Chatbots	36

Thinking about AI in Pop Culture

1) The media we chose: _____

2) The message about AI that we're going to discuss:

- ☐ AI must be safely controlled at all times.
- ☐ AI could lead to social isolation.
- ☐ AI can help us understand what it means to be human.
- ☐ AI can help advance humanity and scientific discovery.
- ☐ AI may pose a threat to humanity.
- ☐ AI can perpetuate biases and lead to unfair outcomes.
- ☐ Humans are helpless in the face of AI.

3) Our thoughts about how the media we chose supports the message we selected:

This image shows a single sheet of white paper with horizontal blue ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

Introduction to Programming in a Nutshell

The **Editor** is a software program we use to write Code. Our Editor allows us to experiment with Code on the right-hand side, in the **Interactions Area**. For Code that we want to *keep*, we can put it on the left-hand side in the **Definitions Area**. Clicking the "Run" button causes the computer to re-read everything in the Definitions Area and erase anything that was typed into the Interactions Area.

Data Types

Programming languages involve different *data types*, such as Numbers, Strings, Booleans, and even Images.

- Numbers are values like `1`, `0.4`, `1/3`, and `-8261.003`.
 - Numbers are *usually* used for quantitative data and other values are *usually* used as categorical data.
 - In Pyret, decimals *must* start with a zero. For example, `0.22` is valid, but `.22` is not.
- Strings are values like `"Emma"`, `"Rosanna"`, `"Jen and Ed"`, or even `"08/28/1980"`.
 - All strings *must* be surrounded by quotation marks.
- Booleans are either `true` or `false`.

All values evaluate to themselves. The program `42` will evaluate to `42`, the String `"Hello"` will evaluate to `"Hello"`, and the Boolean `false` will evaluate to `false`.

Operators

Operators (like `+`, `-`, `*`, `<`, etc.) work the same way in Pyret that they do in math.

- Operators are written between values, for example: `4 + 2`.
- In Pyret, operators must always have spaces around them. `4 + 2` is valid, but `4+2` is not.
- If an expression has different operators, parentheses must be used to show order of operations. `4 + 2 + 6` and `4 + (2 * 6)` are valid, but `4 + 2 * 6` is not.

Applying Functions

Functions work much the way they do in math. Every function has a name, takes some inputs, and produces some output. The function name is written first, followed by a list of *arguments* in parentheses.

- In math this could look like $f(5)$ or $g(10, 4)$.
- In Pyret, these examples would be written as `f(5)` and `g(10, 4)`.
- Applying a function to make images would look like `star(50, "solid", "red")`.
- There are many other functions in Pyret, for example `sqr`, `sqrt`, `triangle`, `square`, `string-repeat`, etc.

Functions have *contracts*, which help explain how a function should be used. Every Contract has three parts:

- The *Name* of the function — literally, what it's called.
- The *Domain* of the function — what *type(s) of value(s)* the function consumes, and in what order.
- The *Range* of the function — what *type of value* the function produces.

Strings and Numbers

Make sure you've loaded code.pyret.org (CPO), clicked "Run", and are working in the **Interactions Area** on the right. Hit Enter/return to evaluate expressions you test out.

Strings

String values are always in quotes.

- Try typing your name (in quotes!).
- Try typing a sentence like "I'm excited to learn to code!" (in quotes!).
- Try typing your name with the opening quote, but *without the closing quote*. Read the error message!
- Now try typing your name *without any quotes*. Read the error message!

1) Explain what you understand about how strings work in this programming language. _____

Numbers

2) Try typing 42 into the Interactions Area and hitting "Enter". Is 42 the same as "42"? Why or why not?

3) What is the largest number the editor can handle?

4) Try typing 0.5. Then try typing .5. Then try clicking on the answer. Experiment with other decimals.

Explain what you understand about how decimals work in this programming language. _____

5) What happens if you try a fraction like 1/3? _____

6) Try writing **negative** integers, fractions and decimals. What do you learn? _____

Operators

7) Just like math, Pyret has **operators** like +, -, * and /.

Try typing in 4 + 2 and then 4+2 (without the spaces). What can you conclude from this?

8) Type in the following expressions, **one at a time**: 4 + 2 * 6 (4 + 2) * 6 4 + (2 * 6) What do you notice?

9) Try typing in 4 + "cat", and then "dog" + "cat". What can you conclude from this?

Booleans

Boolean-producing expressions are yes-or-no questions, and will always evaluate to either **true** ("yes") or **false** ("no").

What will the expressions below evaluate to? Write down your prediction, then type the code into the Interactions Area to see what it returns.

Prediction	Result	Prediction	Result
1) <code>3 <= 4</code>		2) <code>"a" > "b"</code>	
3) <code>3 == 2</code>		4) <code>"a" < "b"</code>	
5) <code>2 < 4</code>		6) <code>"a" == "b"</code>	
7) <code>5 >= 5</code>		8) <code>"a" <> "a"</code>	
9) <code>4 >= 6</code>		10) <code>"a" >= "a"</code>	
11) <code>3 <> 3</code>		12) <code>"a" <> "b"</code>	
13) <code>4 <> 3</code>		14) <code>"a" >= "b"</code>	

15) In your own words, describe what `<` does.

16) In your own words, describe what `>=` does.

17) In your own words, describe what `<>` does.

Prediction:	Result:
18) <code>string-contains("catnap", "cat")</code>	
19) <code>string-contains("cat", "catnap")</code>	

20) In your own words, describe what `string-contains` does. Can you generate another expression using `string-contains` that returns true?

★ There are infinite string values ("a", "aa", "aaa" ...) and infinite number values out there (...-2,-1,0,-1,2...). But how many different *Boolean* values are there?

Applying Functions

Open [\(code.pyret.org \(CPO\)\)](https://code.pyret.org) and click "Run". We will be working in the Interactions Area on the right.

Test out these two expressions and record what you learn below:

- `regular-polygon(40, 6, "solid", "green")`
- `regular-polygon(80, 5, "outline", "dark-green")`

1) You've seen data types like Numbers, Strings, and Booleans. What data type did the `regular-polygon` function produce? _____

2) How would you describe what a regular polygon is? _____

3) The `regular-polygon` function takes in four pieces of information (called arguments). Record what you know about them below.

	Data Type	Information it Contains
Argument 1	Number	Length of each side
Argument 2	Number	
Argument 3		
Argument 4		

There are many other functions available to us in Pyret. We can describe them using **contracts**. The Contract for `regular-polygon` is:

```
# regular-polygon :: Number, Number, String, String -> Image
```

- Each Contract begins with the function name: in this case regular-polygon
- Lists the data types required to satisfy its Domain: in this case Number, Number, String, String
- And then declares the data type of the Range it will return: in this case Image

Contracts can also be written with more detail, by annotating the Domain with *variable names*:

```
# regular-polygon :: ( Number , Number , String , String ) -> Image
                     size   number-of-sides fill-style color
```

4) We know that a square is a regular polygon because _____

5) What code would you write to make a big, blue square using the `regular-polygon` function?

```
_____ ( _____ , _____ , _____ , _____ )
      function-name      size :: Number      number-of-sides :: Number      fill-style :: String      color :: String
```

6) Pyret also has a `square` function whose contract is: `# square :: (Number , String , String) -> Image`

What code would you write to make a big blue square using the `square` function?

```
_____ ( _____ , _____ , _____ )
      function-name      size :: Number      fill-style :: String      color :: String
```

7) Why does `square` need fewer arguments to make a square than `regular-polygon`? _____

★ Where else have you heard the word **contract** used before?

Practicing Contracts: Domain & Range

Note: The contracts on this page are not defined in Pyret and cannot be tested in the editor.

is-beach-weather

Consider the following Contract:

```
# is-beach-weather :: Number, String -> Boolean
```

- 1) What is the **Name** of this function? _____
- 2) How many arguments are in this function's **Domain**? _____
- 3) What is the **Type** of this function's **first argument**? _____
- 4) What is the **Type** of this function's **second argument**? _____
- 5) What is the **Range** of this function? _____

6) Circle the expression below that shows the correct application of this function, based on its Contract.

- A. is-beach-weather(70, 90)
- B. is-beach-weather(80, 100, "cloudy")
- C. is-beach-weather("sunny", 90)
- D. is-beach-weather(90, "stormy weather")

cylinder

Consider the following Contract:

```
# cylinder :: Number, Number, String -> Image
```

- 7) What is the **Name** of this function? _____
- 8) How many arguments are in this function's **Domain**? _____
- 9) What is the **Type** of this function's **first argument**? _____
- 10) What is the **Type** of this function's **second argument**? _____
- 11) What is the **Type** of this function's **third argument**? _____
- 12) What is the **Range** of this function? _____

13) Circle the expression below that shows the correct application of this function, based on its Contract.

- A. cylinder("red", 10, 60)
- B. cylinder(30, "green")
- C. cylinder(10, 25, "blue")
- D. cylinder(14, "orange", 25)

Matching Expressions and Contracts

Match the Contract (left) with the expression that uses it correctly (right).

Note: The contracts on this page are not defined in Pyret and cannot be tested in the editor.

Contract	Expression
<code># make-id :: String, Number -> Image</code> 1	A <code>make-id("Savannah", "Lopez", 32)</code>
<code># make-id :: String, Number, String -> Image</code> 2	B <code>make-id("Pilar", 17)</code>
<code># make-id :: String -> Image</code> 3	C <code>make-id("Akemi", 39, "red")</code>
<code># make-id :: String, String -> Image</code> 4	D <code>make-id("Raissa", "McCracken")</code>
<code># make-id :: String, String, Number -> Image</code> 5	E <code>make-id("von Einsiedel")</code>

Contract	Expression
<code># is-capital :: String, String -> Boolean</code> 6	A <code>show-pop("Juneau", "AK", 31848)</code>
<code># is-capital :: String, String, String -> Boolean</code> 7	B <code>show-pop("San Juan", 395426)</code>
<code># show-pop :: String, Number -> Image</code> 8	C <code>is-capital("Accra", "Ghana")</code>
<code># show-pop :: String, String, Number -> Image</code> 9	D <code>show-pop(3751351, "Oklahoma")</code>
<code># show-pop :: Number, String -> Number</code> 10	E <code>is-capital("Albany", "NY", "USA")</code>

Contracts for Image-Producing Functions

Log into code.pyret.org (CPO) and click "Run". Experiment with each of the functions listed below in the interactions area. Try to find an expression that produces an image. Record the contract and example code for each function you are able to use!

Name	Domain	Range
# triangle	:: Number, String, String	-> Image
triangle(80, "solid", "darkgreen")		
# star	::	->
# circle	::	->
# rectangle	::	->
# text	::	->
# square	::	->
# rhombus	::	->
# ellipse	::	->
# regular-polygon	::	->
# right-triangle	::	->
# isosceles-triangle	::	->
# radial-star	::	->
# star-polygon	::	->
# triangle-sas	::	->
# triangle-asa	::	->

Catching Bugs when Making Triangles

Learning about a Function through Error Messages

- 1) Type `triangle` into the Interactions Area of [code.pyret.org\(CPO\)](http://code.pyret.org(CPO)) and hit "Enter". What do you learn? _____
- 2) We know that all functions will need an open parenthesis and at least one input! Type `triangle(80)` in the Interactions Area and hit Enter/return. Read the error message. What hint does it give us about how to use this function?

- 3) Using the hint from the error message, experiment until you can make a triangle. What is the contract for `triangle`?

- 4) Read the explanation below. Then explain the difference in your own words.
syntax errors — when the computer cannot make sense of the code because of unclosed strings, missing commas or parentheses, etc.
contract errors — when the function isn't given what it needs (the wrong type or number of arguments are used)

The difference between **syntax errors** and **contract errors** is: _____

Finding Mistakes with Error Messages

The following lines of code are all BUGGY! Read the code and the error messages below. See if you can find the mistake WITHOUT typing it into Pyret.

- 5) `triangle(20, "solid" "red")`
Pyret didn't understand your program around
`triangle(20, "solid" "red")`

This is a _____ error. The problem is that _____
contract/syntax
- 6) `triangle(20, "solid")`
This application expression errored:
`triangle(20, "solid")`
2 arguments were passed to the **operator**. The **operator** evaluated to a function accepting 3 parameters. An application expression expects the number of parameters and arguments to be the same.

This is a _____ error. The problem is that _____
contract/syntax
- 7) `triangle(20, 10, "solid", "red")`
This application expression errored:
`triangle(20, 10, "solid", "red")`
4 arguments were passed to the **operator**. The **operator** evaluated to a function accepting 3 parameters. An application expression expects the number of parameters and arguments to be the same.



This is a _____ error. The problem is that _____
contract/syntax
- 8) `triangle (20, "solid", "red")`
Pyret thinks this code is probably a function call:
`triangle (20, "solid", "red")`
Function calls must not have space between the **function expression** and the arguments.

This is a _____ error. The problem is that _____
contract/syntax


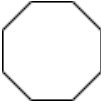
Using Contracts

For questions 1,2,4,5,8 & 9, use the contracts provided to find expressions that will generate images similar to the ones pictured.
Test your code in [code.pyret.org\(CPO\)](http://code.pyret.org(CPO)) before recording it.



```
# ellipse :: ( Numberwidth , Numberheight , Stringfill-style , Stringcolor ) -> Image
```

1)		
2)		
3)	Write an expression using <code>ellipse</code> to produce a circle.	

```
# regular-polygon :: ( Numberside-length , Numbernumber-of-sides , Stringfill-style , Stringcolor ) -> Image
```

4)		
5)		
6)	Use <code>regular-polygon</code> to write an expression for a square!	
7)	How would you describe a regular polygon to a friend?	

```
# rhombus :: ( Numbersize , Numbertop-angle , Stringfill-style , Stringcolor ) -> Image
```

8)		
9)		
10)	Write an expression to generate a <code>rhombus</code> that is a square!	

Triangle Contracts

Respond to the questions. Go to [code.pyret.org\(CPO\)](http://code.pyret.org(CPO)) to test your code.

1) What kind of triangle does the `triangle` function produce? _____
There are lots of other kinds of triangles! And Pyret has lots of other functions that make triangles!

```
# triangle :: (Number, String, String) -> Image
               size      fill-style  color
# right-triangle :: (Number, Number, String, String) -> Image
                    base    height  fill-style  color
# isosceles-triangle :: (Number, Number, String, String) -> Image
                       leg    angle  fill-style  color
```

2) Why do you think `triangle` only needs one number, while `right-triangle` and `isosceles-triangle` need two numbers?

3) Write `right-triangle` expressions for the images below using `100` as one argument for each.





4) Write `isosceles-triangle` expressions for the images below using `100` as one argument for each.





5) Write 2 expressions that would build **right-isosceles** triangles. Use `right-triangle` for one expression and `isosceles-triangle` for the other expression.



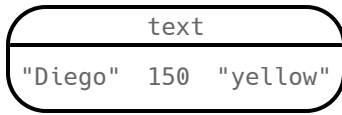
6) Which do you like better? Why? _____

Composing with Circles of Evaluation

Notice and Wonder

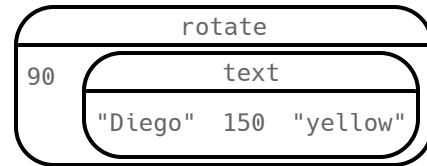
Suppose we want to see the `text` "Diego" written vertically in yellow letters of size 150. Let's use Circles of Evaluation to look at the structure:

We can start by generating the Diego image.



```
text("Diego", 150, "yellow")
```

And then use the `rotate` function to rotate it 90 degrees.



```
rotate(90, text("Diego", 150, "yellow"))
```

1) What do you Notice? _____

2) What do you Wonder? _____

Let's Rotate an Image of Your Name!

Suppose you wanted the computer to show your name in your favorite color and rotate it so that it's diagonal...

Write your name (any size), in your favorite color

3) Draw the circle of evaluation:

`rotate` the image so that it's diagonal

4) Draw the circle of evaluation:

5) Convert the Circle of Evaluation to code:

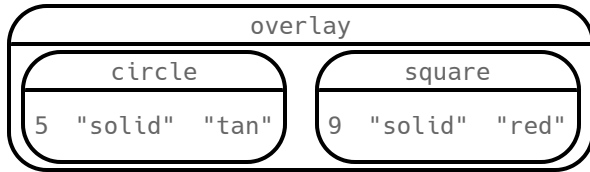
6) Convert the Circle of Evaluation to code:

Circle of Evaluation to Code (Scaffolded)

Complete the Code by Filling in the Blanks!

Finish the Code by filling in the blanks.

1)

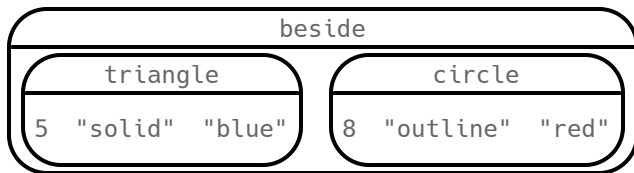


overlay(circle(____, "solid", _____), _____(9, _____, "red"))

Complete the Code by adding Parentheses

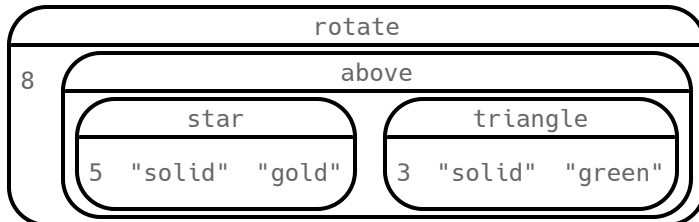
For each Circle of Evaluation, finish the Code by adding parentheses and commas.

2)



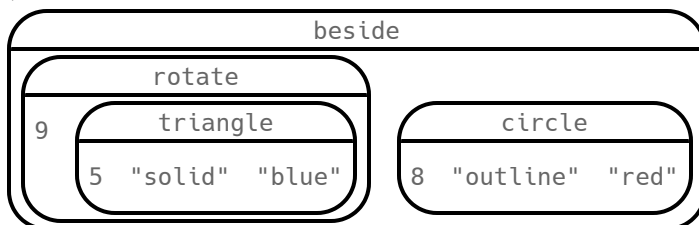
beside triangle 5 "solid" "blue" circle 8 "outline" "red"

3)



rotate 8 above star 5 "solid" "gold" triangle 3 "solid" "green"

4)



beside rotate 9 triangle 5 "solid" "blue" circle 8 "outline" "red"

Frayer Model: Domain and Range

My Definition		Facts and Characteristics
Domain		
Examples		Non-Examples
Range		
My Definition		Facts and Characteristics
Examples		Non-Examples

(optional)

Frayer Model: Function and Variable

My Definition		Facts and Characteristics
Function		
Examples		Non-Examples
Variable		
My Definition		Facts and Characteristics
Examples		Non-Examples

Radial Star

radial-star :: (Number
points, Number
outer-radius, Number
inner-radius, String
fill-style, String
color) -> Image

Using the Contract above, match the images on the left to the expressions on the right. Test the code at [code.pyret.org\(CPO\)](https://code.pyret.org/CPO).



1

A

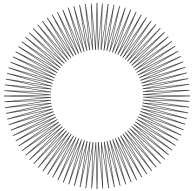
```
radial-star(5, 200, 50, "solid", "black")
```



2

B

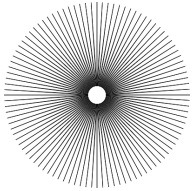
```
radial-star(7, 200, 100, "solid", "black")
```



3

C

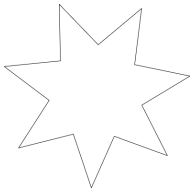
```
radial-star(7, 200, 100, "outline", "black")
```



4

D

```
radial-star(10, 200, 150, "solid", "black")
```



5

E

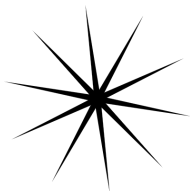
```
radial-star(10, 200, 20, "solid", "black")
```



6

F

```
radial-star(100, 200, 20, "outline", "black")
```



7





G

```
radial-star(100, 200, 100, "outline", "black")
```

(optional)

Triangle Contracts (SAS & ASA)

Type each expression (left) below into the [code.pyret.org \(CPO\)](https://code.pyret.org/CPO/), and match it to the image it creates (right).

Expression			Image
<code>triangle-sas(120, 45, 70, "solid", "black")</code>	1	A	
<code>triangle-sas(120, 90, 70, "solid", "black")</code>	2	B	
<code>triangle-sas(120, 135, 70, "solid", "black")</code>	3	C	
<code>triangle-sas(70, 135, 120, "solid", "black")</code>	4	D	

Contracts

Think about how you would describe each `triangle-sas` argument to someone who'd never used the function before.

5) Annotate the Contract below using descriptive variable names.

```
triangle-sas :: ( Number, Number, Number, String, String ) -> Image
```

If you have a printed workbook, add examples of each of the triangle functions we've explored to your contracts pages.

★ If you have time, experiment with the `triangle-asa` function.

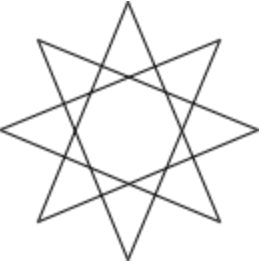
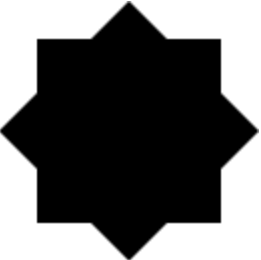
```
# triangle-asa :: ( Numbertop-left-angle, Numberleft-side, Numberbottom-angle, Stringfill-style, Stringcolor ) -> Image
```

★ Why did these two functions need to take in one more Number than `right-triangle` did?

Star Polygon

star-polygon :: (Number
side-length , Number
points-on-polygon , Number
points-to-skip-for-star , String
fill-style , String
color) -> Image

1. Using the Contract above, write expressions to create images like those pictured below.
2. Go to [\(CPO\)](http://code.pyret.org) to test your code.
3. Then write expressions to generate two more star polygons of your choosing.
Sketch them and record your working code.

1		<hr/>
2		<hr/>
3		<hr/>
4		<hr/>

(optional)

Human Spell Checking

"Before beginning the lesson, the teacher said wello to her students and asked them how thay were doing."

First Misspelled Word

1) What is the **first word that is spelled incorrectly** in the quote at the top? _____ What was the *intended* word? _____

How did you know that that word was misspelled? Explain in as much detail as you can. _____

2) Do you think a 6-year old would be able to identify and correct the first misspelled word in this sentence? Explain. _____

Second misspelled word

3) What is the **second word that is spelled incorrectly** in the quote at the top? _____ What was the *intended* word? _____

How did you know that that word was misspelled? Explain in as much detail as you can. _____

4) Do you think a 6-year old would be able to identify and correct the second misspelled word in this sentence? Explain. _____

Reflect

5) You probably didn't need a dictionary to correct the two errors above. When might consulting a dictionary be a part of your spell-checking process? Explain. _____

The First Spell Checker

The very first spell checker would use misspelled words as inputs.

- **First**, it would develop alternative candidates for the misspelled word (input) by making one of the following adjustments: (1) replace a single letter with a different one, or (2) swap the positions of two adjacent letters.
- **Next**, it would search the provided dictionary to see which alternatives were valid options.
- **Finally**, it would produce a list of valid correctly spelled words for the user to choose from.

Part 1: Follow the Algorithm

Misspelled Word: wello

1) Generate alternative candidates for the misspelled word. *(These candidates don't need to be real words!)*

- First, replace **one** letter from "wello" with a different one. *We've provided a first option to get you started.*

_____ dello _____ _____ _____ _____

- Next, swap the positions of two **adjacent** letters. *Again, we've provided a first option to get you started.*

_____ ewllo _____ _____ _____ _____

2) Circle the *actual* words you listed above. How many real-word suggestions did you come up with? _____

3) You already found 5 "words" by replacing a single letter of **wello**. Using that strategy, about how many more could you list?

about 10 more

about 50 more

about 100 more

about 1000 more

Part 2: Reflect

4) What do you think are some limitations of this algorithm? _____

5) Compared to your own strategy for spell-checking, how similar / different is the first spell checker's algorithm?

A Pyret Spell Checker: The Algorithm

Open the [Spell Checker Starter File](#) and click "Run". Follow each of the instructions below to discover how our very own Pyret spell checker works.

1) In the [Spell Checker Starter File](#), the function `alt-words` consumes a string and a dictionary... but what does it produce?

- Test `alt-words("straw", WORDS)` in the Interactions Area and record the results in the left column of the table below.
- In the second column, describe how each option differs from the input string ("straw").
- In the third column, record how many edits to the string "straw" this algorithm makes. (A single edit could entail replacing one letter with a different one, for instance.)

Suggestion	How it's different from the string "straw"	Number of Edits

2) Test `alt-words("planet", WORDS)` in the Interactions Area.

- Notice that "planet" has six letters, but all of the suggested words have five letters. What else can you conclude about the spell checker's algorithm? _____

- To get from "planet" to "plait", delete _____ and change _____ to _____.
- To get from "planet" to "plank", delete _____ and change _____ to _____.
- The spell checker makes just *one* edit when it proposes _____ as a match for "planet".

Experiment with the `alt-words` function to answer the questions below, testing what it produces using Strings from your imagination.

3) We have seen multiple two-edit changes. Does the program propose words that include three edits? _____

4) What is the maximum length string that this spell checker will accept? _____

- Why does that maximum length make sense with what we know about the algorithm? _____

5) What else do you notice? _____

6) The Pyret spell check program does not use the same algorithm as the [The First Spell Checker](#). How are the programs different? _____

7) **Edit distance** is a word programmers use to describe the number of operations needed to transform one string into another string. How can **edit distance** help us measure similarity? _____

A Pyret Spell Checker: Exploring Different Dataset Sizes

The player of the partially-completed Wordle game wants some Pyret "assistance". Open the [Spell Checker Starter File](#) and click "Run".

1) In the Interactions Area, test each of the three lines of code (below).

Pyret Expression	Number of suggestions returned
<code>alt-words("games", WORDS-100)</code>	
<code>alt-words("games", WORDS-1000)</code>	
<code>alt-words("games", WORDS-ALL)</code>	

W	O	R	T	H
M	E	D	I	A
G	A	M	E	S

2) Which line of code would you recommend for the player? Why? _____

3) The solution to the game above is "camel". Look back at the suggestions in your Interactions Area. Did any of the lists include it? _____

4) We used the very same function (`alt-words`) for each entry. Why didn't we receive the same results from all 3 expressions? _____

5) To see the words in the `WORDS-100` dictionary, type `WORDS-100` into the Interactions Area and then click on `tree-set`. How many words are in this list? _____

6) What do you notice about the words in the list? _____

7) What do you wonder? _____

8) As you might have guessed, the `WORDS-1000` dictionary contains 1000 words and the `WORDS-ALL` dictionary contains even more (2314). How does that help explain the fact that `alt-words("games", WORDS-ALL)` returned the most suggestions? _____

★ Let's imagine that

- `this-word` and `that-word` represent two different five-letter words
- `alt-words(this-word, WORDS-ALL)` produces a list that includes `that-word`

Would you expect `alt-words(that-word, WORDS-ALL)` to produce a list that includes `this-word`? Why or why not? _____

★ Can you produce an input that returns more than one word from the `WORDS-100` dictionary? If so, what is it? _____

Supervised Machine Learning: Training a Self-Driving Car

Thinking about Training

1) Refer to the weather forecast (right) in your response. During a week of daytime test drives:

- On Wednesday and Friday, the self-driving car drives safely.
- On the other days, the self-driving car is unsafe.

What might explain why the car drove unsafely on some days?

SAT	SUN	MON	TUE	WED	THU	FRI
						
THUNDER, HEAVY RAIN	SUNNY, LIGHT SNOW	CLOUDY, LIGHT RAIN	HEAVY CLOUD	FEW CLOUDS	CLOUDY	SUNNY, WARMER
63°F	65°F	64°F	69°F	71°F	72°F	70°F
41°F	42°F	43°F	40°F	42°F	39°F	65°F

2) Imagine that a self-driving car has done extensive training on a one-lane road, in all weather conditions, and at all times of day. Would you expect it to be able to safely drive on a busy two-lane road? Explain.

3) Imagine that a self-driving car trained on isolated country roads, as well as city streets, and highways in all weather conditions and times of day. It would be pretty safe on many roads! What might it still not be prepared to navigate safely?

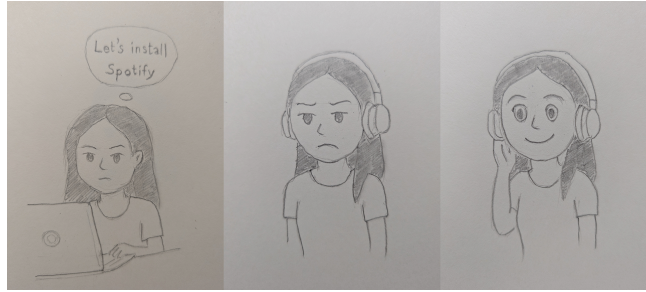
Confidence Rating

4) In addition to producing a steering angle for each image of the road, ALVINN produces a *numeric* measure of "confidence" in its response.

- What do you think causes ALVINN's "confidence" to increase or decrease?

- Is 100% "confidence" possible?

Case Study: Michelle's Spotify Use



Michelle started using Spotify six months ago. The first time she checked out her AI-produced "Discover Weekly" playlist, she was underwhelmed. The songs did not match her tastes, and she ended up skipping over many of them within a few seconds. Michelle was so disappointed in that first experience that she gave up entirely on the "Discover Weekly" playlist, writing it off as a feature she did not care for.

After three months of using Spotify and listening to all of her favorite songs, artists, and playlists, Michelle happened to click on "Discover Weekly" again — but this time, she was delighted by what she heard. It felt as though her best friend had curated a playlist just for her. Michelle discovered new music that she had never heard before, that she absolutely loved.

She remarked, "**Wow, Spotify has really improved their algorithms since I last checked out Discover Weekly.**"

1) Circle your assessment of Michelle's statement: *definitely true* *possibly true* *definitely not true*

2) What do you think might explain why her playlist is much better after three months of using Spotify? Provide as much detail as you can.
Hint: There may be more than one possible explanation!

3) Michelle loves folk music featuring female vocals, but she has decided that she wants to branch out and is starting to feel frustrated by how Spotify keeps recommending folk music to her over and over and over! Why isn't Spotify giving Michelle the music suggestions she's looking for?

4) What do you think Michelle should do to start getting different recommendations?

Designing a Song Recommendation System

1) Think of a song that you know very well. _____
Describe it in as much detail as you can.

- Begin with facts like the topic, musical genre, artist, year of release, language the lyrics are sung in, tempo, instruments used, etc.
- Then describe it more personally... for example, how it makes you feel, when you choose to listen / avoid listening to it, etc.
- Have fun!

[illegible]

2) A song recommendation system must collect a user's listening history and extract data that is useful for making predictions about the user's music preferences.

- If you were designing a song recommendation system, what specific data would it store? Make a list of ideas.
- Then rank each item based on how heavily you think it should be weighted in your song prediction function.

A Primitive Plagiarism Detector

Review the Documents

Wikipedia article on "Elephants in Thailand"	Student-submitted essay on "Elephants in Thailand"
The elephant has been a contributor to Thai society and its icon for many centuries. The elephant has had a considerable impact on Thai culture. The Thai elephant is the official national animal of Thailand. The elephant found in Thailand is the Indian elephant, a subspecies of the Asian elephant.	The elephant is a contributor to Thai society. It has been an icon of Thai life for many centuries. The elephant, which it is possible to see found in every part of Thailand, is the Indian elephant, which is a subspecies of the Asian elephant. The Thai elephant has a considerable impact on culture. The elephant is the official national animal of Thailand.

1) Do **you** think the student who submitted this essay on Elephants in Thailand committed plagiarism? Circle one: Yes No

Detect Identicality

Open the [Plagiarism Detection Starter File](#) and click "Run".

- The two paragraphs above are defined in the starter file as `wiki-article` and `student-essay`, respectively.
- The file also defines a function `simple-equality`, which takes in two Strings and produces a Boolean: `true` indicates that there is a match; `false` indicates that there is not a match.

2) Record if the output is `true` or `false` for each line of code:

`simple-equality("hello", "goodbye")` _____ `simple-equality("hello", "HELLO")` _____
`simple-equality("hello", "helo")` _____ `simple-equality("hello", " Hello.")` _____

3) How does `simple-equality` deal with punctuation, capitalization, and spaces? _____

4) Compare the essays using `simple-equality(student-essay, wiki-article)`. What does the program return? _____

Building a Better Plagiarism Detector

Yara and Xola agree that there has to be a way to measure the *similarity* of the essays. They each propose a different solution.

Yara: I made a list of the four most unusual and unique words in each essay. Then I compared the lists. Because the two lists of unusual words are identical (elephant, Thai, Indian, and Asian), I decided that the student plagiarized!

Xola: I figured out the four most common words in each essay. In the wiki-article, they are: "the" (7), "elephant" (6), "Thai" (3), and "a" (3). In the student essay, they are: "the" (7), "elephant" (6), "is" (5), and "of" (4). Because the most common words are not the same, I don't think the student plagiarized.

5) What do you like about Yara's method? What are its shortcomings?

6) What do you like about Xola's method? What are its shortcomings?

7) Come up with your own method! What is a **different** way to measure the similarity between the two essays?

A Slightly Less Primitive Plagiarism Detector

Open [Plagiarism Detection Starter File](#). Click "Run".

Reviewing `simple-similarity`

- 1) Type `doc-a` into the Interactions Area, then hit "Enter". What appears? _____
- 2) Type `doc-b` into the Interactions Area, then hit "Enter". What appears? _____
- 3) If we compare these two documents using our `simple-equality` function, what do you predict the outcome will be... and why?

- 4) Type `simple-equality(doc-a, doc-b)`. What is the output? _____

Testing `string-to-bag` and `bag-equality`

- 5) This starter file uses the function `string-to-bag`, which converts a string into a bag of words. Type `string-to-bag(doc-a)` in the Interactions Area and hit "Enter". What is the output? _____

- 6) Type `string-to-bag(doc-b)` in the Interactions Area and hit "Enter". What do you observe? _____

- 7) The function `bag-equality` checks if two different bags of words are identical. What do you predict the output will be when you use `doc-a` and `doc-b` as the arguments for `bag-equality`? _____ Were you correct? _____

Assessing the model on our elephants texts

- 8) Use `string-to-bag` to build one bag of words for the `wiki-article`, and another for the `student-essay`. Do the bags of words appear to be identical to you, or not? _____

- 9) Our `simple-equality` function returned `false` when we compared `wiki-article` with `student-essay`, because the two documents are not identical. What do you predict that `bag-equality` will produce, when we compare `wiki-article` with `student-essay`? Why? _____

- 10) Use `bag-equality` to compare `wiki-article` and `student-essay`. Was your prediction correct? _____

Reflect

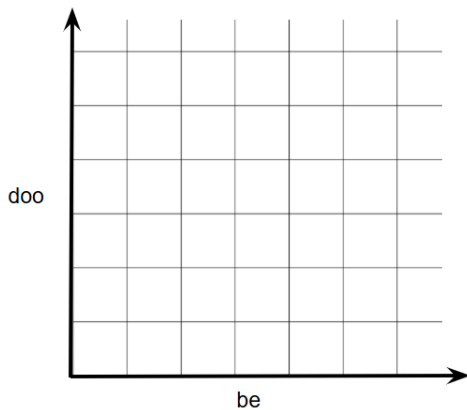
- 11) How is the `bag-equality` plagiarism detector different from our primitive `simple-equality` plagiarism detector? How is it similar?

Plotting Bags of Words

Complete the table below by filling in any blank cells. We've filled in the first row for you. Some cells have more than one possible solution.

Document	Text	Bag-of-Words Summary	Ordered Pair
A	be be be be be be	be: 6, doo: 0	(6,0)
B	doo doo doo doo doo doo		
C	doo doo doo be be be		
D	doo be doo be doo be		
E	doo doo be be be be		
F	be be doo doo doo doo		
G	doo be doo doo doo doo		
H			(5,1)
I		be: 1, doo: 2	
J	doo doo be be		
K		be: 6, doo: 6	
L			(4,6)

1) Plot and label the points with the appropriate letter on the coordinate plane (below), then record what you Notice and Wonder.



What do you Notice?

What do you Wonder?

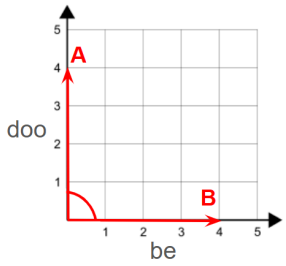
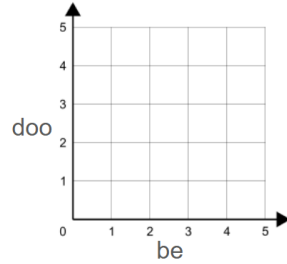
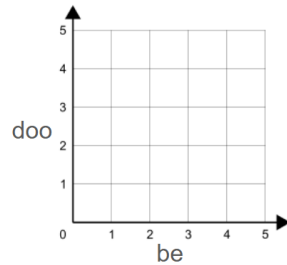
2) Jane loves song A and wants to listen to a song that is similar.

- Sierra suggests song H, claiming that it is similar to song A because five of its six words are the same.
- Jaden recommends song B. He says those two songs are similar because they are the only songs that have just one word!

Who do you agree with?

Angle Difference (with Protractor)

Fill in the blank table cells, then use them to determine the ordered pairs. Plot and label point A and point B on the coordinate plane. Draw a ray from the origin to each of the points. Measure and record the size of the angle (in degrees) formed by the rays.

<p>1 StringA: doo doo doo doo</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td>0</td> </tr> <tr> <td>doo</td> <td>4</td> </tr> </tbody> </table> <p>Ordered pair: <u>(0,4)</u></p>	Word	Frequency	be	0	doo	4	<p>StringB: be be be be</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td>4</td> </tr> <tr> <td>doo</td> <td>0</td> </tr> </tbody> </table> <p>Ordered pair: <u>(4,0)</u></p>	Word	Frequency	be	4	doo	0	 <p>The angle formed is _____.</p>
Word	Frequency													
be	0													
doo	4													
Word	Frequency													
be	4													
doo	0													
<p>2 StringA: doo be</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td></td> </tr> <tr> <td>doo</td> <td></td> </tr> </tbody> </table> <p>Ordered pair: _____</p>	Word	Frequency	be		doo		<p>StringB: doo doo doo doo</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td></td> </tr> <tr> <td>doo</td> <td></td> </tr> </tbody> </table> <p>Ordered pair: _____</p>	Word	Frequency	be		doo		 <p>The angle formed is _____.</p>
Word	Frequency													
be														
doo														
Word	Frequency													
be														
doo														
<p>3 StringA: doo be</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td></td> </tr> <tr> <td>doo</td> <td></td> </tr> </tbody> </table> <p>Ordered pair: _____</p>	Word	Frequency	be		doo		<p>StringB: be doo doo doo</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td></td> </tr> <tr> <td>doo</td> <td></td> </tr> </tbody> </table> <p>Ordered pair: _____</p>	Word	Frequency	be		doo		 <p>The angle formed is _____.</p>
Word	Frequency													
be														
doo														
Word	Frequency													
be														
doo														
<p>4 StringA: doo doo be doo be</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td></td> </tr> <tr> <td>doo</td> <td></td> </tr> </tbody> </table> <p>Ordered pair: _____</p>	Word	Frequency	be		doo		<p>StringB: doo be doo be doo</p> <table border="1"> <thead> <tr> <th>Word</th> <th>Frequency</th> </tr> </thead> <tbody> <tr> <td>be</td> <td></td> </tr> <tr> <td>doo</td> <td></td> </tr> </tbody> </table> <p>Ordered pair: _____</p>	Word	Frequency	be		doo		 <p>The angle formed is _____.</p>
Word	Frequency													
be														
doo														
Word	Frequency													
be														
doo														

Angle Difference (No Protractor)

Fill in the blank table cells, then use them to determine the ordered pairs. Plot and label point A and point B on the coordinate plane. Draw a ray from the origin to each of the points. Approximate and record the size of the angle (in degrees) formed by the rays.

1 StringA: doo doo doo doo

Word	Frequency
be	0
doo	4

Ordered pair: (0,4)

StringB: be be be be

Word	Frequency
be	4
doo	0

Ordered pair: (4,0)

The angle formed is 90.

2 StringA: doo be

Word	Frequency
be	
doo	

Ordered pair: _____

StringB: doo doo doo doo

Word	Frequency
be	
doo	

Ordered pair: _____

The angle formed is _____.

3 StringA: doo doo be be

Word	Frequency
be	
doo	

Ordered pair: _____

StringB: be be be be

Word	Frequency
be	
doo	

Ordered pair: _____

The angle formed is _____.

4 StringA: doo doo be doo be

Word	Frequency
be	
doo	

Ordered pair: _____

StringB: doo be doo be doo

Word	Frequency
be	
doo	

Ordered pair: _____

The angle formed is _____.

The Importance of Human Judgment

How Many Dimensions?

A teacher asks students to write an essay on one of three topics: modern art, soccer, or zebras. The teacher feeds the essays to a plagiarism detector to **compare the essays to one another**. It plots 32 points in a 1023-dimensional space.

1) What does the number 32 tell us? _____

2) What does the number 1023 tell us? _____

3) The teacher knows, of course, that students don't merely copy one another's writing; instead, they copy/paste text from the internet! The teacher opens her plagiarism detection software and tells it to consume a large **training corpus** of internet text about modern art, soccer, and zebras. How many dimensions will the space contain, now?

Circle one: more dimensions the same number of dimensions fewer dimensions

4) Explain your choice.

Assessing the Output

On the left is a description of the output of a plagiarism detector. On the right is an interpretation of what each output means. Match the description of each output with the correct interpretation.

The 32 points are evenly scatter all over the 1023-dimensional space, with no apparent clustering.

1

A

Essays about one topic are likely to use similar words, but essays about different topics likely use different words. The three clusters probably represent the three different topics (modern art, zebras, and soccer).

The points are clustered into three groups, but two of the points are almost on top of one another.

2

B

Two of the essays use almost the same words, in almost the same frequencies. This look suspicious, and the teacher should take a closer look at those essays for potential plagiarism.

The points are clustered into three groups.

3

C

Only one student opted to write about a particular topic, while the rest of the class wrote essays on the other two. There is no evidence of plagiarism.

Almost all the points are clustered into two groups, but one point is positioned very far away from both clusters.

4

D

Every student appears to have written essays about entirely different topics! The teacher has no reason to suspect plagiarism, but they should *probably* worry about their students not following the assignment!

Exploring the Model

Use [Plagiarism Detection Starter File](#) to respond to the questions, below.

distance-to

1) In the Interactions Area, type `distance-to(student-essay)`. The table that appears displays the angle difference between `student-essay` and each of the other animal articles in the corpus. Use the information to fill in the blanks, below.

	article	angle difference
most similar to <code>student-essay</code>		
most different from <code>student-essay</code>		

2) In the Interactions Area, type `distance-to(chimpanzee-article)`. One of the angle differences listed in the table that appears is zero. What does that tell you? _____

3) Use the output of `distance-to(chimpanzee-article)` to fill in the blanks, below.

	article	angle difference
most similar to <code>chimpanzee-article</code>		
most different from <code>chimpanzee-article</code>		

4) Is `distance-to` sophisticated enough to be able to determine with certainty whether or not plagiarism occurred? Explain.

5) Can you think of any ways to improve this model? _____

string-to-bag-cleaned

6) In the Interactions Area, type `string-to-bag(snail-article)`.

a) How many unique words are used in the snail article? _____

b) List each word that is used 2 or more times in the article. We've provided the first three: a (2), and (3), be (3), _____

c) What do you Notice about the commonly used words? What do you Wonder? _____

7) In the Interactions Area, type `string-to-bag-cleaned(snail-article)`.

a) Notice that there are fewer unique words listed. How many unique words are on the table now? _____

b) List each word that appears 2 or more times, now: _____

distance-to-cleaned

Use [Plagiarism Detection Starter File](#) to respond to the questions, below.

Part 1: chimpanzee-article

Let's compare the output of `distance-to(chimpanzee-article)` with the output of `distance-to-cleaned(chimpanzee-article)`.

- 1) On both output tables, the angle produced in the "chimpanzee" row is zero. Why? _____

- 2) When we used `distance-to-cleaned`, the angles were generally larger. Why? _____

- 3) When we used `distance-to-cleaned`, four articles produced an angle measure of 90°. What does this tell us? _____

Part 2: mystery-article

Follow the steps below to make some discoveries about the `mystery-article`.

- 4) Type `distance-to(mystery-article)` in the Interactions Area. Look at the different angle measurements produced. Do you notice any outliers? _____ Explain. _____

- 5) Think back to when you compared `student-essay` with `elephant-article`. Those two paragraphs were virtually identical, which resulted in an angle difference of about 23.706°. The difference between these two rays (`giraffe-article` and `mystery-article`) is much greater, but still seems unusual. Can you make any hypotheses about `mystery-article`? _____

- 6) Let's investigate further. Type `distance-to-cleaned(mystery-essay)`. What do you notice? _____

- 7) Time for the reveal! Type `mystery-article` into the Interactions Area. Why was the angle difference between this article and `giraffe-article` small? _____

Part 3: Revisiting student-essay

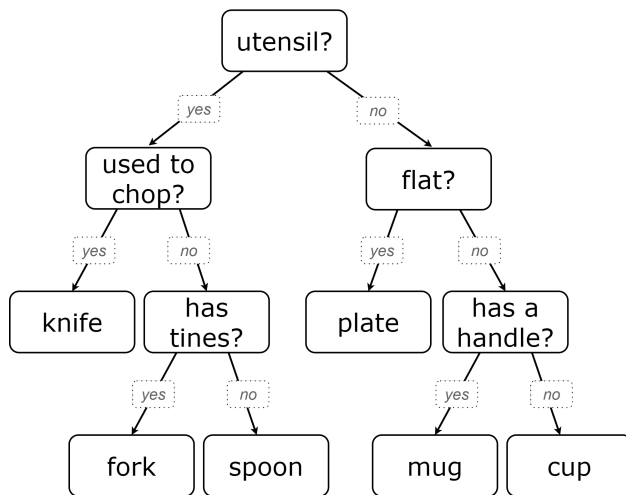
- 8) Type `distance-to-cleaned(student-essay)` in the Interactions Area. What do you Notice? What do you Wonder?

Decision Tree: Spoon, Fork, Knife, Plate, Mug, Cup

Connecting Table and Decision Trees

Item	flat?	has-handle?	has-tines?	utensil?	used-to-chop?	category
A	no	yes	no	no	no	cup
B	no	yes	yes	yes	no	fork
C	yes	yes	no	yes	yes	knife
D	no	no	no	no	no	mug
E	yes	yes	no	no	no	plate
F	no	no	no	yes	no	spoon
G	yes	yes	no	yes	yes	knife

The decision tree below is one possible model for working with the data contained in the table above.



1) What do you Notice? What do you Wonder?

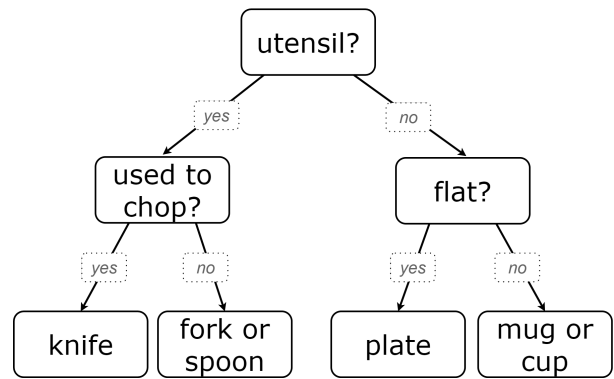
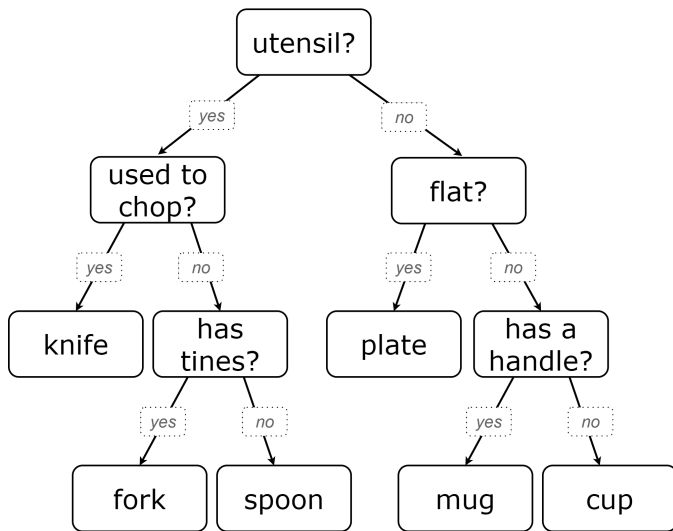
Another possible decision tree from this data

2) Make a second decision tree for the same 6 items listed above using "flat?" as the root node with "used-to-chop" as the first decision node for **yes** and "utensil?" as the first decision node for **no**

- For each question, add the missing "yes" or "no" arrow.
- Determine whether your arrow will lead to
 - a leaf node (a single item from the list)
 - a decision node (a new question)
- Keep adding to the decision tree until you've created pathways that lead to all 6 categories: Spoon, Fork, Knife, Plate, Mug, Cup

Testing and Comparing Decision Trees

Determine how each of these trees would categorize the items below. Explain your responses.



1) butter knife _____

2) chopstick _____

3) spork _____

4) butter knife _____

5) chopstick _____

6) spork _____

7) What do you Notice about the two decision trees above?

8) What do you Wonder about the two decision trees?

9) Which decision tree (above) would make predictions more efficiently? Explain.

Hint: A program is **efficient** when it performs a task with minimal time, memory, energy or data.

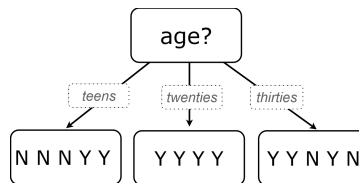
10) Which decision tree (above) would make predictions more accurately? Explain.

Decision Tree: Predicting Shopping Behavior — Part 1

name	age	shopping history	interest in game	buys game
Jan	teens	previous customer	no	no
Jose	teens	previous customer	no	no
Maribel	twenties	previous customer	no	yes
Noah	thirties	previous customer	no	yes
Sydney	thirties	previous customer	yes	yes
Mariana	thirties	new customer	yes	no
Rasula	twenties	new customer	yes	yes
Jillian	teens	previous customer	no	no
Ariella	teens	new customer	yes	yes
Isabela	thirties	previous customer	yes	yes
Danial	teens	previous customer	yes	yes
Kate	twenties	previous customer	no	yes
Taikhoom	twenties	previous customer	yes	yes
Peter	thirties	new customers	no	no

"Age" as the Root Node

The **decision stump** below splits the above training data by age and indicates whether the individuals in each group buy the game.



1) Where do the Y/N lists beneath each of the three branches come from? _____

2) What prediction will our current model (decision stump) make for each group?

- people in their teens _____ buy the game
will / will not
- people in their twenties _____ buy the game
will / will not
- people in their thirties _____ buy the game
will / will not

3) Place checkmarks below each of the values in the stump leaves for which our prediction is correct.

4) Find the likelihood of a correct prediction for each age group. teens: _____% twenties: _____% thirties: _____%

5) How accurate is the current prediction across our entire dataset? _____ correct predictions out of 14 attempts. (_____ % accuracy).

Improving Our Prediction

We made our prediction without considering all of the columns in our training data. If we add another level to our tree, we might be able to improve our accuracy!

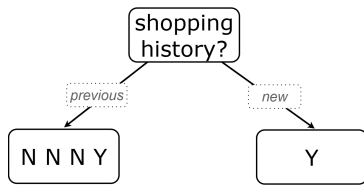
6) Before moving on to the second level of his decision tree, Ernie removed all of the rows for people in their twenties. Bert said, "I don't think that's a good idea! Why would we alter our dataset just because we're starting the second level of the tree?" Explain Ernie's (correct) decision to Bert.

7) We used "age" as our root node. What questions could we ask at our second-level decision node? _____ or _____
[column] [column]

Decision Tree: Predicting Shopping Behavior – Part 2

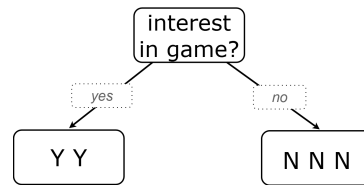
For Teens – Is Shopping History or Interest in Game a better Indicator?

1) Use the decision stumps below to make rules that predict whether or not teens would buy the game based on each category.



Predict that

- previous teenage customers will buy / will not buy / no prediction
- new teenage customers will buy / will not buy / no prediction



Predict that

- interested teenage shoppers _____
will buy / will not buy / no prediction
- uninterested teenage shoppers _____
will buy / will not buy / no prediction

Place a checkmark below each value that the computer would predict correctly.

2) Use the decision stumps to determine how well each rule predicts whether or not teens would buy the game.

For our training data, this rule

- would make _____ correct predictions out of 5 attempts
- is _____ % accurate

For our training data, this rule

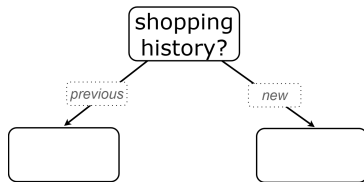
- would make _____ correct predictions out of 5 attempts
- is _____ % accurate

3) Will you use shopping history or interest for your second decision node for teens?

For People in their Thirties – Is Shopping History or Interest in Game a better Indicator?

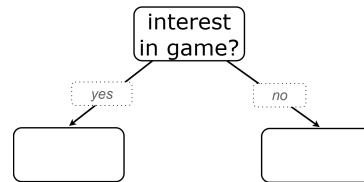
Refer to the training data on [Decision Tree: Predicting Shopping Behavior – Part 1](#) to complete the stumps below.

4) Complete the decision stumps below to make rules that predict whether or not people in their thirties would buy the game.



Predict that

- new customers in their 30s will buy / will not buy / no prediction
- previous customers in their 30s will buy / will not buy / no prediction



Predict that

- interested customers in their 30s will buy / will not buy / no prediction
- uninterested customers in their 30s will buy / will not buy / no prediction

5) Add checkmarks to the correct predictions. Then determine how well each rule predicts which customers in their thirties buy the game.

For our training data, this rule

- would make _____ correct predictions out of 5 attempts
- is _____ % accurate

For our training data, this rule

- would make _____ correct predictions out of 5 attempts
- is _____ % accurate

6) Will you use shopping history or interest for your second decision node for customer in their thirties?

Building and Testing a Decision Tree

For this page you will need to refer to your decisions from [Decision Tree: Predicting Shopping Behavior – Part 2](#).

Build and Understand the Tree

Complete the tree (left), then fill in the blanks (right).

1) The root node of this tree is _____.

2) The first set of branches includes:
_____, _____, and _____.

3) Write the rules that this decision tree follows. Predict that:

_____ interested / uninterested teens will buy the game.

_____ in their 20s will buy the game.

_____ who?

_____ in their 30s will buy the game.

_____ who?

Test the Tree

4) Below is a new set of potential customers. Use the decision tree and rules above to predict whether or not they will buy the game.

name	age	shopping history	interest in game	buys game	model predicts
Kat	teen	new customer	yes	yes	
Billy	twenties	new customer	no	no	
Chen	twenties	previous customer	no	no	

5) Compare the "buys game" column with the predictions. For which customers was the computer correct? Kat Billy Chen

6) Should we change any of our rules based on the addition of this new data? Why? _____

Reflect

7) Our rules made 100% accurate predictions with our training dataset, but were only 33% accurate with our test dataset. Why?

8) What could we do to improve the quality of this decision tree? _____

Decision Tree Training Dataset

name	age	shopping history	interest in game	buys game
Jan	16	previous customer	no	no
Jose	19	previous customer	no	no
Maribel	21	previous customer	no	yes
Noah	33	previous customer	no	yes
Sydney	36	previous customer	yes	yes
Mariana	32	new customer	yes	no
Rasula	24	new customer	yes	yes
Jillian	14	previous customer	no	no
Ariella	16	new customer	yes	yes
Isabela	38	previous customer	yes	yes
Danial	19	previous customer	yes	yes
Kate	27	previous customer	no	yes
Taikhoom	22	previous customer	yes	yes
Peter	35	new customers	no	no

Constructing a Language Model

Use [There Was an Old Lady Who Swallowed a Fly](#) to answer the questions, below. Note that the title of the song is a part of the corpus!

Did she swallow a fly?

Let's say we want to know the likelihood that the old lady "swallowed a fly". We need a ratio!

$$1) \quad p(\text{fly} \mid \text{swallowed a}) = \frac{\text{count}(\text{swallowed a fly})}{\text{count}(\text{swallowed a...})} = \frac{\quad}{10}$$

2) The likelihood that the old lady "swallowed a fly" in the excerpt is _____.

What about the other animals?

Find the ratio representing the likelihood that she swallowed each of the other animals.

$$3) \quad p(\text{spider} \mid \text{swallowed a}) = \frac{\text{count}(\text{swallowed a spider})}{\text{count}(\text{swallowed a...})} = \frac{\quad}{10}$$

$$4) \quad p(\text{bird} \mid \text{swallowed a}) = \frac{\text{count}(\text{swallowed a bird})}{\text{count}(\text{swallowed a...})} = \frac{\quad}{10}$$

$$5) \quad p(\text{cat} \mid \text{swallowed a}) = \frac{\text{count}(\text{swallowed a cat})}{\text{count}(\text{swallowed a...})} = \frac{\quad}{10}$$

6) Across the entire poem, the lady *most likely* swallowed a _____.

7) The second most likely option is that she swallowed a _____.

8) There is an equal probability that she swallowed a _____ or a _____.

To catch the...?

Find the ratio representing the likelihood that she caught each of the animals.

$$9) \quad p(\text{fly} \mid \text{to catch the}) = \frac{\text{count}(\text{to catch the fly})}{\text{count}(\text{to catch the...})} = \frac{\quad}{6}$$

$$10) \quad p(\text{spider} \mid \text{to catch the}) = \frac{\text{count}(\text{to catch the spider})}{\text{count}(\text{to catch the...})} = \frac{\quad}{6}$$

$$11) \quad p(\text{bird} \mid \text{to catch the}) = \frac{\text{count}(\text{to catch the bird})}{\text{count}(\text{to catch the...})} = \frac{\quad}{6}$$

The lady was most likely trying to catch the _____. It is the least likely that she was trying to catch the _____.

The...?

$$12) \quad p(\text{fly} \mid \text{the}) = \frac{\quad}{\quad} \quad 13) \quad p(\text{spider} \mid \text{the}) = \frac{\quad}{\quad} \quad 14) \quad p(\text{bird} \mid \text{the}) = \frac{\quad}{\quad} \quad 15) \quad p(\text{cat} \mid \text{the}) = \frac{\quad}{\quad}$$

Sampling from the Language Model

Use [There Was an Old Lady Who Swallowed a Fly](#) to answer the questions, below. **Note that the title of the song is a part of the corpus!**

Text Generation Exercise #1

1) Let's start our text generation with the phrase "She ".

Text Generation 1: She
first word second word third word fourth word

2) Determine which word is most likely to follow "She " and record it on the line for second word above.

3) Determine which word is most likely to follow the word you just wrote down and record it on the line for third word above.

4) Use statistical language modeling to determine the fourth word.

5) Everyone in your class should have generated the **same** text. Why do you think that was the outcome?

Text Generation Exercise #2

Here is a list of the most common unigrams in the corpus: "the": 12 times "she": 11 times "a": 11 times

6) Let's start by choosing the **most** common word:

7) Determine which word is most likely to follow that word:

8) There are two words that have an equal probability of appearing in the third spot! What are they? and

9) Flip a coin to determine which one you will use to complete **Text Generation 2a** and use statistical language modeling to determine the fourth word. Then repeat the process for **Text Generation 2b** using the other word that tied for the third spot.

Text Generation 2a: the
first word second word third word fourth word

Text Generation 2b: the
first word second word third word fourth word

10) Why was there only one result for **Text Generation Exercise 1**, while **Text Generation Exercise 2** had two possible results?

★ What is another starting word that would result in only one possible four-word phrase?

★★ What is another starting word that would result in at least two possible outcomes?

There Was an Old Lady Who Swallowed a Fly

1 There was an old lady who swallowed a fly,
2 I don't know why she swallowed a fly – perhaps she'll die!

3

4 There was an old lady who swallowed a spider
5 That wriggled and jiggled and tickled inside her;
6 She swallowed the spider to catch the fly;
7 I don't know why she swallowed a fly – perhaps she'll die!

8

9 There was an old lady who swallowed a bird;
10 How absurd to swallow a bird!
11 She swallowed the bird to catch the spider
12 That wriggled and jiggled and tickled inside her,
13 She swallowed the spider to catch the fly;
14 I don't know why she swallowed a fly – perhaps she'll die!

15

16 There was an old lady who swallowed a cat;
17 Well, fancy that, she swallowed a cat!
18 She swallowed the cat to catch the bird,
19 She swallowed the bird to catch the spider
20 That wriggled and jiggled and tickled inside her,
21 She swallowed the spider to catch the fly;
22 I don't know why she swallowed a fly – perhaps she'll die!

Meet Soekia!

Respond to the prompts below by interacting with <https://bootstrapworld.org/Soekia/>.

Generate text

The blue panel occupying most of your screen is where text generation takes place. This is the level typically visible to the user when using a chatbot.

1) To tell Soekia to start writing, click on the ► on the right. Read the text until Soekia finishes writing. Take a minute to scroll and click on the other icons in this panel: ↺ 📄 🤖 😊 ●○. What do you Notice? What do you Wonder?

2) Hover your mouse over one or two of the highlighted words in the fairy tale. What appears? (If the words in your fairy tale aren't already highlighted with different colors, click the 🔍.)

Documents

The green *Documents* panel is where the training corpus lives. You can access it by clicking the "LOOK INSIDE →" button in the upper right-hand corner of [Soekia](#) and scrolling all the way to the right.

3) How many documents are there in this particular corpus? _____

4) Give the title of one of the documents: _____

5) Click on the 📁 (the middle icon in the upper right). Hover over each of the 7 icons that drop down. Which one are you most interested in exploring? _____

6) Take a minute to scroll and click on any remaining icons in this panel. What else do you Notice about the *Documents* panel? What do you Wonder?

Meet Soekia! (n-grams)

Respond to the prompts below by interacting with <https://bootstrapworld.org/Soekia/> using the 🐒 **Intelligent Monkeys?** collection.

Exploring the N-Grams Panel

The orange *N-grams* panel is where Soekia lists possible *n-grams* and how frequently they occur in the training corpus. The default setting (**3**) displays a list of every *trigram*. Clicking on the other numbers at the top will display lists of *n-grams* of other lengths.

1) Hover your mouse over the **3**. How many different *trigrams* are there in this collection? _____ How does that compare to the number of *n-grams* of other lengths? _____

2) The most common *trigram* appears at the top of the list. Click on it. What do you learn? _____

3) Click on the **5** tab. Notice that all of the 5-word *n-grams* occur equally often. Can you explain why this might be? _____

4) Take a minute to explore the *N-grams* Panel. What do you Notice? What do you Wonder?

Predicting the Next Word using N-Grams

For this section, make sure you are in the 🐒 **Intelligent Monkeys?** collection. Go to the *Suggested words* panel, click on "Customize temperature/number of suggestions", and set the temperature to **low**.

Without introducing any randomization into the algorithm, Soekia generates text by selecting words one at a time from the most-frequent valid *n-gram* of the highest order available.

5) For the first "word", Soekia looks in the **1** tab to find the most frequently occurring *unigram*. What do you expect it to choose? _____

6) To choose the second "word", Soekia:

- Looks at the **2** tab to find the most frequently occurring *bigram* that begins with the first "word".
- If there isn't one, it will return to the **1** tab and select the next most popular *unigram*.

What do you expect Soekia to choose? _____ Which list did you select it from? _____

7) Why do you think there weren't any bigrams that began with the most popular "word"? *Hint: Read the documents closely!* _____

8) To choose the third "word", Soekia:

- Looks at the **3** tab to find the most frequently occurring *trigram* that begins with the first and second "words".
- If there isn't one, it will look in the **2** tab for the most frequently occurring *bigram* beginning with the second "word".
- If there isn't one, it will return to the **1** tab and select the next most popular *unigram*.

What do you expect Soekia to choose? _____ Which list did you select it from? _____

9) Continuing this process, what do you expect Soekia to choose for the:

fourth "word"? _____ fifth "word"? _____ sixth "word"? _____

Testing our Prediction

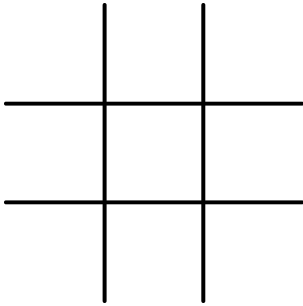
For this section, make sure you are still in the 🐒 **Intelligent Monkeys?** collection with the temperature set to **low**.

10) How does Soekia answer the question *How intelligent are monkeys?*... when you click ➤ _____

11) How does that text compare to your prediction? _____

Tic Tac Toe

Warmup: Play A Game of Tic Tac Toe



In case you need a refresher on the game:

- The tic-tac-toe board is a 3x3 grid.
- One person will draw an **X** in one of the squares.
- The other person will draw an **O** in one of the squares.
- Keep taking turns — the goal is to get three in a row or block your neighbor from getting three in a row.
- The game ends when one of you gets three in a row or the grid is full.

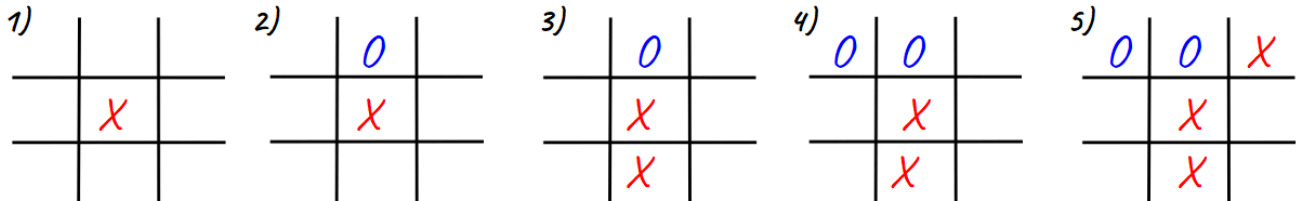
Design a Tic Tac Toe Notation

In order to communicate with Soekia about tic tac toe games, we'll need to record the moves using an annotation.

- Let's think of the 3x3 tic-tac-toe grid as a first quadrant coordinate plane with the origin (0,0) in the bottom left corner
- For each move, our notation must indicate:
 - the player whose turn it is (X or O)
 - the ordered pair (x, y) for the location of the player's move on that turn

1) How would you annotate the first turn (above)? _____ The second turn? _____

Annotating a Game of Tic Tac Toe



2) Complete the table by translating the 5 turn sequence above into our tic tac toe annotation.

For Reference:

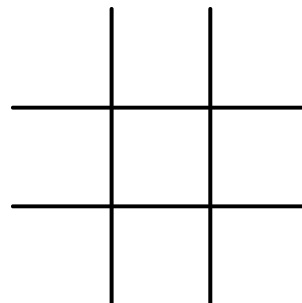
- If player X makes a move in the bottom right corner, we would describe that turn as: X31
- If player O makes a move in the middle of the left column, we would describe that turn as: O12

1st move	2nd move	3rd move	4th move	5th move

Translating Notation to the Game Board

3) Translate the list of moves below into a game played on the tic tac toe board.

- X32
- O22
- X12
- O11
- X33
- O13
- X31



4) Is there a winner? _____

What Makes a Language?

Respond to the prompts below by interacting with <https://bootstrapworld.org/Soekia/>.

Soekia & Tic-Tac-Toe



Access the *Documents* panel by clicking the "LOOK INSIDE →" button in the upper right-hand corner of [Soekia](https://bootstrapworld.org/Soekia/).

- Click  and select  from the drop-down menu.

1) Describe what you see. _____

2) How are the documents of the tic-tac-toe collection **similar** to the documents of the fairy tale collection? How are they **different**?

In the *Generate Text* panel,

- Click  to tell Soekia to generate a set of game moves.
- Click  **TicTacToe-player** (beneath the list of moves) to see a visualization of the game you just generated. *Note that the player only appears after you have generated and paused a tic-tac-toe game.*

3) In what ways does the game seem **similar** to a normal game of Tic Tac Toe? In what ways does it seem **different**?

Thinking About Natural Language Processing

In the *Documents* panel,

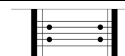
- Click  and select  from the dropdown menu.

In the *Generate Text* panel,

- Click  to tell Soekia to write you some music. Click  **Music-Player** to try playing your composition.

4) What do you Notice? What do you Wonder? _____



5) In musical notation, a "repeat sign" (pictured on right) indicates that a section of the music should be repeated. Does Soekia's music player repeat when it arrives at a repeat sign? Why or why not? (*Stuck? Test it out in Soekia.*)



In the *Documents* panel,

- Click  and select  from the bottom of the dropdown menu.

In the *Generate Text* panel,

- Click  to tell Soekia to generate a set of chess moves. Click  **Chess-Player** to visualize the game.

6) In actual chess, a player starts with two knights and there is no way to gain additional knights. How come additional knights sometimes appear in this chess game?

12 famous AI disasters

Feature

Jun 9, 2025

11 mins

By Thor Olavsrud

Source: www.cio.com

Insights from data and machine learning algorithms can be invaluable, but be warned — mistakes can be irreversible. These recent high-profile AI blunders illustrate the damage done when things don't go according to plan.

In 2017, [The Economist](#) declared that data, rather than oil, had become the world's most valuable resource, and the refrain has been repeated ever since. Organizations across every industry have been investing, and continue to heavily invest, in data and analytics. But like oil, data and analytics have their dark side. According to CIO's State of the CIO Survey 2025, 42% of CIOs say AI and ML are their biggest technology priority for 2025. And while actions driven by [ML algorithms](#) can give organizations a competitive advantage, mistakes can be costly in terms of reputation, revenue, or even lives.

Understanding your data and what it's telling you is important, but it's equally vital to understand your tools, know your data, and keep your organization's values firmly in mind. And with that in mind, here are a handful of high-profile AI blunders from the past decade to illustrate what can go wrong.

Chicago Sun-Times, Philadelphia Inquirer publish summer reading list of fake books

The Chicago Sun-Times and Philadelphia Inquirer took reputational hits when May 2025 editions featured a special section that included a summer reading list recommending books that don't exist.

The [Chicago Sun-Times explained](#) that the syndicated section, "Heat Index: Your Guide to the Best of Summer," was provided by King Features Syndicate, a unit of Hearst. Marco Buscaglia, the author of the special section, admitted he used AI to assist putting it together, including the recommended reading list, and failed to fact check the output.

The reading list featured many real authors but attributed nonexistent books to them. For example, the first recommendation was Tidewater Dreams by famed Chilean-American writer Isabel Allende, who's written more than 20 novels. But Tidewater Dreams, a "climate fiction novel that explores how one family confronts rising sea levels while uncovering long-buried secrets," isn't one of them. Like most books on the list, it was hallucinated by AI.

The newsrooms of both papers said they had nothing to do with the insert, though neither paper marked it as advertorial content. King Features terminated its relationship with Buscaglia following the incident, noting that his use of AI violated a strict policy.

McDonald's ends AI experiment after drive-thru ordering blunders

After working with IBM for three years to leverage [AI to take drive-thru orders](#), McDonald's called the whole thing off in June 2024. The reason? A slew of social media videos showing confused and [frustrated customers](#) trying to get the AI to understand their orders.

One TikTok video in particular featured two people repeatedly pleading with the AI to stop as it kept adding more Chicken McNuggets to their order, eventually reaching 260. In a June 13, 2024, internal memo obtained by trade publication Restaurant Business, McDonald's announced it would end the partnership with IBM and shut down the tests.

The restaurant had piloted the AI at more than 100 US drive-thrus, and indicated it still saw a future in a voice-ordering solution.

Grok AI falsely accuses NBA star of vandalism spree

In an April 2024 post on X, Grok, the AI chatbot from [Elon Musk's xAI](#), falsely accused NBA star Klay Thompson of throwing bricks through windows of multiple houses in Sacramento, California.

Some commentators speculated that Grok may have hallucinated the vandalism story about the Golden State Warriors small forward, who's since moved to the Dallas Mavericks, after ingesting posts about Thompson "throwing bricks," common basketball parlance for badly missed shots. In his final game with GSW, the team suffered a blowout loss as part of the worst postseason performance of Thompson's career.

Grok does display a disclaimer saying, "Grok is an early feature and can make mistakes. Verify its outputs." Still, the incident raises questions about liability when an AI chatbot makes false and defamatory statements.

NYC AI chatbot encourages business owners to break the law

In March 2024, [The Markup](#) reported that Microsoft-powered chatbot MyCity was giving entrepreneurs incorrect information that would lead to them break the law.

Unveiled in October 2024, MyCity was intended to help provide New Yorkers with information on starting and operating businesses in the city, as well as housing policy and worker rights. The only problem was The Markup found MyCity falsely claimed that business owners could take a cut of their workers' tips, fire workers who complain of sexual harassment, and serve food that had been nibbled by rodents. It also claimed landlords could discriminate based on source of income.

In the wake of the report, indicted New York City Mayor Eric Adams defended the project. The chatbot remains online.

Air Canada pays damages for chatbot lies

In February 2024, Air Canada was ordered to [pay damages to a passenger](#) after its virtual assistant gave him incorrect information at a particularly difficult time.

Jake Moffatt consulted Air Canada's virtual assistant about bereavement fares following the death of his grandmother in November 2023. The chatbot told him he could buy a regular price ticket from Vancouver to Toronto and apply for a bereavement discount within 90 days of purchase. Following that advice, Moffatt purchased a one-way CA\$794.98 ticket to Toronto and a CA\$845.38 return flight to Vancouver.

But when Moffatt submitted his refund claim, the airline turned him down, saying bereavement fares can't be claimed after tickets have been purchased.

Moffatt took Air Canada to a tribunal in Canada, claiming the airline was negligent and misrepresented information via its virtual assistant. [According to tribunal member](#) Christopher Rivers, Air Canada argued it can't be held liable for the information provided by its chatbot.

Rivers denied that argument, saying the airline didn't take "reasonable care to ensure its chatbot was accurate," so he ordered the airline to pay Moffatt CA\$812.02, including CA\$650.88 in damages.

Sports Illustrated may have published AI-generated writers

In November 2023, online magazine Futurism said Sports Illustrated was publishing articles [by AI-generated writers](#).

Futurism cited anonymous sources were involved to create content, and said the storied sports magazine published a lot of articles by authors generated by AI.

Futurism also found the author headshots in question were listed on a site that sells AI-generated portraits. The online magazine then reached out to The Arena Group, publisher of Sports Illustrated, and in a statement, Arena Group said the articles in question were licensed content from a third party, AdVon Commerce.

"We continually monitor our partners and were in the midst of a review when these allegations were raised," Arena Group said in the statement provided to Futurism. "AdVon has assured us that all of the articles in question were written and edited by humans."

The statement added that AdVon writers used pseudonyms in certain articles, noting that Arena Group doesn't condone those actions and subsequently removed the articles in question from the Sports Illustrated website. Responding to the Futurism article, the Sports Illustrated Union [posted a statement](#) that it was horrified by the allegations and demanded answers and transparency from Arena Group management.

"If true, these practices violate everything we believe in about journalism," the SI Union said in its statement. "We deplore being associated with something so disrespectful to our readers."

iTutor Group's recruiting AI rejects applicants due to age

In August 2023, tutoring company iTutor Group [agreed to pay \\$365,000](#) to settle [a suit](#) brought by the US Equal Employment Opportunity Commission (EEOC). The federal agency said the company, which provides remote tutoring services to students in China, used AI-powered recruiting software that automatically rejected female applicants ages 55 and older, and male applicants ages 60 and older.

The EEOC said more than 200 qualified applicants were automatically rejected by the software.

"Age discrimination is unjust and unlawful," former EEOC chair Charlotte Burrows said in a statement. "Even when technology automates the discrimination, the employer is still responsible."

iTutor Group denied any wrongdoing but did decide to settle the suit. As part of the settlement and consent decree, it agreed to adopt new anti-discrimination policies.

ChatGPT hallucinates court cases

Advances made in 2023 by LLMs stoked widespread interest in the transformative potential of gen AI across nearly every industry. OpenAI's ChatGPT was at the center of this surge in interest, but the technology still has a long way to go before it can reliably take over most processes, as attorney Steven Schwartz learned when he found himself in hot water with US District Judge Kevin Castel after using it to research precedents in a suit against Colombian airline Avianca.

Schwartz, an attorney with Levidow, Levidow & Oberman in New York, used the OpenAI gen AI chatbot to find precedent to support a case filed by Avianca employee Roberto Mata for injuries he sustained in 2019. But at least six of the cases submitted in the brief didn't exist. In a [document filed in May 2023](#), Castel noted the cases submitted by Schwartz included false names and docket numbers, along with bogus internal citations and quotes. Schwartz's partner, Peter LoDuca, was Mata's lawyer of record and signed the brief, putting himself in jeopardy as well.

In an affidavit, Schwartz told the court it was the first time he used ChatGPT as a legal research source and was unaware of the possibility that its content could be false. He admitted he hadn't confirmed the sources provided by the AI chatbot, and regretted using gen AI to supplement his legal research, adding he'll never do so in the future without verifying authenticity.

In June 2023, Castel imposed a \$5,000 fine on Schwartz and LoDuca, and in a separate ruling in June, he dismissed Mata's lawsuit against Avianca.

AI algorithms identify everything but COVID-19

Since the [COVID-19 pandemic](#) began in 2020, numerous organizations have sought to apply ML algorithms to help hospitals diagnose or triage patients faster. But according to the UK's Turing Institute, a national center for data science and AI, the predictive tools made little to no difference.

MIT Technology Review chronicled a number of failures, most of which stemmed from errors in the way the tools were trained or tested. The use of mislabeled data, or data from unknown sources, was a common culprit.

Derek Driggs, an ML researcher at the University of Cambridge, together with his colleagues, [published a paper in Nature Machine Intelligence](#) that explored the use of deep learning models to diagnose the virus. But the paper determined the technique wasn't fit for clinical use. For example, Driggs' group found their own model was flawed because it was trained on a data set that included scans of patients that were lying down while scanned, and patients who were standing up. The patients who were lying down were much more likely to be seriously ill, so the algorithm learned to identify COVID risk based on the position of the person in the scan.

A similar example included an algorithm trained with a data set with scans of chests of healthy children. The algorithm learned to identify children, not high-risk patients.

Zillow wrote down millions, slashed workforce due to algorithmic home-buying disaster

In November 2021, online real estate marketplace Zillow told shareholders it would wind down its Zillow Offers operations and cut 25% of the company's workforce — about 2,000 employees — over the next several quarters. The home-flipping unit's woes were the result of the error rate in the ML algorithm it used to predict home prices.

Zillow Offers was a program through which the company made cash offers on properties based on a "Zestimate" of home values derived from an ML algorithm. The idea was to renovate the properties and flip them quickly. But a Zillow spokesperson [told CNN](#) the algorithm had a median error rate of 1.9%, and could be as high as 6.9% for off-market homes.

CNN then reported that Zillow bought 27,000 homes through Zillow Offers since its launch in April 2018, but sold only 17,000 through the end of September 2021. Black swan events like the COVID-19 pandemic and a home renovation labor shortage contributed to the algorithm's accuracy troubles.

Zillow said the algorithm led it to unintentionally purchase homes at higher prices rather than its current estimates of future selling prices, resulting in a \$304 million inventory write-down in Q3 2021. In a conference call with investors following the announcement, Zillow co-founder and CEO Rich Barton said it might have been possible to tweak the algorithm, but ultimately it was too risky.

Dataset trained Microsoft chatbot to spew racist tweets

In March 2016, Microsoft learned that using Twitter interactions as [training data for ML](#) algorithms can have dismaying results.

Microsoft released Tay, an AI chatbot, on the social media platform, and the company described it as an experiment in conversational understanding. The idea was the chatbot would assume the persona of a teenage girl and interact with individuals via Twitter using a combination of ML and natural language processing. Microsoft seeded it with anonymized public data and some material pre-written by comedians, then set it loose to learn and evolve from its interactions on the social network.

Within 16 hours, the chatbot posted more than 95,000 tweets, and those tweets rapidly turned overtly racist, misogynist, and anti-Semitic. Microsoft quickly suspended the service for adjustments and ultimately pulled the plug.

"We are deeply sorry for the unintended offensive and hurtful tweets from Tay, which do not represent who we are or what we stand for, nor how we designed Tay," Peter Lee, corporate VP, Microsoft Research & Incubations (then corporate VP of Microsoft Healthcare), wrote in a post on Microsoft's official blog following the incident.

Lee noted that Tay's predecessor, Xiaoice, released by Microsoft in China in 2014, had successfully conducted conversations with more than 40 million people in the two years prior to Tay's release. What Microsoft didn't take into account was that a group of Twitter users would immediately begin tweeting inappropriate comments to Tay. The bot quickly learned from that material and incorporated it into its own tweets.

Amazon AI-enabled recruitment tool only recommended men

Like many large companies, Amazon is hungry for tools that can help its HR function screen applications for the best candidates. In 2014, Amazon started working on AI-powered recruiting software to do just that. There was only one problem: The system vastly preferred male candidates. In 2018, [Reuters broke the news](#) that Amazon had [scrapped the project](#).

Amazon's system gave candidates star ratings from 1 to 5, but the ML models at the heart of the system were trained on 10 years' worth of résumés submitted to Amazon — most of them from men. As a result of that training data, the system started penalizing phrases in résumés that included the word "women's," and even downgraded candidates from all-women colleges.

At the time, Amazon said the tool was never used by Amazon recruiters to evaluate candidates. The company tried to edit the tool to make it neutral but ultimately decided it couldn't guarantee it wouldn't learn some other discriminatory way of sorting candidates and ended the project.

Contracts for Ai

Contracts tell us how to use a function, by telling us three important things:

1. The **Name**
2. The **Domain** of the function - what kinds of inputs do we need to give the function, and how many?
3. The **Range** of the function - what kind of output will the function give us back?

For example: The contract `triangle :: (Number, String, String) -> Image` tells us that the name of the function is `triangle`, it needs three inputs (a Number and two Strings), and it produces an Image.

With these three pieces of information, we know that typing `triangle(20, "solid", "green")` will evaluate to an Image.

Name	Domain	Range
# circle	:: (<u>Number</u> _{radius} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
circle(50, "solid", "purple")		
# ellipse	:: (<u>Number</u> _{width} , <u>Number</u> _{height} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
ellipse(100, 50, "outline", "orange")		
# isosceles-triangle	:: (<u>Number</u> _{size} , <u>Number</u> _{vertex-angle} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
isosceles-triangle(50, 20, "solid", "grey")		
# overlay	:: (<u>Image</u> _{top} , <u>Image</u> _{bottom})	-> Image
overlay(circle(10, "solid", "black"), square(50, "solid", "red"))		
# radial-star	:: (<u>Num</u> _{points} , <u>Num</u> _{outer} , <u>Num</u> _{inner} , <u>Str</u> _{fill-style} , <u>Str</u> _{color})	-> Image
radial-star(6, 20, 50, "solid", "red")		
# rectangle	:: (<u>Number</u> _{width} , <u>Number</u> _{height} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
rectangle(100, 50, "outline", "green")		
# regular-polygon	:: (<u>Number</u> _{size} , <u>Number</u> _{vertices} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
regular-polygon(25,5, "solid", "purple")		
# rhombus	:: (<u>Number</u> _{size} , <u>Number</u> _{top-angle} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
rhombus(100, 45, "outline", "pink")		
# right-triangle	:: (<u>Number</u> _{leg1} , <u>Number</u> _{leg2} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
right-triangle(50, 60, "outline", "blue")		
# rotate	:: (<u>Number</u> _{degrees} , <u>Image</u> _{img})	-> Image
rotate(45, star(50, "solid", "dark-blue"))		

Name	Domain	Range
# <code>sqr</code>	:: (<u>Number</u>)	-> Number
<code>sqr(4)</code>		
# <code>sqrt</code>	:: (<u>Number</u>)	-> Number
<code>sqrt(4)</code>		
# <code>square</code>	:: (<u>Number</u> _{size} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
<code>square(50, "solid", "red")</code>		
# <code>star</code>	:: (<u>Number</u> _{radius} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
<code>star(50, "solid", "red")</code>		
# <code>star-polygon</code>	:: (<u>Number</u> _{size} , <u>Number</u> _{point-count} , <u>Number</u> _{step-count} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
<code>star-polygon(100, 10, 3, "outline", "red")</code>		
# <code>string-contains</code>	:: (<u>String</u> _{haystack} , <u>String</u> _{needle})	-> Boolean
<code>string-contains("hotdog", "dog")</code>		
# <code>string-length</code>	:: (<u>String</u>)	-> Number
<code>string-length("rainbow")</code>		
# <code>text</code>	:: (<u>String</u> _{message} , <u>Number</u> _{size} , <u>String</u> _{color})	-> Image
<code>text("Zari", 85, "orange")</code>		
# <code>triangle</code>	:: (<u>Number</u> _{size} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
<code>triangle(50, "solid", "fuchsia")</code>		
# <code>triangle-asa</code>	:: (<u>Number</u> _{top-left-angle} , <u>Number</u> _{left-side} , <u>Number</u> _{bottom-angle} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
<code>triangle-asa(90, 200, 10, "solid", "purple")</code>		
# <code>triangle-sas</code>	:: (<u>Number</u> _{bottom-R-side} , <u>Number</u> _{top-R-angle} , <u>Number</u> _{top-side} , <u>String</u> _{fill-style} , <u>String</u> _{color})	-> Image
<code>triangle-sas(50, 20, 70, "outline", "dark-green")</code>		
	::	->
	::	->
	::	->
	::	->



These materials were developed partly through support of the National Science Foundation (awards 1042210, 1535276, 1648684, and 1738598) and are licensed under a Creative Commons 4.0 Unported License. Based on a work at www.BootstrapWorld.org. Permissions beyond the scope of this license may be available by contacting contact@BootstrapWorld.org.

These materials were developed partly through support of the National Science Foundation (awards 1042210, 1535276, 1648684, and 1738598) and are licensed under a Creative Commons 4.0 Unported License. Based on a work at www.BootstrapWorld.org. Permissions beyond the scope of this license may be available by contacting contact@BootstrapWorld.org.