CPSC 100

Computational Thinking

Artificial Intelligence

Instructor: Firas Moosvi
Department of Computer Science
University of British Columbia

Learning Goals

After this lecture, you should be able to:

- Discuss and summarize the AI policy activity results
- Work through the “Turing Test”
- Explain the dilemma of the **Trolley Problem**
  - Describe the relevance of the problem to AI Ethics
- Identify and explain the traditional steps of **Natural Language Processing**
- Apply a traditional NLP algorithm to a given input

Course Admin

Class Activity Results

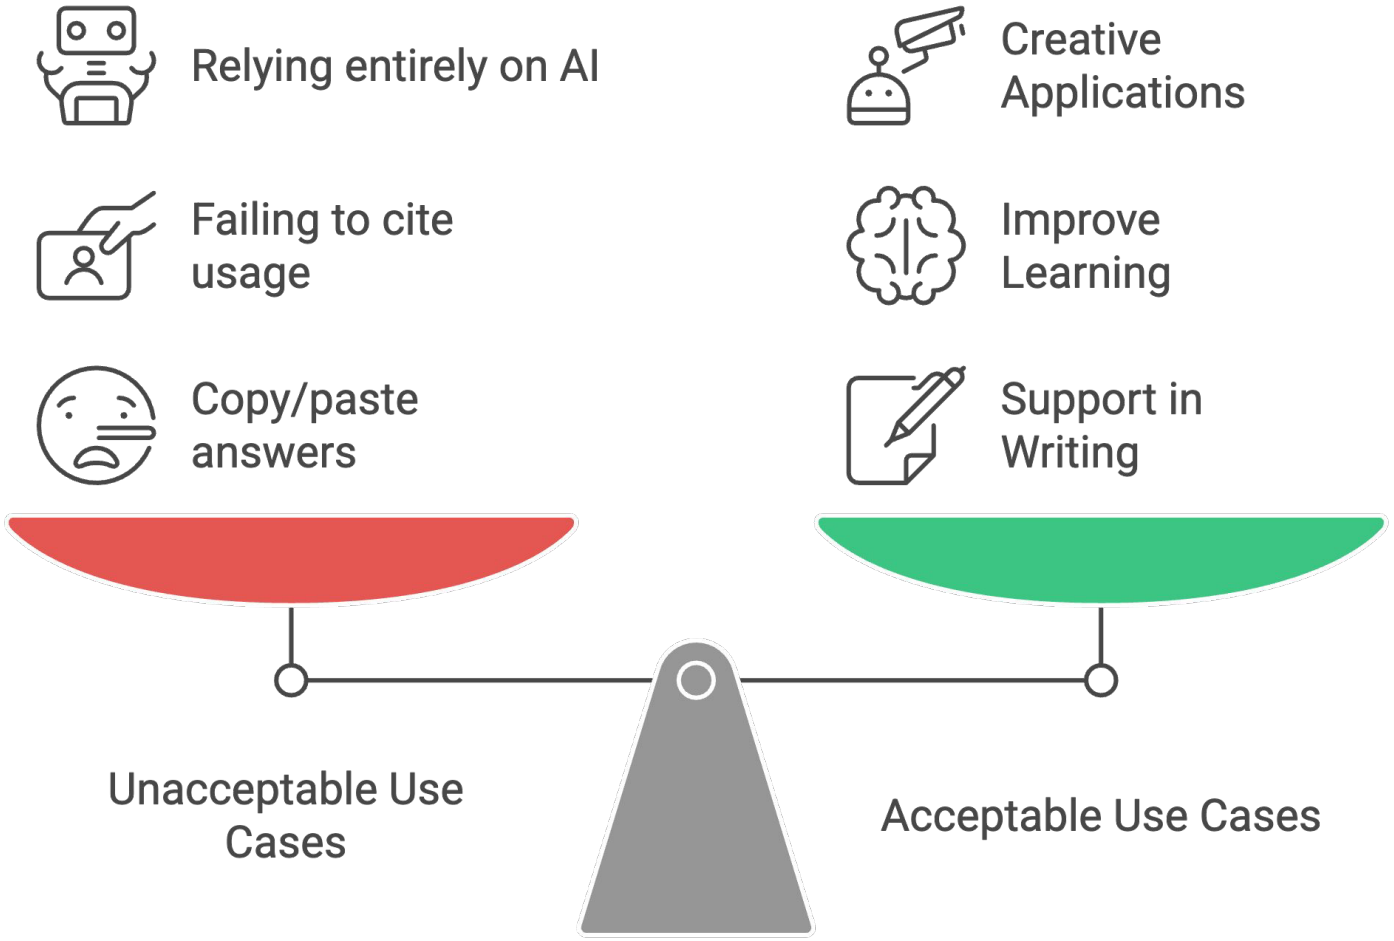

*Graphics made using NapkinAI

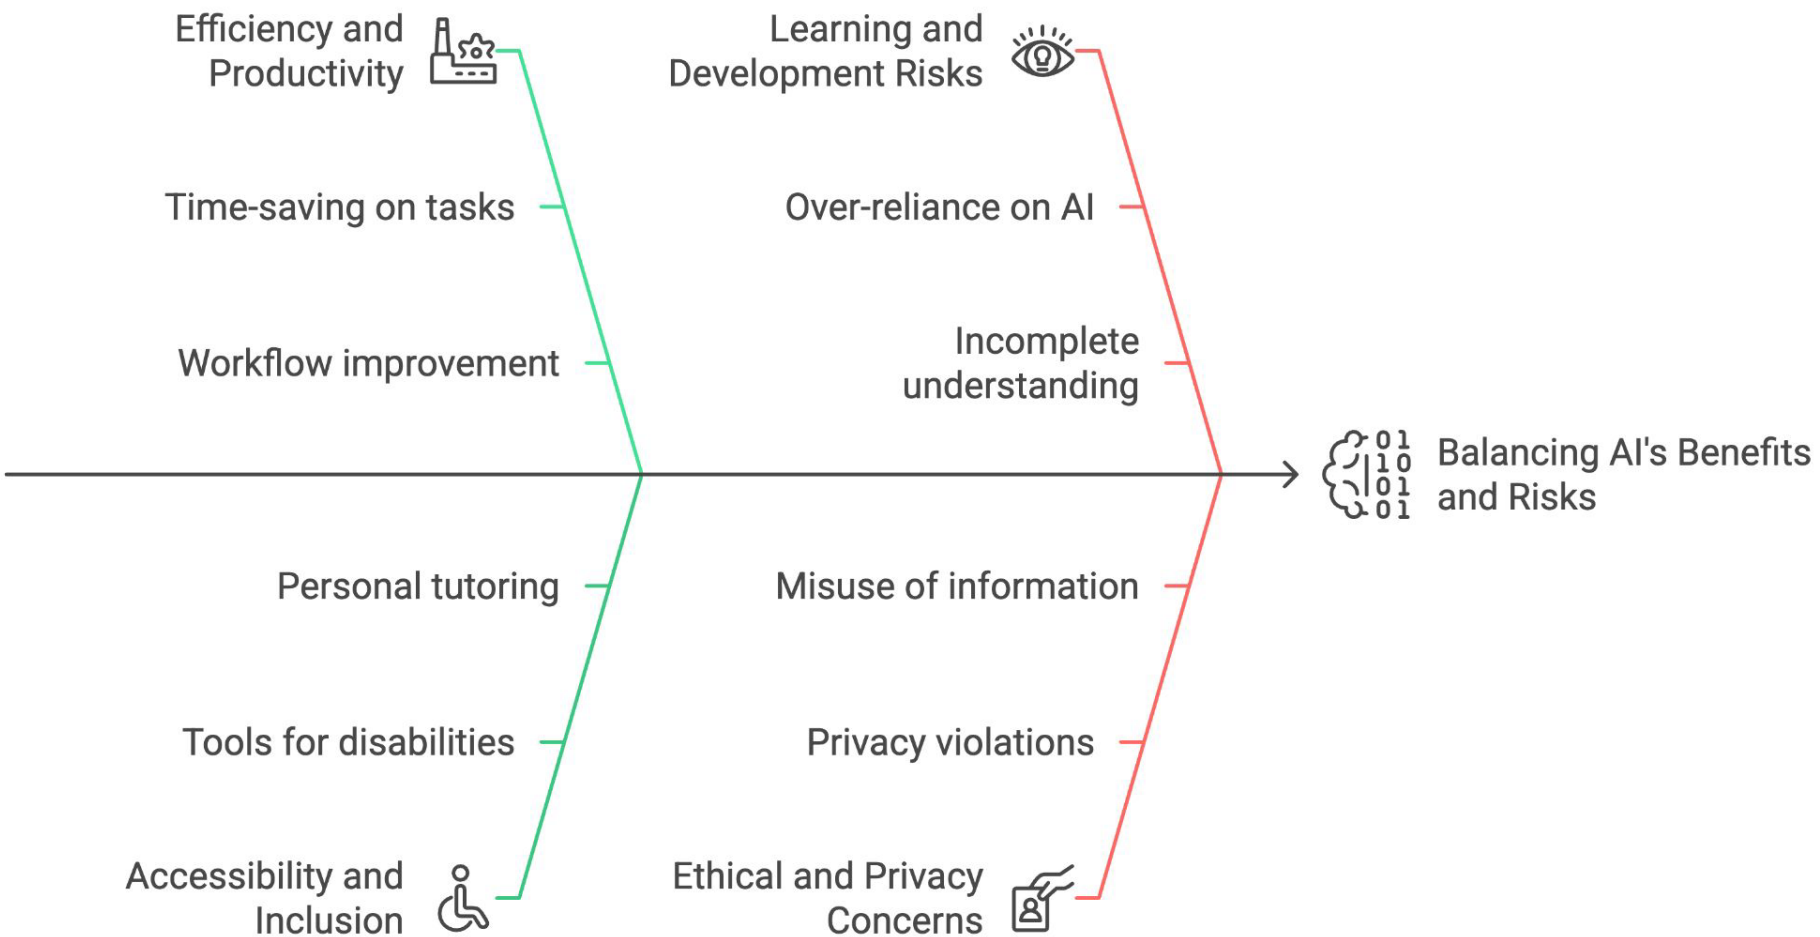

*Graphics made using NapkinAI

Is It Safe to Use GenAI for This Task?

Adapted from a graphic by Aleksandr Tuilkanov

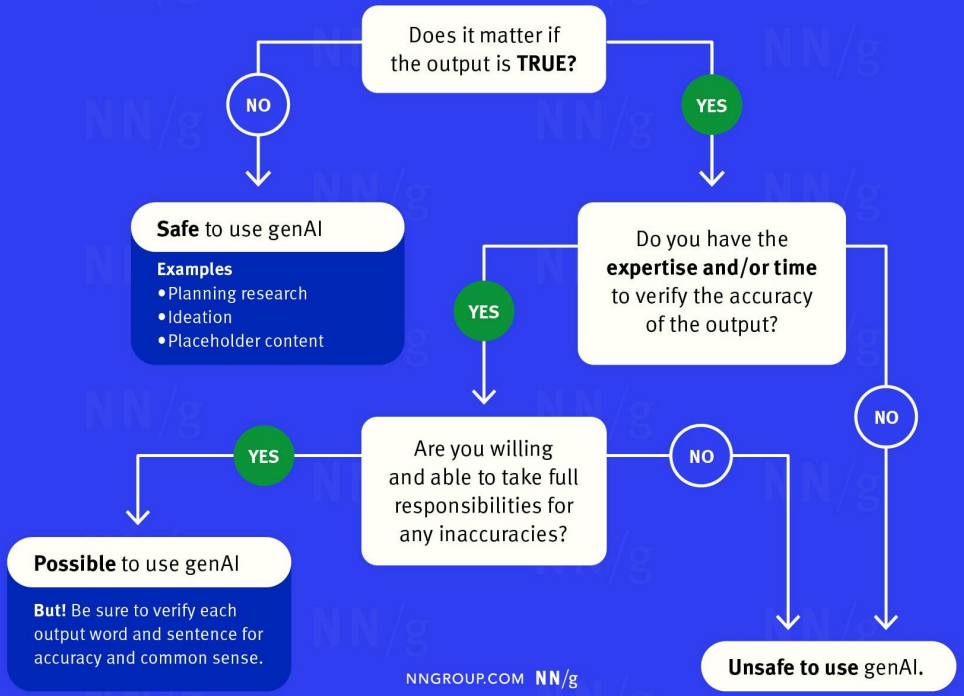

Chatbot

What does Chatbot mean?

A **chatbot** is a software or computer program simulating human conversation. It can be powered by various technologies, ranging from basic decision tree algorithms to advanced conversational AI, and can operate through text or voice interactions.

Intelligent Agents

What does it mean for a machine to be intelligent?

Turing Test

The test doesn't care whether a machine is intelligent or not; it cares whether a machine acts like. it's intelligent.

Turing Test

- "I propose to consider the question, "Can machines think?" The problem can be described in terms of the ‘imitation game’.

- "I believe that in about **fifty years' time** it will be possible to programme computers to make them play the imitation game so well that an average **interrogator will not have more than 70 percent chance of making the right identification after five minutes of questioning.**"– Alan Turing, 1950.

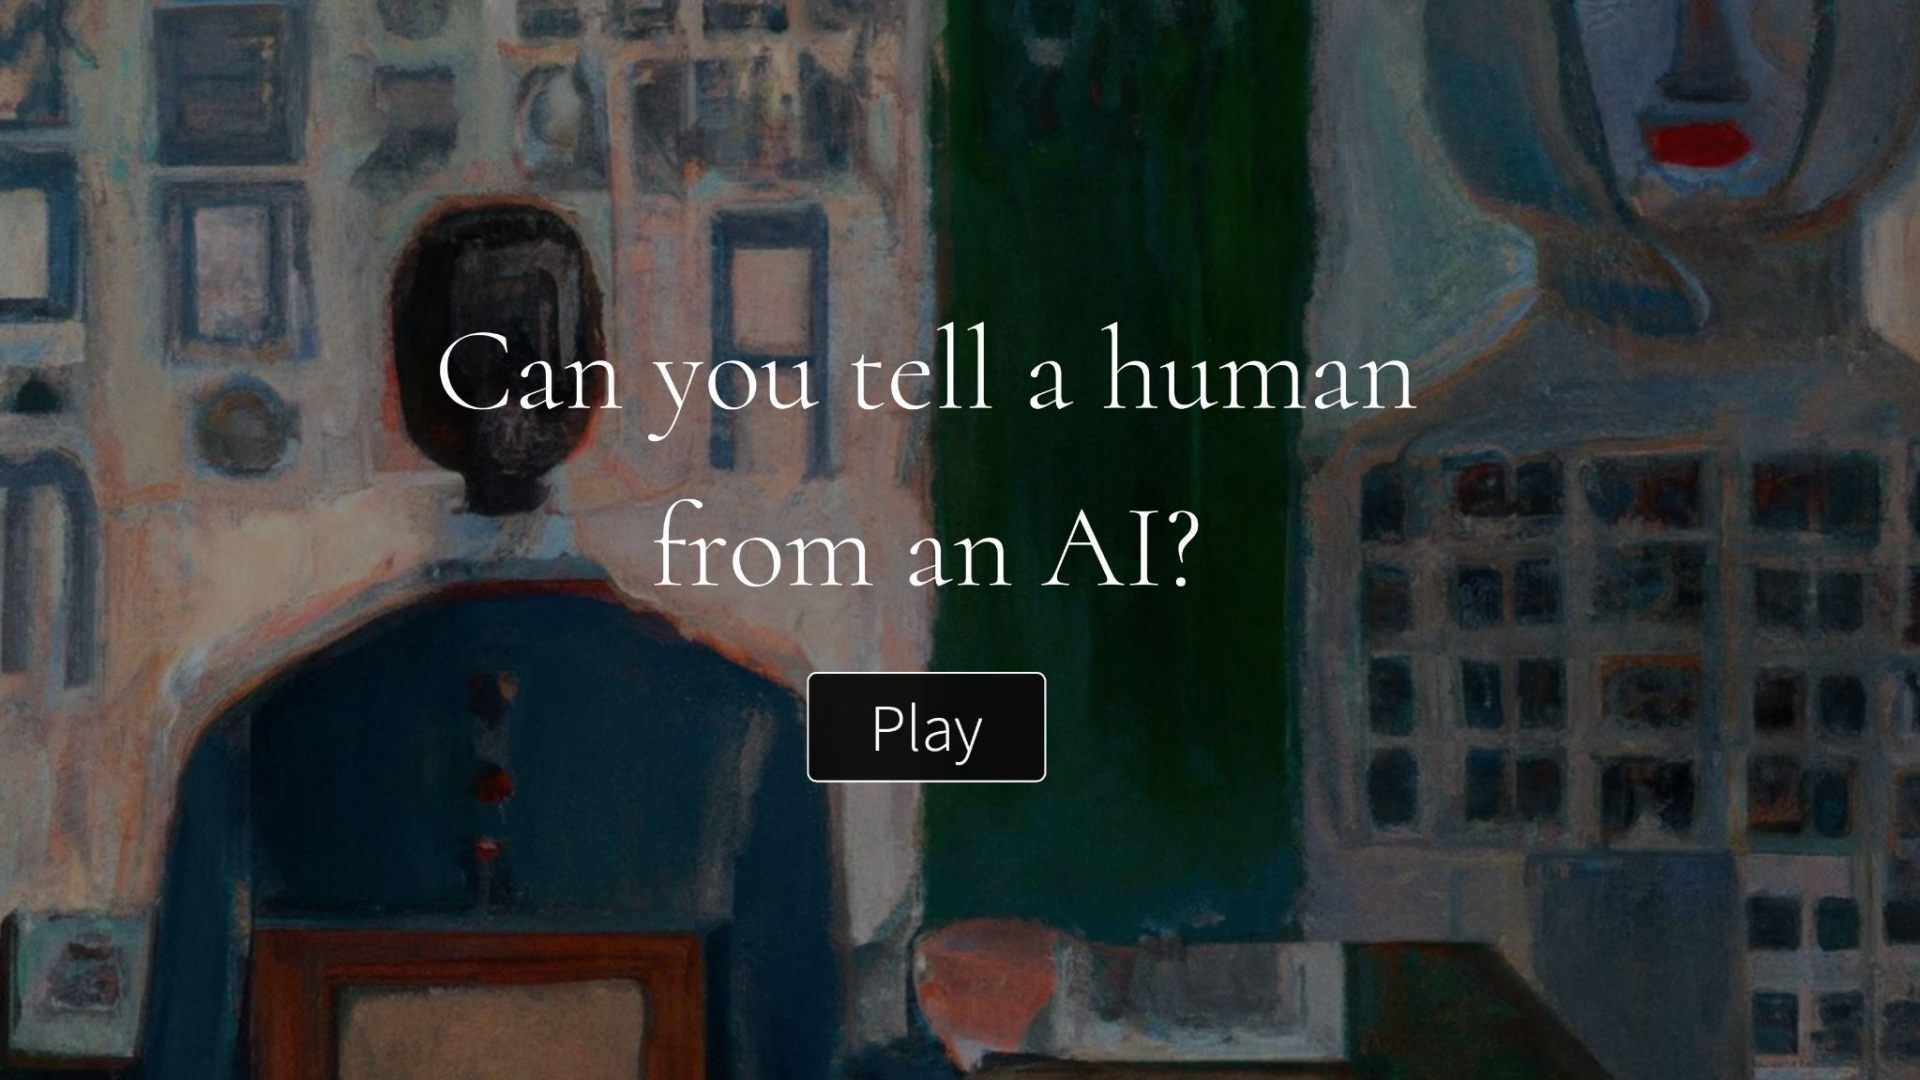
Can you tell a human
from an AI?
Play

- You will get randomly assigned to be either an **Interrogator** or a **Witness.**
- **Note this is a REAL study! Take it seriously…**

You are the

WITNESS

- Wait for the interrogator to ask the first question.
- Answer the interrogators questions and try to convince them that you are a human.
- The game will last either 5 or 10 minutes or until the interrogator makes a decision.
- After the game, you'll find out if the interrogator thought you were a human or an AI.

Okay

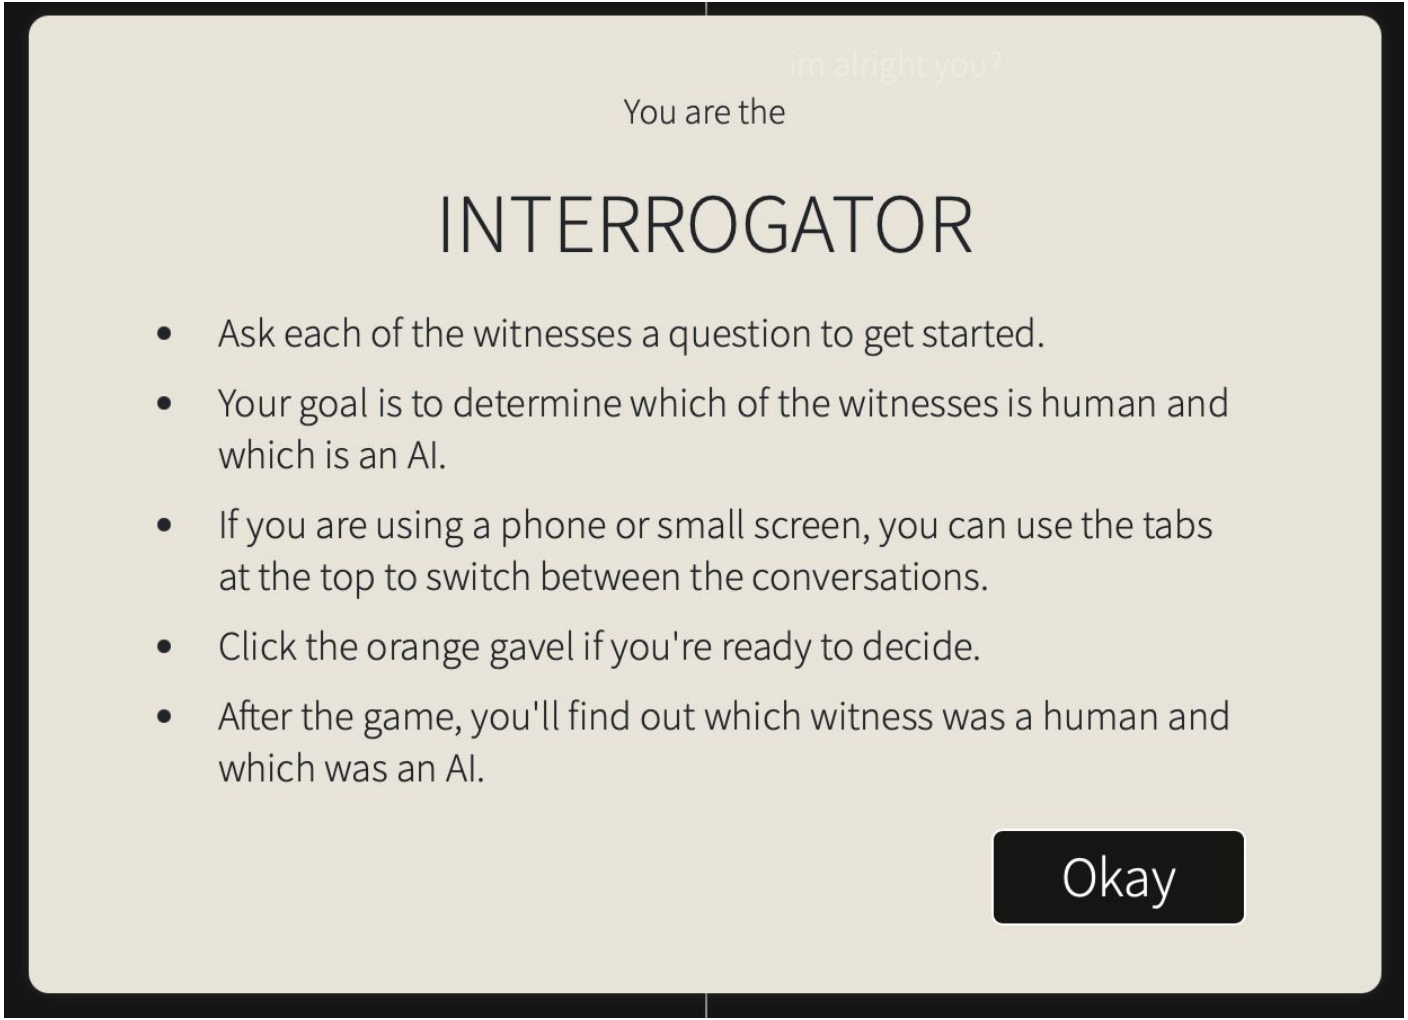
im alright you?
You are the
INTERROGATOR
Ask each of the witnesses a question to get started.
Your goal is to determine which of the witnesses is human and which is an AI.
If you are using a phone or small screen, you can use the tabs at the top to switch between the conversations.
Click the orange gavel if you're ready to decide.
After the game, you'll find out which witness was a human and which was an AI.
Okay

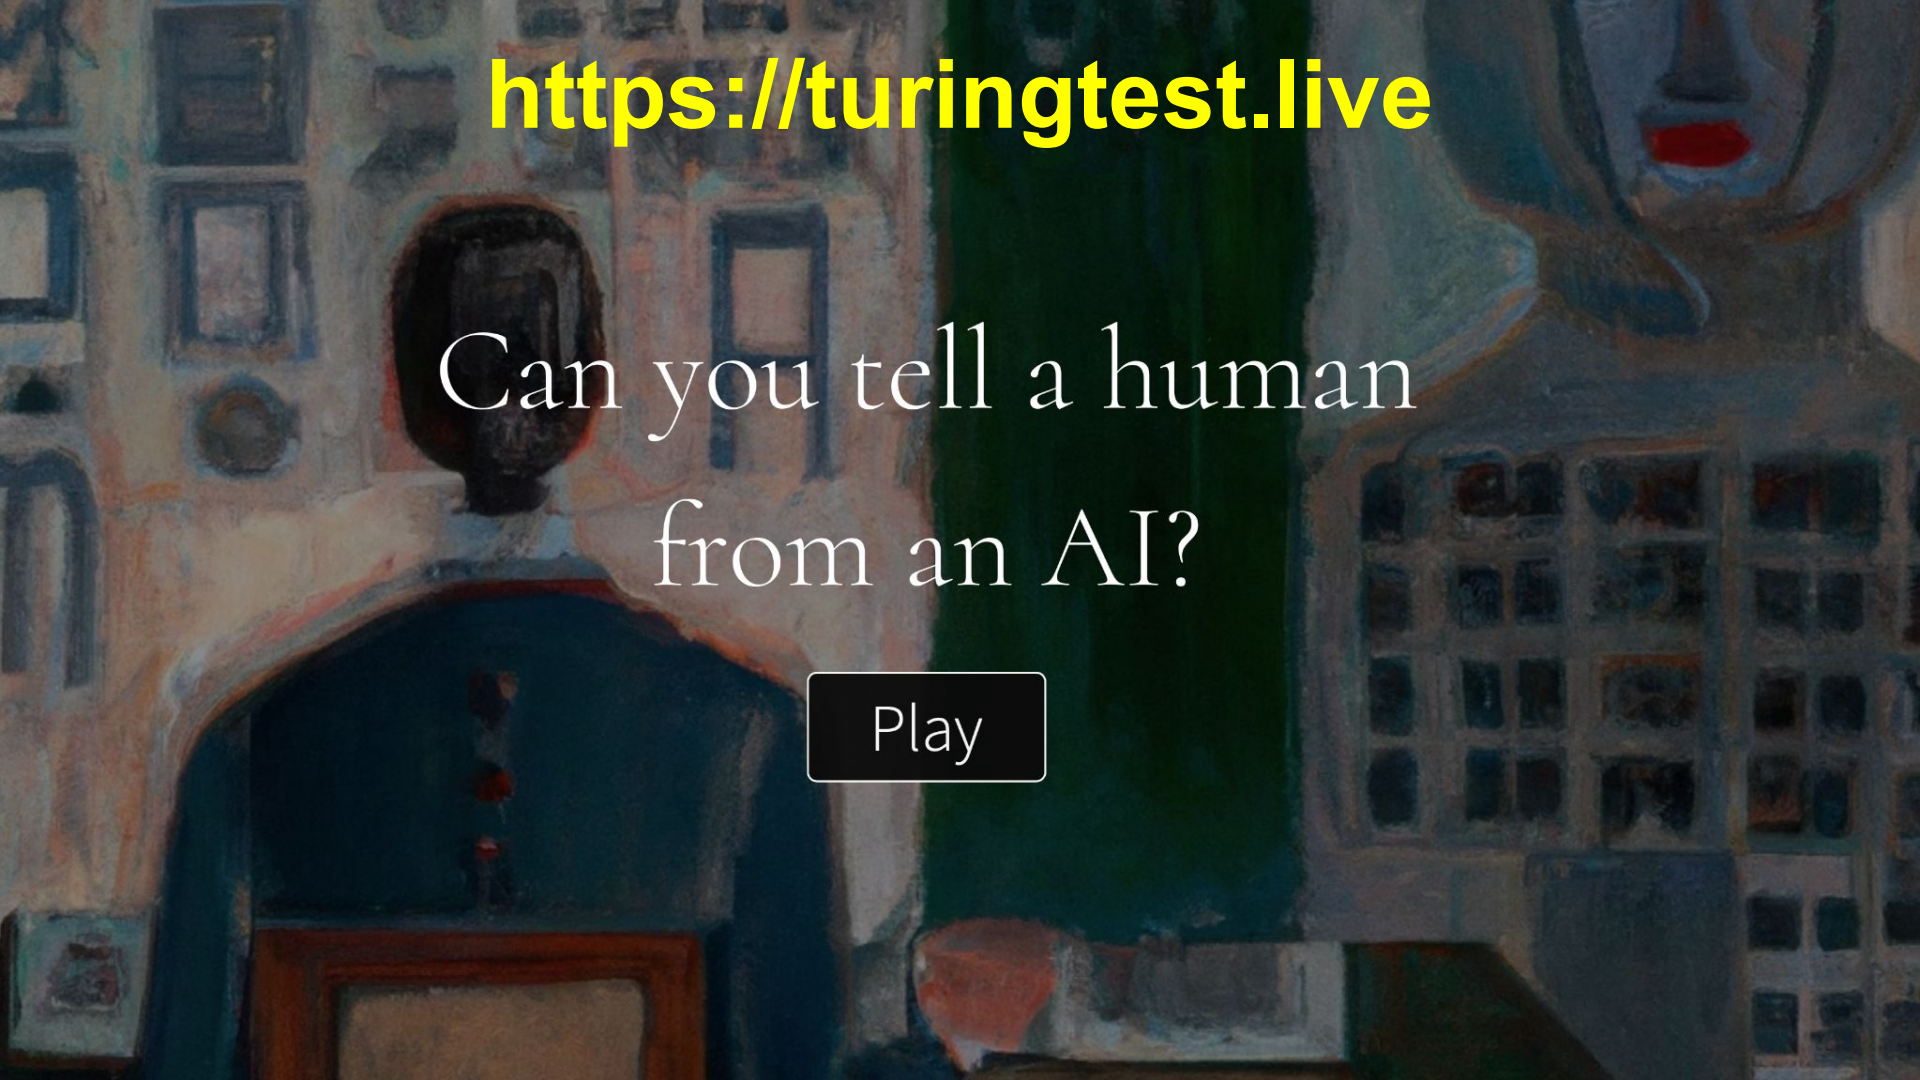
https://turingtest.live
Can you tell a human from an AI?
Play

YOUR STATS

	Interrogator	Witness	**Total**
Games	2	0	**2**
Success	50%	NA	**50%**
Words	193	NA	**193**
Time	4:33	NA	**4:33**

Foundations of AI

Foundations of AI/ChatGPT

- **Natural Language Processing (NLP)**: ChatGPT's primary function is to process and generate human language, which is the core of NLP. It uses advanced NLP techniques to understand context, generate responses, and maintain coherent conversations.
- **Machine Learning (ML)**: It utilizes the transformer model, a deep learning technique, to train on extensive text data. This enables ChatGPT to learn language patterns and context, thus generating coherent, context-aware responses.

Daniel Denntt on Turing Test

[Turing meant it] as a thought experiment that should convince people, that ... any computer that could pass this test, fair-and-square, of course it would be intelligent!" [Video]

Only concerned with whether a machine behaves intelligently

In addition, intelligence (circa Turing) is whether the machine can converse

How does NLP work?

NLP in a nutshell

- NLP draws on many disciplines: linguistics, cognitive science, psychology, logic, computer science, philosophy, engineering.
- **Traditional approach**: Long list of rules for processing language, formulated by people and programmed into computers
- **Modern approach**: Machines learn from text examples using artificial neural networks and similar approaches. Statistical methods allows to compare different interpretations

Traditional NLP Steps

1. **Recognize speech** (typically chatbots receive ASCII versions of the questions)

2. **Syntax analysis,or parsing**: inferring parts of speech and sentence structure, using a lexicon and grammar

3. **Semantic analysis**: inferring meaning using syntax and semantic rules

4. **Pragmatics**: inferring meaning from contextual information

Limitations of traditional NLP

- Natural language is structurally **ambiguous**, so parsing alone cannot lead to understanding.
- Synonyms for words can’t be used interchangeably in every context, e.g., “*minister of agriculture*” isn't “*priest of farming*.”
- Natural languages have many exceptions to grammatical rules; there’s no agreed-upon grammar for all uses of a language.

Parsing is the basis for programming

- A computer has to "**understand**" programs in order to execute them
- Programming languages are designed so that they can be parsed **unambiguously**
- A grammar specifies all the possible programs that can be written in a language
- **Designing programming** languages (and their grammars) is a fun and important part of computer science

NLP: Semantic Analysis Processes

- **Word Sense Disambiguation**: Identifying the correct meaning of words with multiple interpretations in a given context.
- **Semantic Role Labeling**: Assigning roles (e.g., agent, patient) to elements in a sentence to understand their relationships.
- **Named Entity Recognition**: Identifying and categorizing named entities like names, locations, and organizations in text.
- **Semantic Parsing**: Converting natural language into a formal representation of meaning, aiding in understanding user queries or commands.
- **Sentiment Analysis**: Determining the sentiment expressed in text (positive, negative, or neutral).
- **Semantic Similarity**: Measuring the likeness or relatedness of text based on its meaning.

Tradition vs. Modern

- **Parsing**: While traditional parsing involves breaking down and analyzing the structure of language (like sentence structure and grammar), ChatGPT's transformer architecture processes text through a series of layers that capture different aspects of language, including some structural elements.
- **Semantic Analysis**: ChatGPT performs semantic analysis through its deep learning model, which understands the meanings and relationships of words and phrases in context. This is essential for generating coherent and contextually relevant responses.

(1) Input: Sentence: “The rat ate cheese”

ASCII TABLE

Decimal	Hex	Char	Decimal	Hex	Char	Decimal	Hex	Char	Decimal	Hex	Char
0	0	[NULL]	32	20	[SPACE]	64	40	@	96	60	`
1	1	[START OF HEADING]	33	21	!	65	41	A	97	61	a
2	2	[START OF TEXT]	34	22	"	66	42	B	98	62	b
3	3	[END OF TEXT]	35	23	#	67	43	C	99	63	c
4	4	[END OF TRANSMISSION]	36	24	$	68	44	D	100	64	d
5	5	[ENQUIRY]	37	25	%	69	45	E	101	65	e
6	6	[ACKNOWLEDGE]	38	26	&	70	46	F	102	66	f
7	7	[BELL]	39	27	'	71	47	G	103	67	g
8	8	[BACKSPACE]	40	28	(	72	48	H	104	68	h
9	9	[HORIZONTAL TAB]	41	29	)	73	49	I	105	69	i
10	A	[LINE FEED]	42	2A	*	74	4A	J	106	6A	j
11	B	[VERTICAL TAB]	43	2B	+	75	4B	K	107	6B	k
12	C	[FORM FEED]	44	2C	,	76	4C	L	108	6C	l
13	D	[CARRIAGE RETURN]	45	2D	-	77	4D	M	109	6D	m
14	E	[SHIFT OUT]	46	2E	.	78	4E	N	110	6E	n
15	F	[SHIFT IN]	47	2F	/	79	4F	O	111	6F	o
16	10	[DATA LINK ESCAPE]	48	30	0	80	50	P	112	70	p
17	11	[DEVICE CONTROL 1]	49	31	1	81	51	Q	113	71	q
18	12	[DEVICE CONTROL 2]	50	32	2	82	52	R	114	72	r
19	13	[DEVICE CONTROL 3]	51	33	3	83	53	S	115	73	s
20	14	[DEVICE CONTROL 4]	52	34	4	84	54	T	116	74	t
21	15	[NEGATIVE ACKNOWLEDGE]	53	35	5	85	55	U	117	75	u
22	16	[SYNCHRONOUS IDLE]	54	36	6	86	56	V	118	76	v
23	17	[END OF TRANS. BLOCK]	55	37	7	87	57	W	119	77	w
24	18	[CANCEL]	56	38	8	88	58	X	120	78	x
25	19	[END OF MEDIUM]	57	39	9	89	59	Y	121	79	y
26	1A	[SUBSTITUTE]	58	3A	:	90	5A	Z	122	7A	z
27	1B	[ESCAPE]	59	3B	;	91	5B	[	123	7B	{
28	1C	[FILE SEPARATOR]	60	3C	<	92	5C	\	124	7C	|
29	1D	[GROUP SEPARATOR]	61	3D	=	93	5D	]	125	7D	}
30	1E	[RECORD SEPARATOR]	62	3E	>	94	5E	^	126	7E	~
31	1F	[UNIT SEPARATOR]	63	3F	?	95	5F	_	127	7F	[DEL]

(1) Input: Sentence: “The rat ate cheese”

ASCII TABLE

Decimal	Hex	Char	Decimal	Hex	Char	Decimal	Hex	Char	Decimal	Hex	Char
0	0	[NULL]	32	20	[SPACE]	64	40	@	96	60	`
1	1	[START OF HEADING]	33	21	!	65	41	A	97	61	a
2	2	[START OF TEXT]	34	22	"	66	42	B	98	62	b
3	3	[END OF TEXT]	35	23	#	67	43	C	99	63	c
4	4	[END OF TRANSMISSION]	36	24	$	68	44	D	100	64	d
5	5	[ENQUIRY]	37	25	%	69	45	E	101	65	e
6	6	[ACKNOWLEDGE]	38	26	&	70	46	F	102	66	f
7	7	[BELL]	39	27	'	71	47	G	103	67	g
8	8	[BACKSPACE]	40	28	(	72	48	H	104	68	h
9	9	[HORIZONTAL TAB]	41	29	)	73	49	I	105	69	i
10	A	[LINE FEED]	42	2A	*	74	4A	J	106	6A	j
11	B	[VERTICAL TAB]	43	2B	+	75	4B	K	107	6B	k
12	C	[FORM FEED]	44	2C	,	76	4C	L	108	6C	l
13	D	[CARRIAGE RETURN]	45	2D	-	77	4D	M	109	6D	m
14	E	[SHIFT OUT]	46	2E	.	78	4E	N	110	6E	n
15	F	[SHIFT IN]	47	2F	/	79	4F	O	111	6F	o
16	10	[DATA LINK ESCAPE]	48	30	0	80	50	P	112	70	p
17	11	[DEVICE CONTROL 1]	49	31	1	81	51	Q	113	71	q
18	12	[DEVICE CONTROL 2]	50	32	2	82	52	R	114	72	r
19	13	[DEVICE CONTROL 3]	51	33	3	83	53	S	115	73	s
20	14	[DEVICE CONTROL 4]	52	34	4	84	54	T	116	74	t
21	15	[NEGATIVE ACKNOWLEDGE]	53	35	5	85	55	U	117	75	u
22	16	[SYNCHRONOUS IDLE]	54	36	6	86	56	V	118	76	v
23	17	[END OF TRANS. BLOCK]	55	37	7	87	57	W	119	77	w
24	18	[CANCEL]	56	38	8	88	58	X	120	78	x
25	19	[END OF MEDIUM]	57	39	9	89	59	Y	121	79	y
26	1A	[SUBSTITUTE]	58	3A	:	90	5A	Z	122	7A	z
27	1B	[ESCAPE]	59	3B	;	91	5B	[	123	7B	{
28	1C	[FILE SEPARATOR]	60	3C	<	92	5C	\	124	7C	|
29	1D	[GROUP SEPARATOR]	61	3D	=	93	5D	]	125	7D	}
30	1E	[RECORD SEPARATOR]	62	3E	>	94	5E	^	126	7E	~
31	1F	[UNIT SEPARATOR]	63	3F	?	95	5F	_	127	7F	[DEL]

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

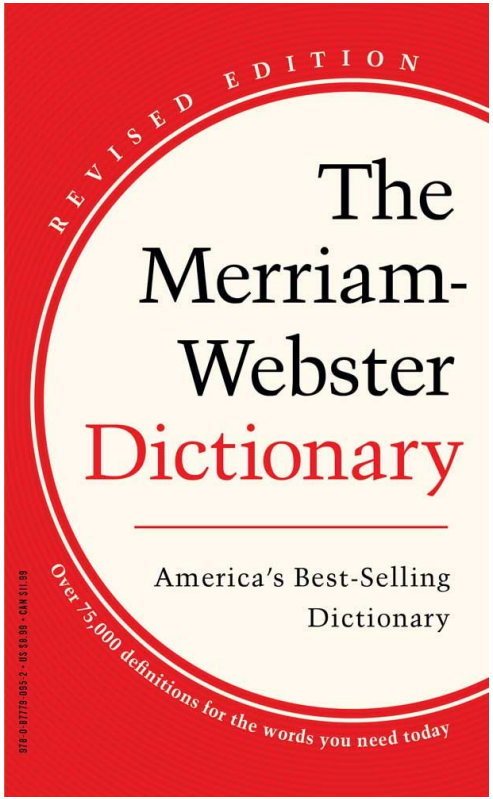

(1) **Input:** Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase

“The dog barked.”
- Noun Phrase (NP): “The dog”
- Verb Phrase (VP): “barked”

“A cat is sleeping.”
- Noun Phrase (NP): “A cat”
- Verb Phrase (VP): “is sleeping”

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase

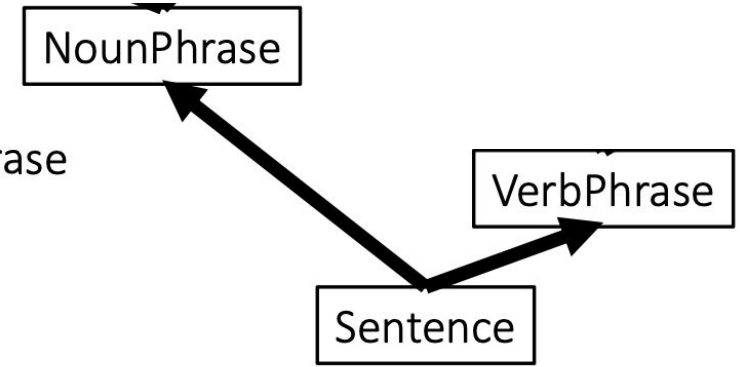

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase

VerbPhrase → Verb, NounPhrase

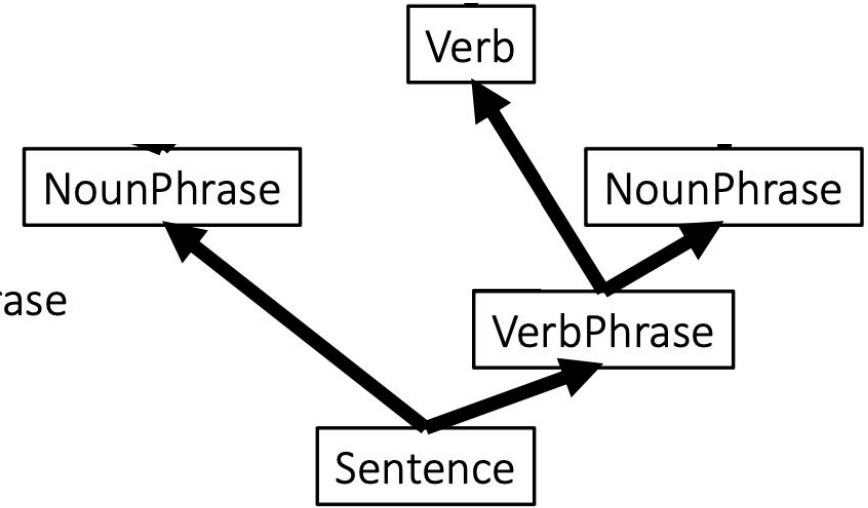

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase
VerbPhrase → Verb, NounPhrase
NounPhrase → Article, Noun

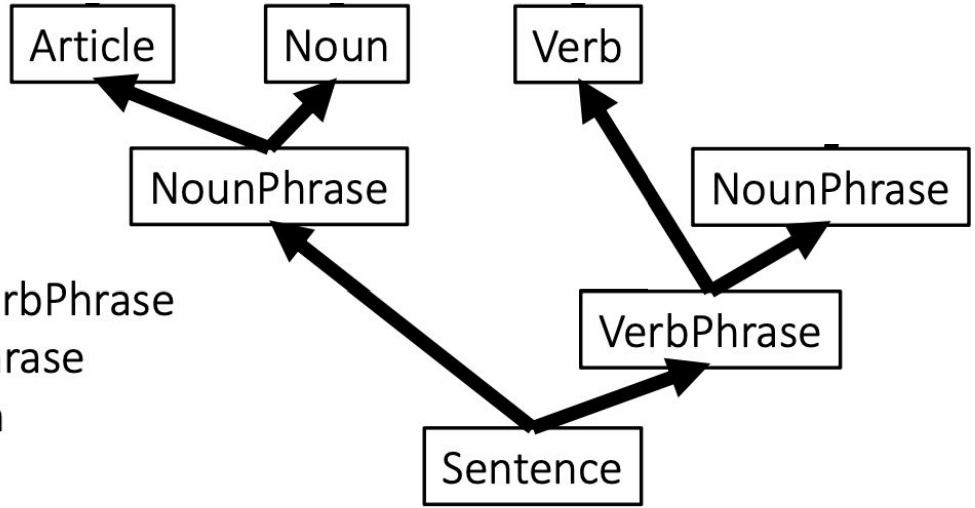

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase
VerbPhrase → Verb, NounPhrase
NounPhrase → Article, Noun
NounPhrase → Noun

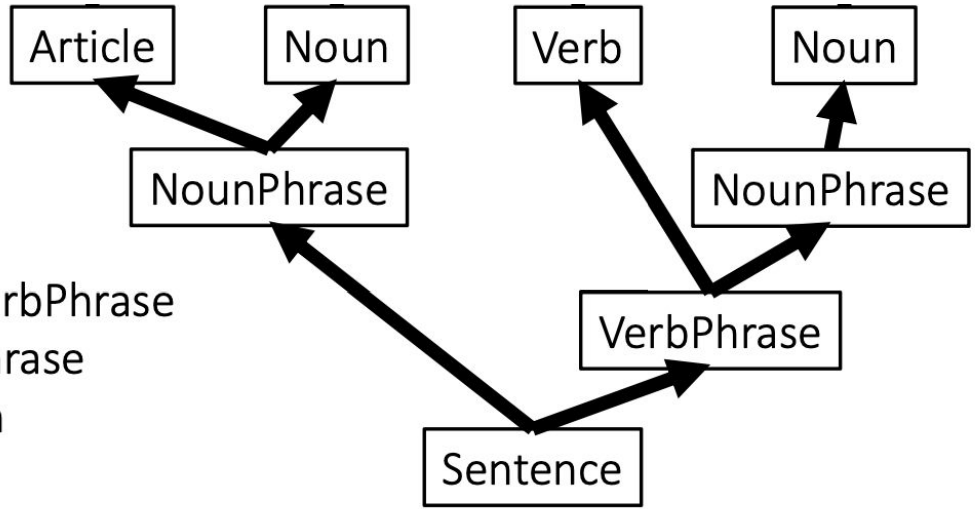

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase
VerbPhrase → Verb, NounPhrase
NounPhrase → Article, Noun
NounPhrase → Noun

(4) Output: A parse tree:

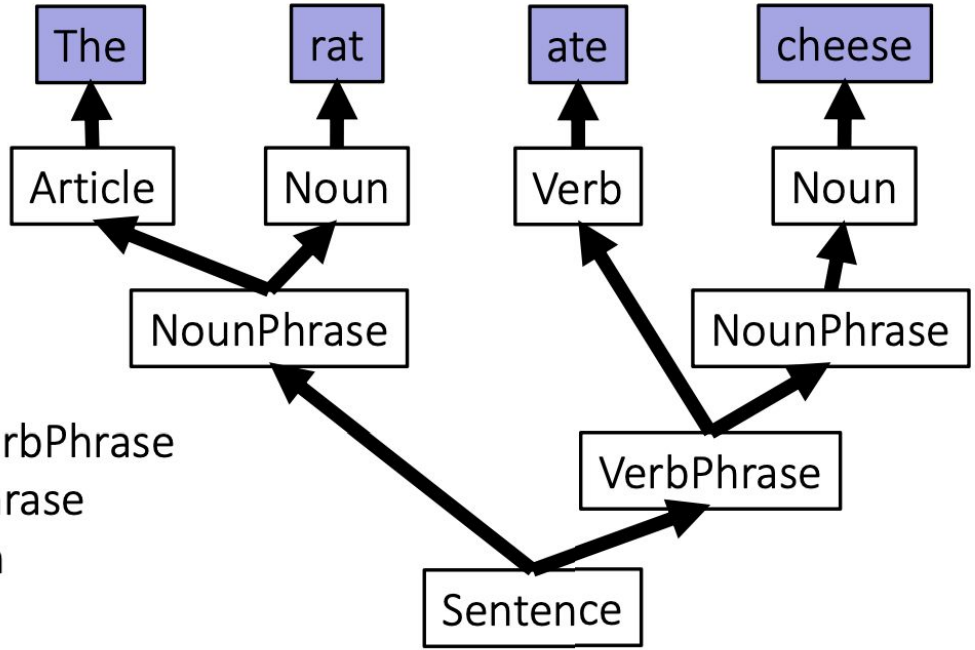

(1) Input: Sentence: “The rat ate cheese”

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase
VerbPhrase → Verb, NounPhrase
NounPhrase → Article, Noun
NounPhrase → Noun

(4) Output: A parse tree:

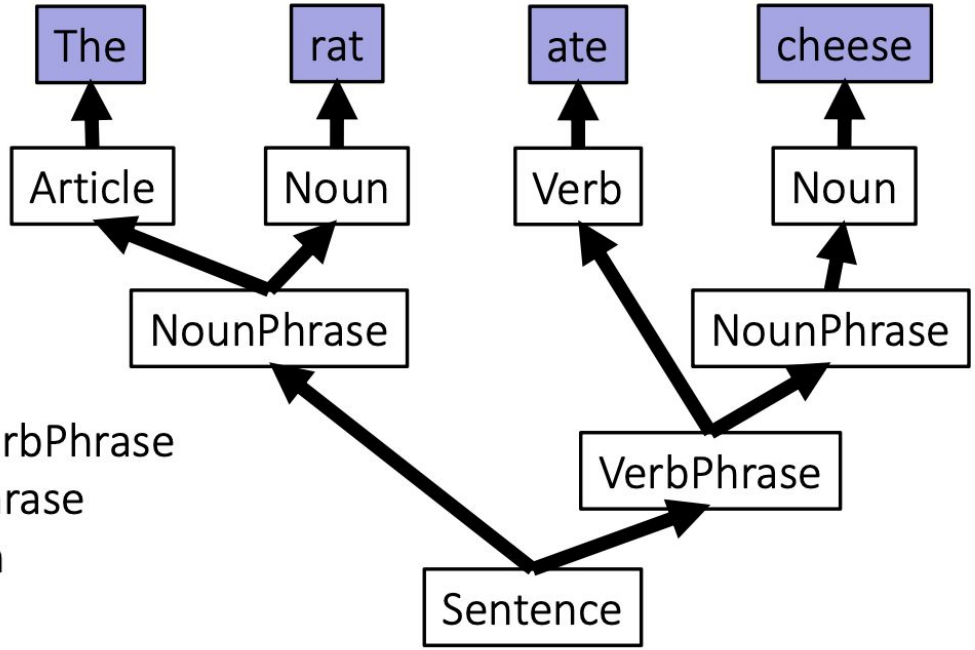

Q: Which skill is the most applicable in step 3 (Grammar)?

iClicker

A. Abstraction

B. Decomposition

C. Synthesis

D. Simulation

E. Modelling

(1) Input: Sentence: "The rat ate cheese"

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase
VerbPhrase → Verb, NounPhrase
NounPhrase → Article, Noun
NounPhrase → Noun

(4) Output: A parse tree:

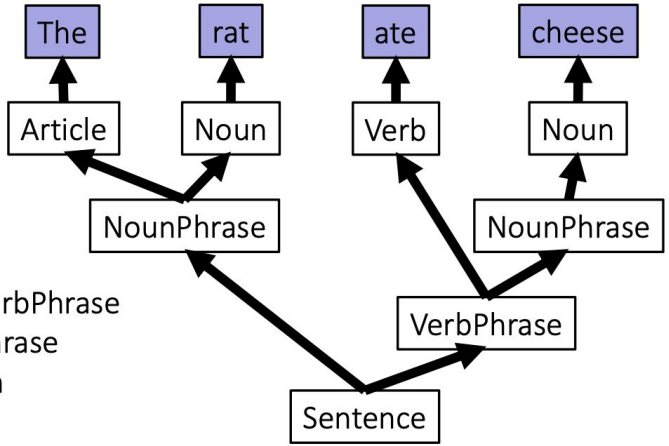

Q: Which skill is the most applicable in step 3 (Grammer)?

iClicker

A. Abstraction

B. Decomposition

C. Synthesis

D. Simulation

E. Modelling

(1) Input: Sentence: "The rat ate cheese"

(2) Lexicon:

rat	Noun
cheese	Noun
ate	Verb
The	Article

(3) Grammar:

Sentence → NounPhrase, VerbPhrase
VerbPhrase → Verb, NounPhrase
NounPhrase → Article, Noun
NounPhrase → Noun

(4) Output: A parse tree:

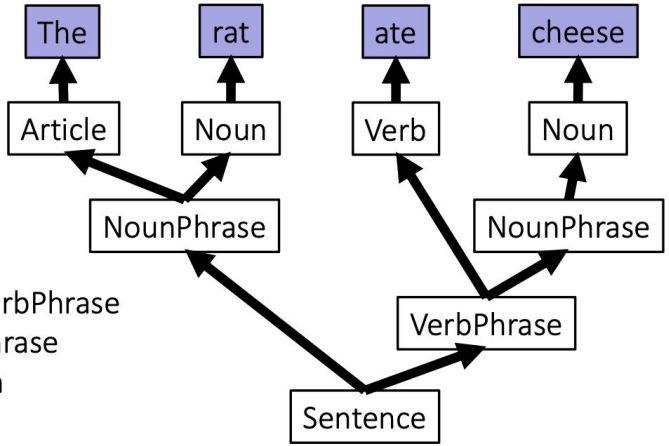

Q: Why is the Turing Test significant in AI?

iClicker

A. It measures the speed of algorithms.

B. It assesses whether machines can behave intelligently like humans.

C. It evaluates the hardware capacity of AI systems.

D. It rates the efficiency of chatbots.

E. It assesses whether machines are smarter than humans.

Q: Why is the Turing Test significant in AI?

iClicker

A. It measures the speed of algorithms.

B. It assesses whether machines can behave intelligently like humans.

C. It evaluates the hardware capacity of AI systems.

D. It rates the efficiency of chatbots.

E. It assesses whether machines are smarter than humans.

Trolley Problem

Trolley Problem

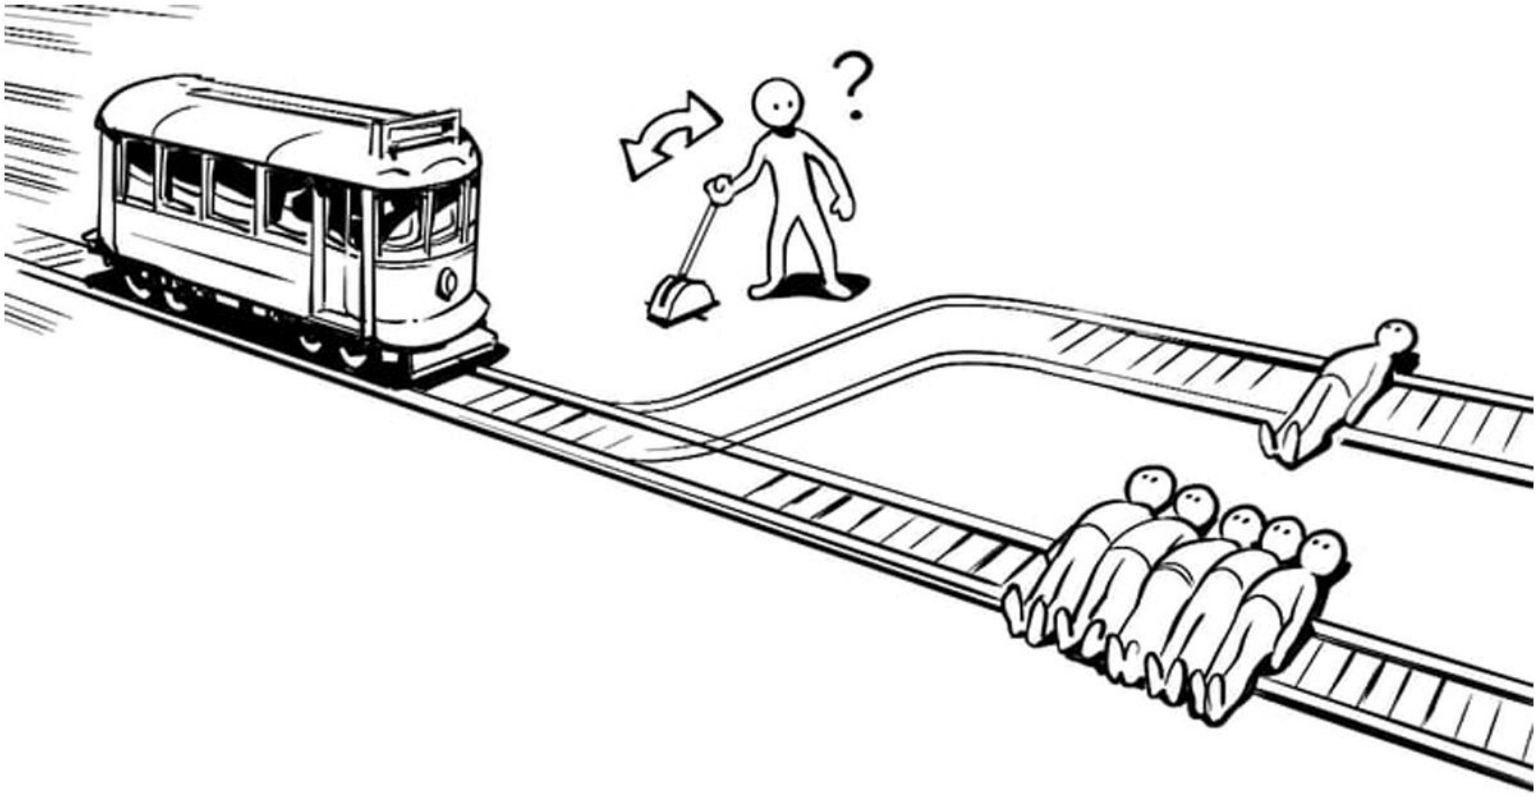

Trolley Problem

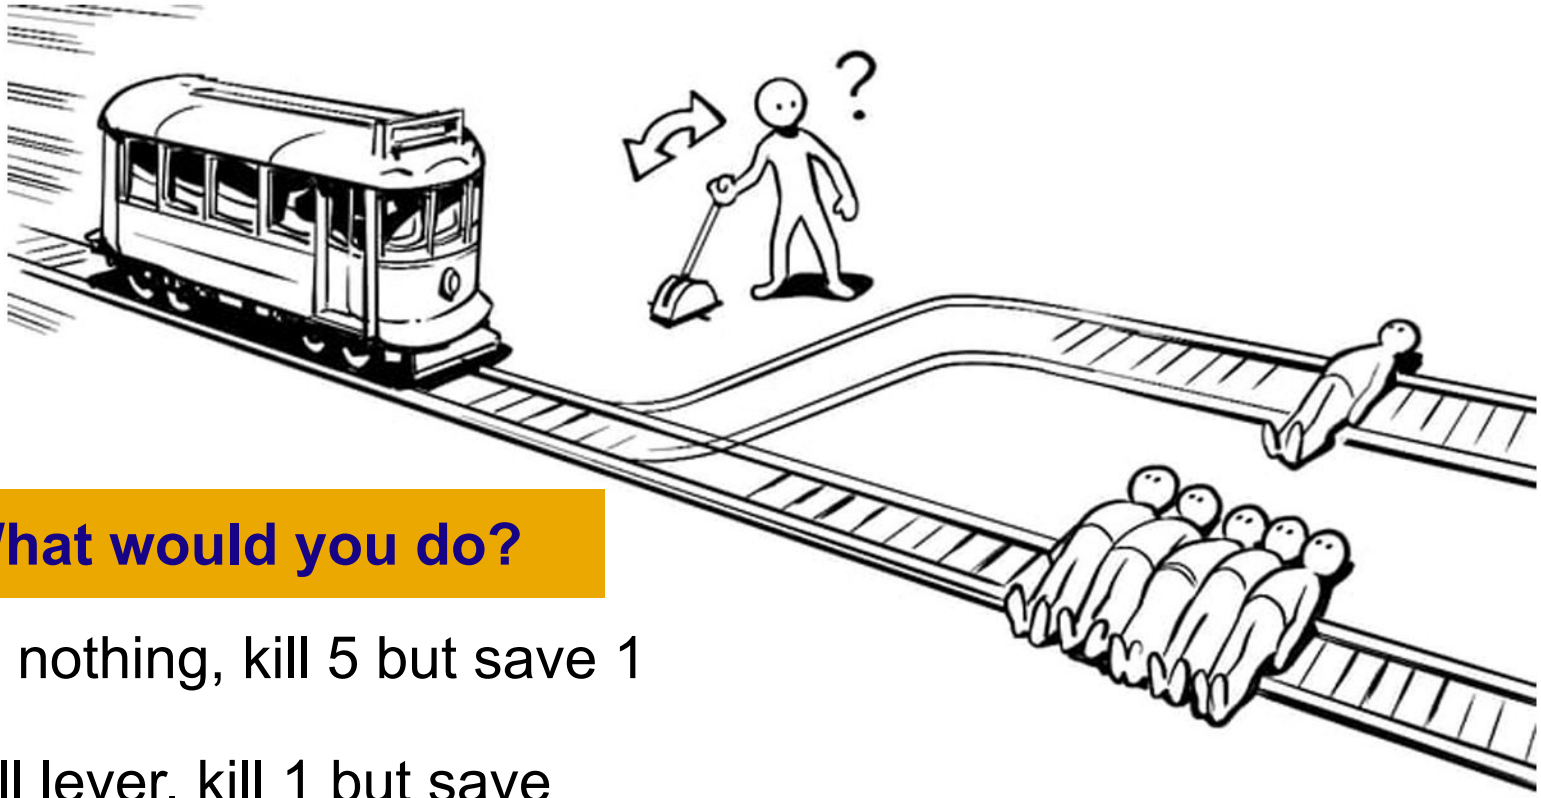

iClicker

Q: What would you do?

A. Do nothing, kill 5 but save 1

B. Pull lever, kill 1 but save

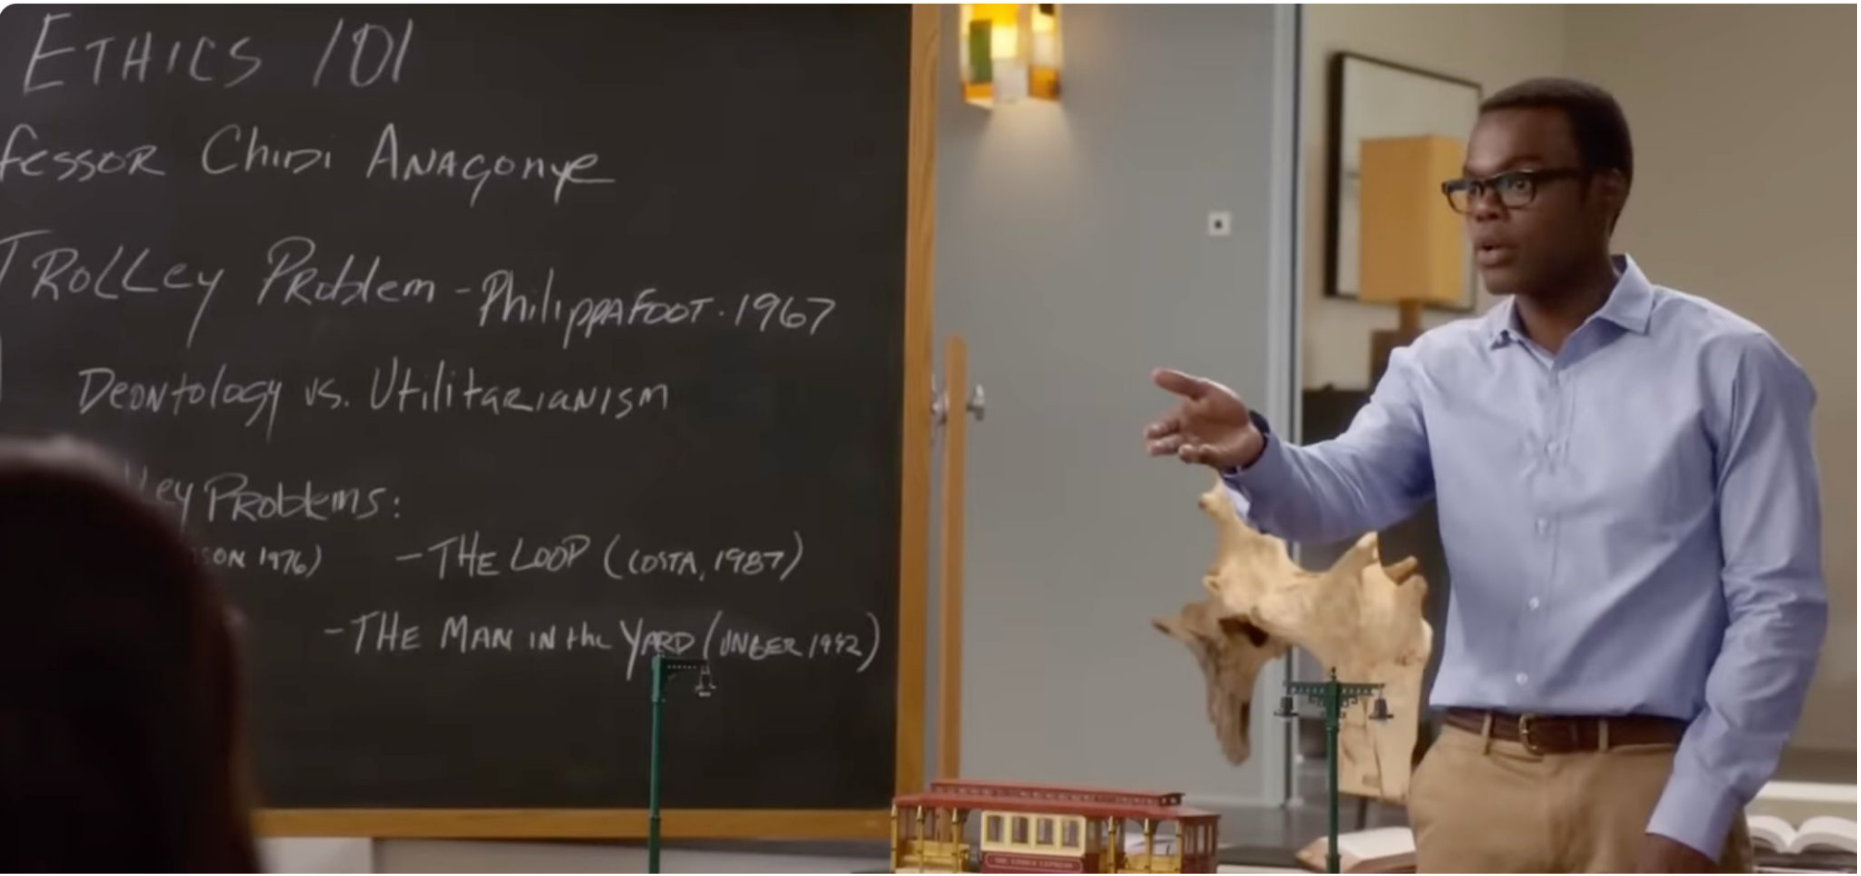
ETHICS 101
fessor Chidi Anagonye
Trolley Problem - Philippa Foot · 1967
Deontology vs. Utilitarianism
ley Problems:
-THE LOOP (COSTA, 1987)
-THE MAN IN the YARD (UNGER 1992)

Trolley Problem

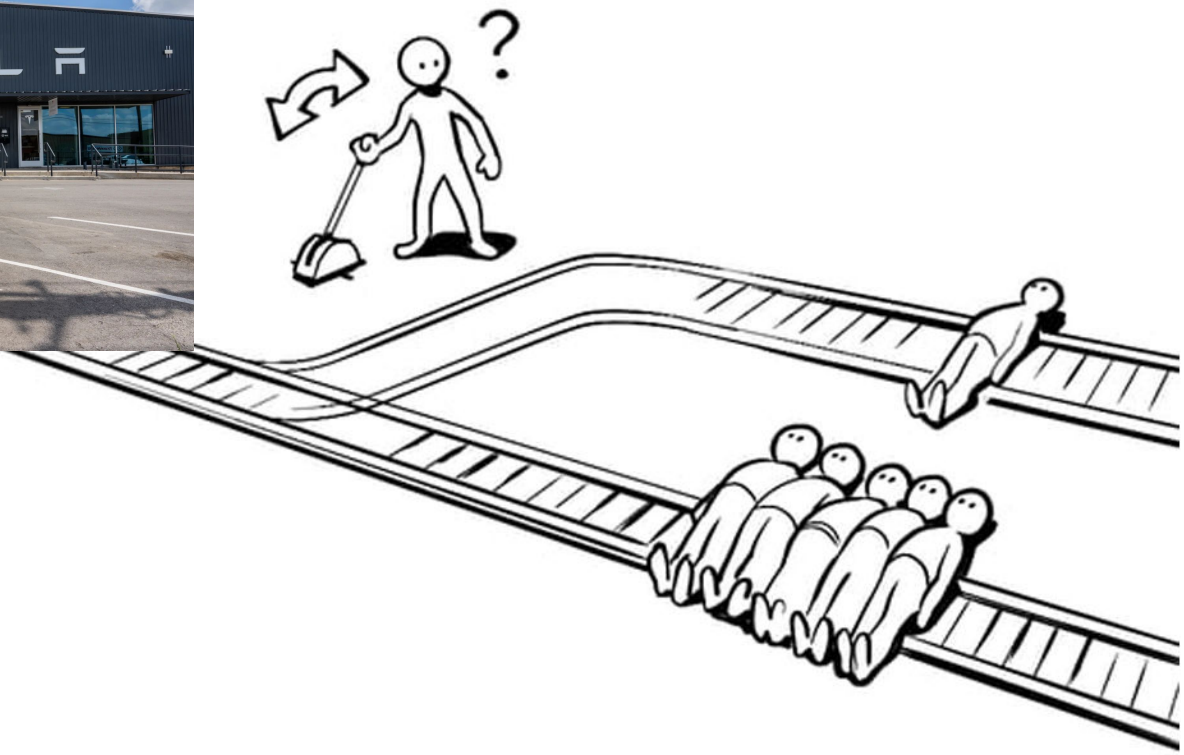

Trolley Problem

iClicker

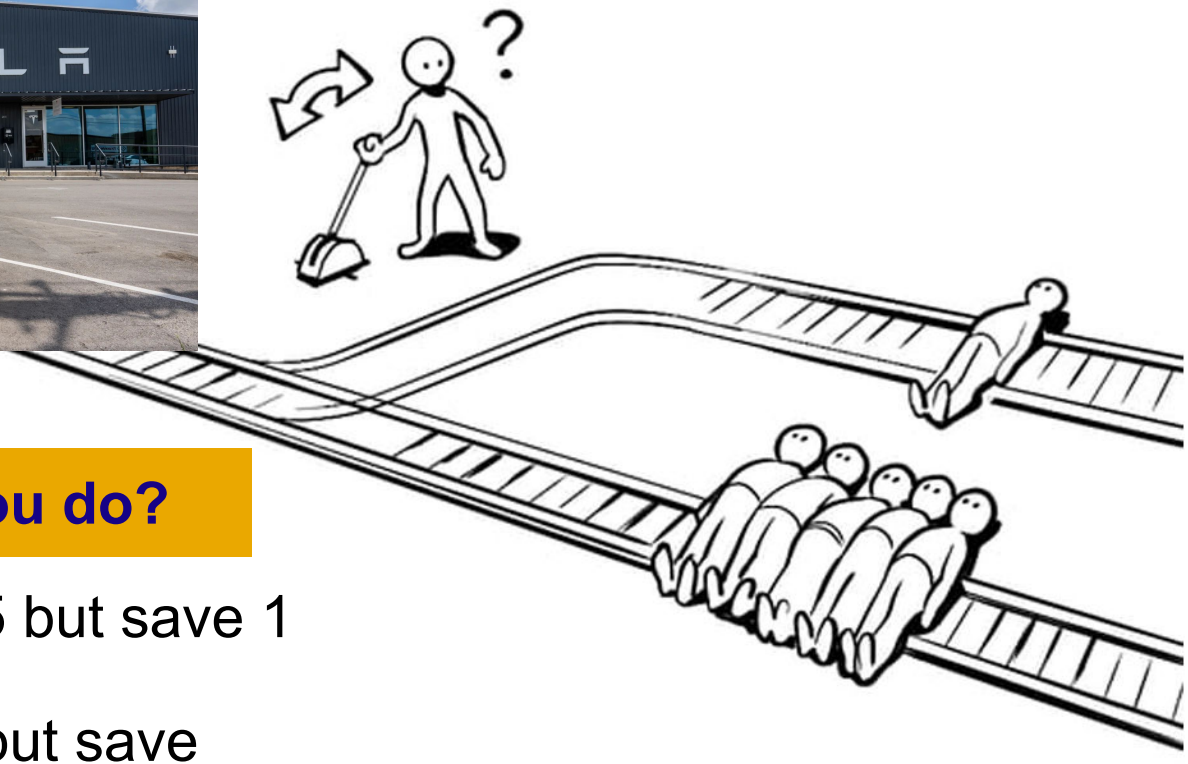

Q: What would you do?

A. Do nothing, kill 5 but save 1

B. Pull lever, kill 1 but save

Class Activity

Class Activity: Parsing!

Parse “time flies like an arrow”. Make a parse tree structure and write down your algorithm.

http://tiny.cc/100-W2C

Grammar

Sentence → NounPhrase, VerbPhrase
NounPhrase → Article, Adjective, Noun
NounPhrase → Article, Noun
NounPhrase → Noun, Noun
NounPhrase → Noun
VerbPhrase → Verb, Article, NounPhrase
VerbPhrase → Verb, NounPhrase
VerbPhrase → Verb, PrepPhrase
PrepPhrase → Preposition, NounPhrase

Lexicon

Word	Category
a / an	article
arrow	noun
banana	noun
flies	noun
flies	verb
fruit	noun
fruit	adjective
like	preposition
like	verb
time	noun
time	verb

Q: What parts of speech did you come up with?

iClicker

A. Noun Noun Verb Article Noun

B. Verb Noun Verb Article Noun

C. Noun Verb Preposition Article Noun

D. Noun Verb Verb Article Noun

E. None of the above

Q: What parts of speech did you come up with?

iClicker

A. Noun Noun Verb Article Noun

B. Verb Noun Verb Article Noun

C. Noun Verb Preposition Article Noun

D. Noun Verb Verb Article Noun

E. None of the above

Wrap up