

# Creating an Artificial Intelligence to Play Trivia Games

---

Pedro Rodriguez  
CU Boulder Computer Science PhD Student  
Artificial Intelligence Group advised by Jordan Boyd-Graber



Department of Computer Science  
UNIVERSITY OF COLORADO **BOULDER**

# About Me

---

- 2nd year C.S. PhD Student at CU, Data Scientist
- Previously at: Oracle, Trulia, and the AMPLab
- UC Berkeley 2014 Graduate in Computer Science
- Interests: large-scale machine learning, deep learning, NLP
- Personal Interests: Ski, climb, hike, games

# QANTA Project Collaborators

---

- Jordan Boyd-Graber, CU Professor of Computer Science
- Mohit Iyyer, PhD Student at University of Maryland
- Hal Daumé III, Anupam Guha, He He, Brianna Satinoff, Manjhunath Ravi, Danny Bouman, Alvin Grissom



University of Colorado  
Boulder



UNIVERSITY OF  
MARYLAND



# Outline

---


- Goal: broad strokes overview for intuition
- What is Artificial Intelligence?
- Quiz Bowl Introduction
- Text Representation
- Deep Learning for content models (guessing)
- Feature Extraction
- Reinforcement Learning for buzzing



# What is Artificial Intelligence?

---

- Wiki: AI is the **intelligence** exhibited by machines or software
- John McCarthy: science and engineering of making **intelligent** machines
- What is intelligence?
  - Ability to learn or understand the world to make decisions in new or difficult situations



Machine Learning: design and development of algorithms to evolve behavior based on data



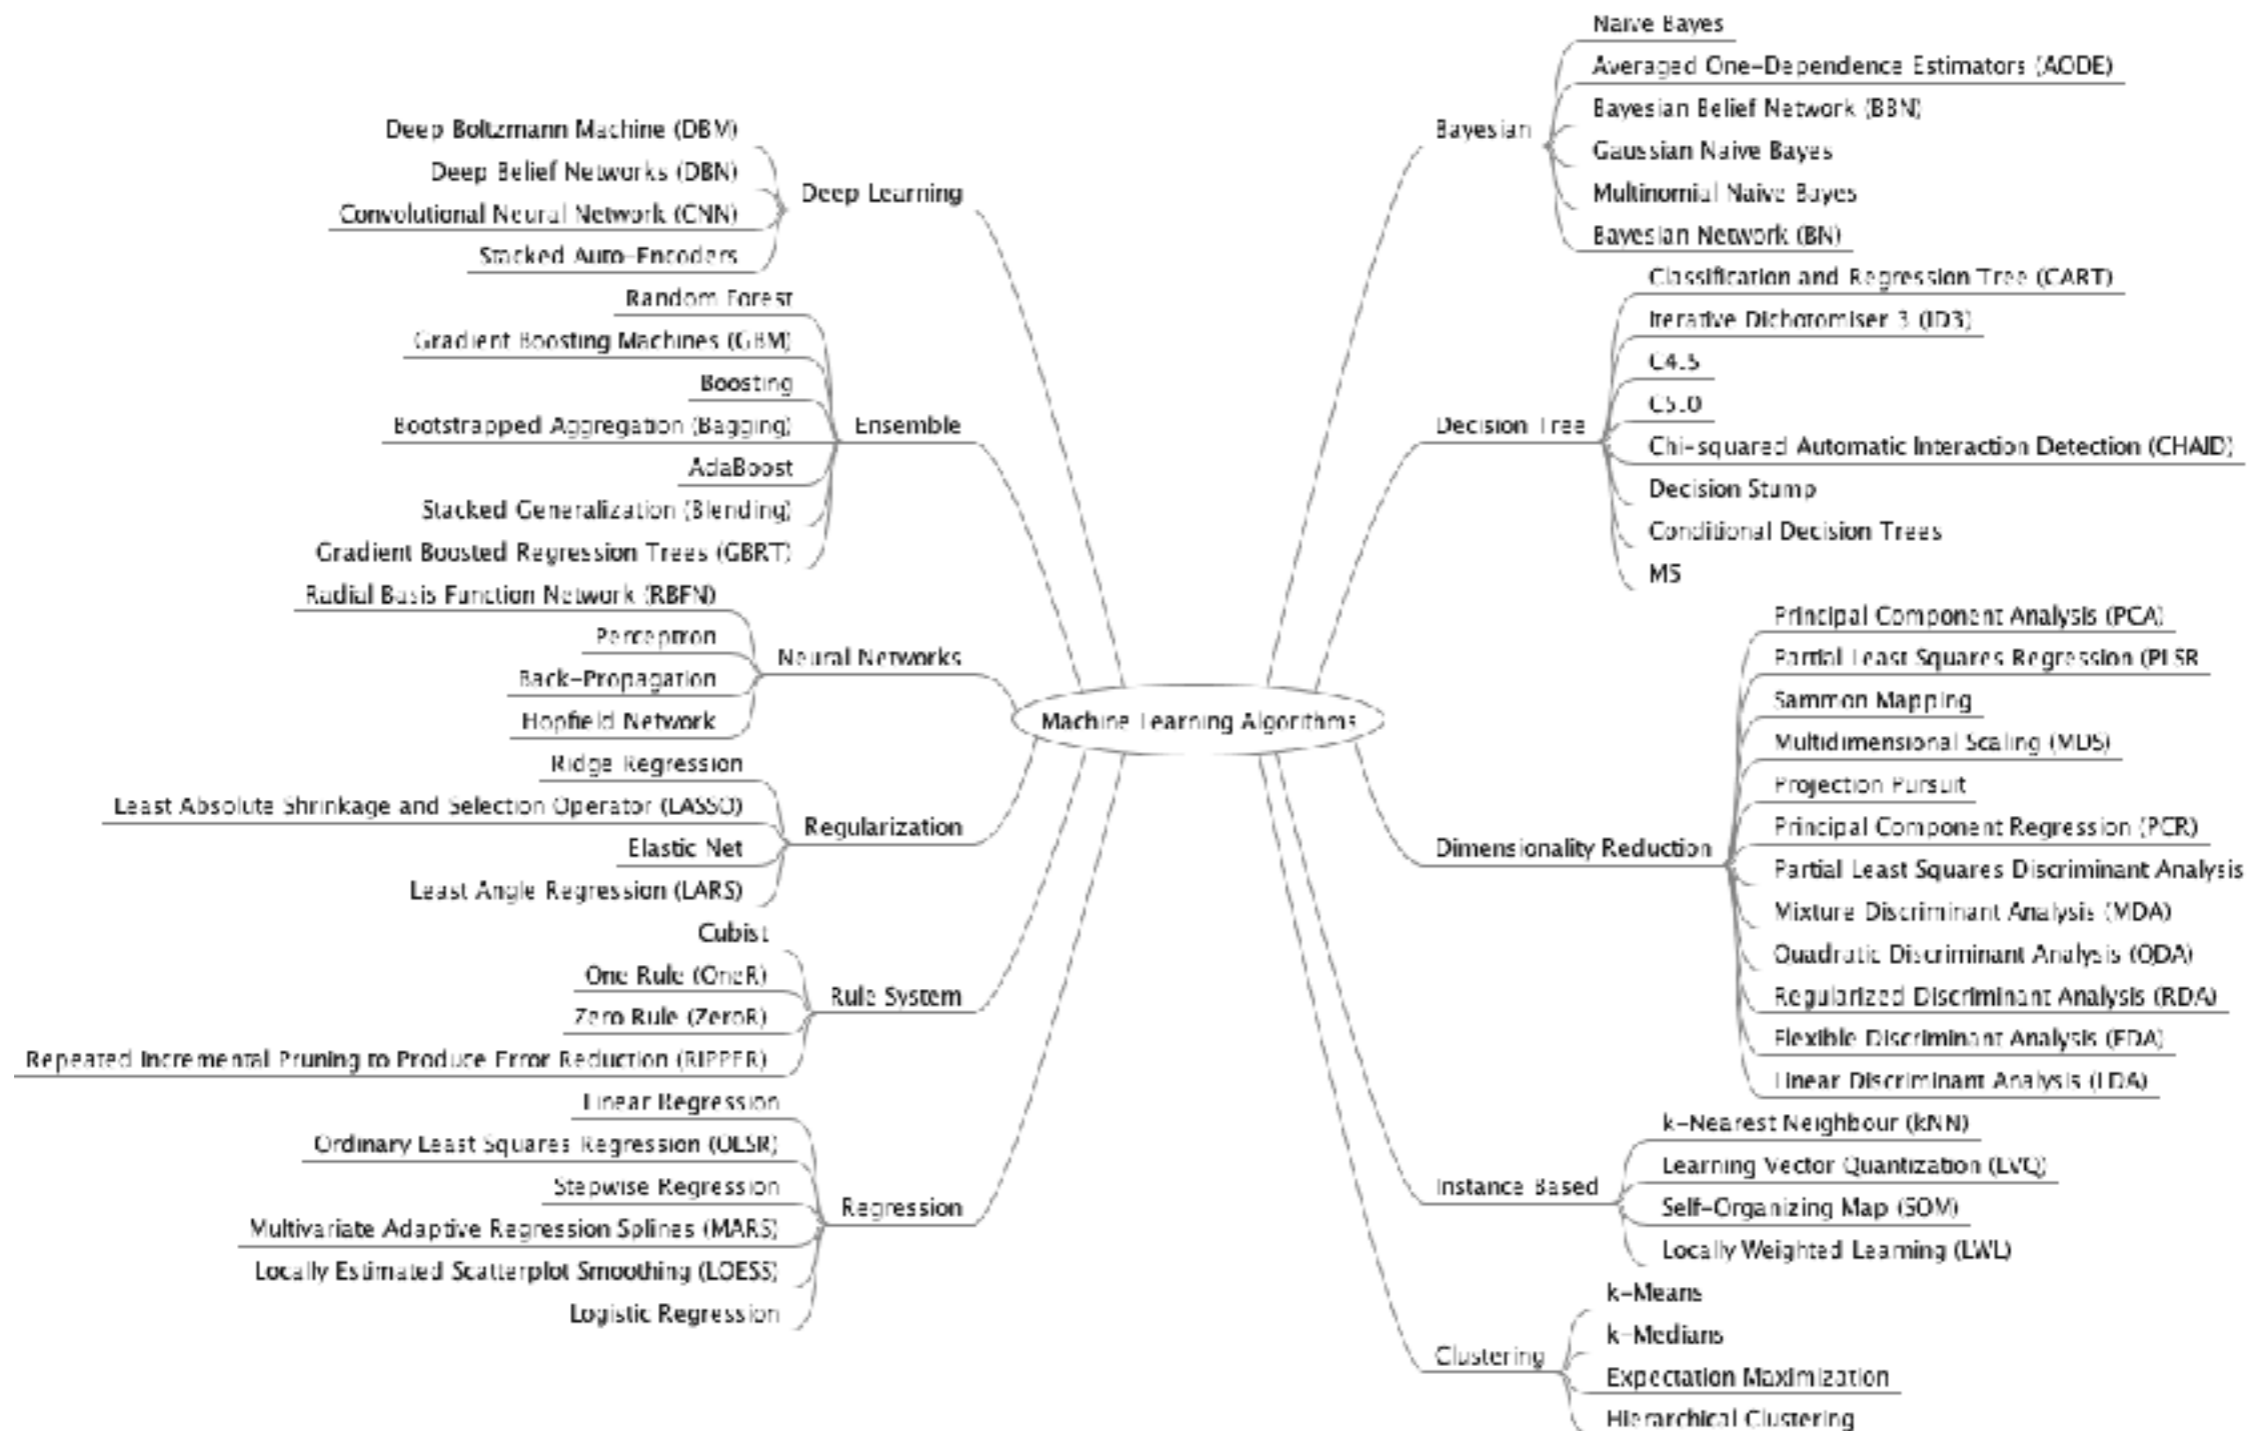
# Classification: Spam Email

---





# Machine Learning



# Games: Intelligent Agents

---

Given knowledge of gameplay and current game  
How do you play?



# Quiz Bowl

---

- Two teams play against each other
  - Moderator reads question
  - When team knows the answer “buzz” in
  - Correct guesses award points, wrong guesses let other team see entire question
- Thousands of teams in US



# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous

# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of

# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so

# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients



# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients for spontaneous and stimulated emission, the subject of one of his multiple groundbreaking 1905 papers. He further developed the model of statistics sent to him by

# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients for spontaneous and stimulated emission, the subject of one of his multiple groundbreaking 1905 papers. He further developed the model of statistics sent to him by Bose to describe particles with integer spin. For 10 points, who is this German physicist best known for formulating the

# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients for spontaneous and stimulated emission, the subject of one of his multiple groundbreaking 1905 papers. He further developed the model of statistics sent to him by Bose to describe particles with integer spin. For 10 points, who is this German physicist best known for formulating the special and general theories of relativity?

# Who is this question talking about?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients for spontaneous and stimulated emission, the subject of one of his multiple groundbreaking 1905 papers. He further developed the model of statistics sent to him by Bose to describe particles with integer spin. For 10 points, who is this German physicist best known for formulating the special and general theories of relativity?
- Albert Einstein

# Pyramidal Clues

---

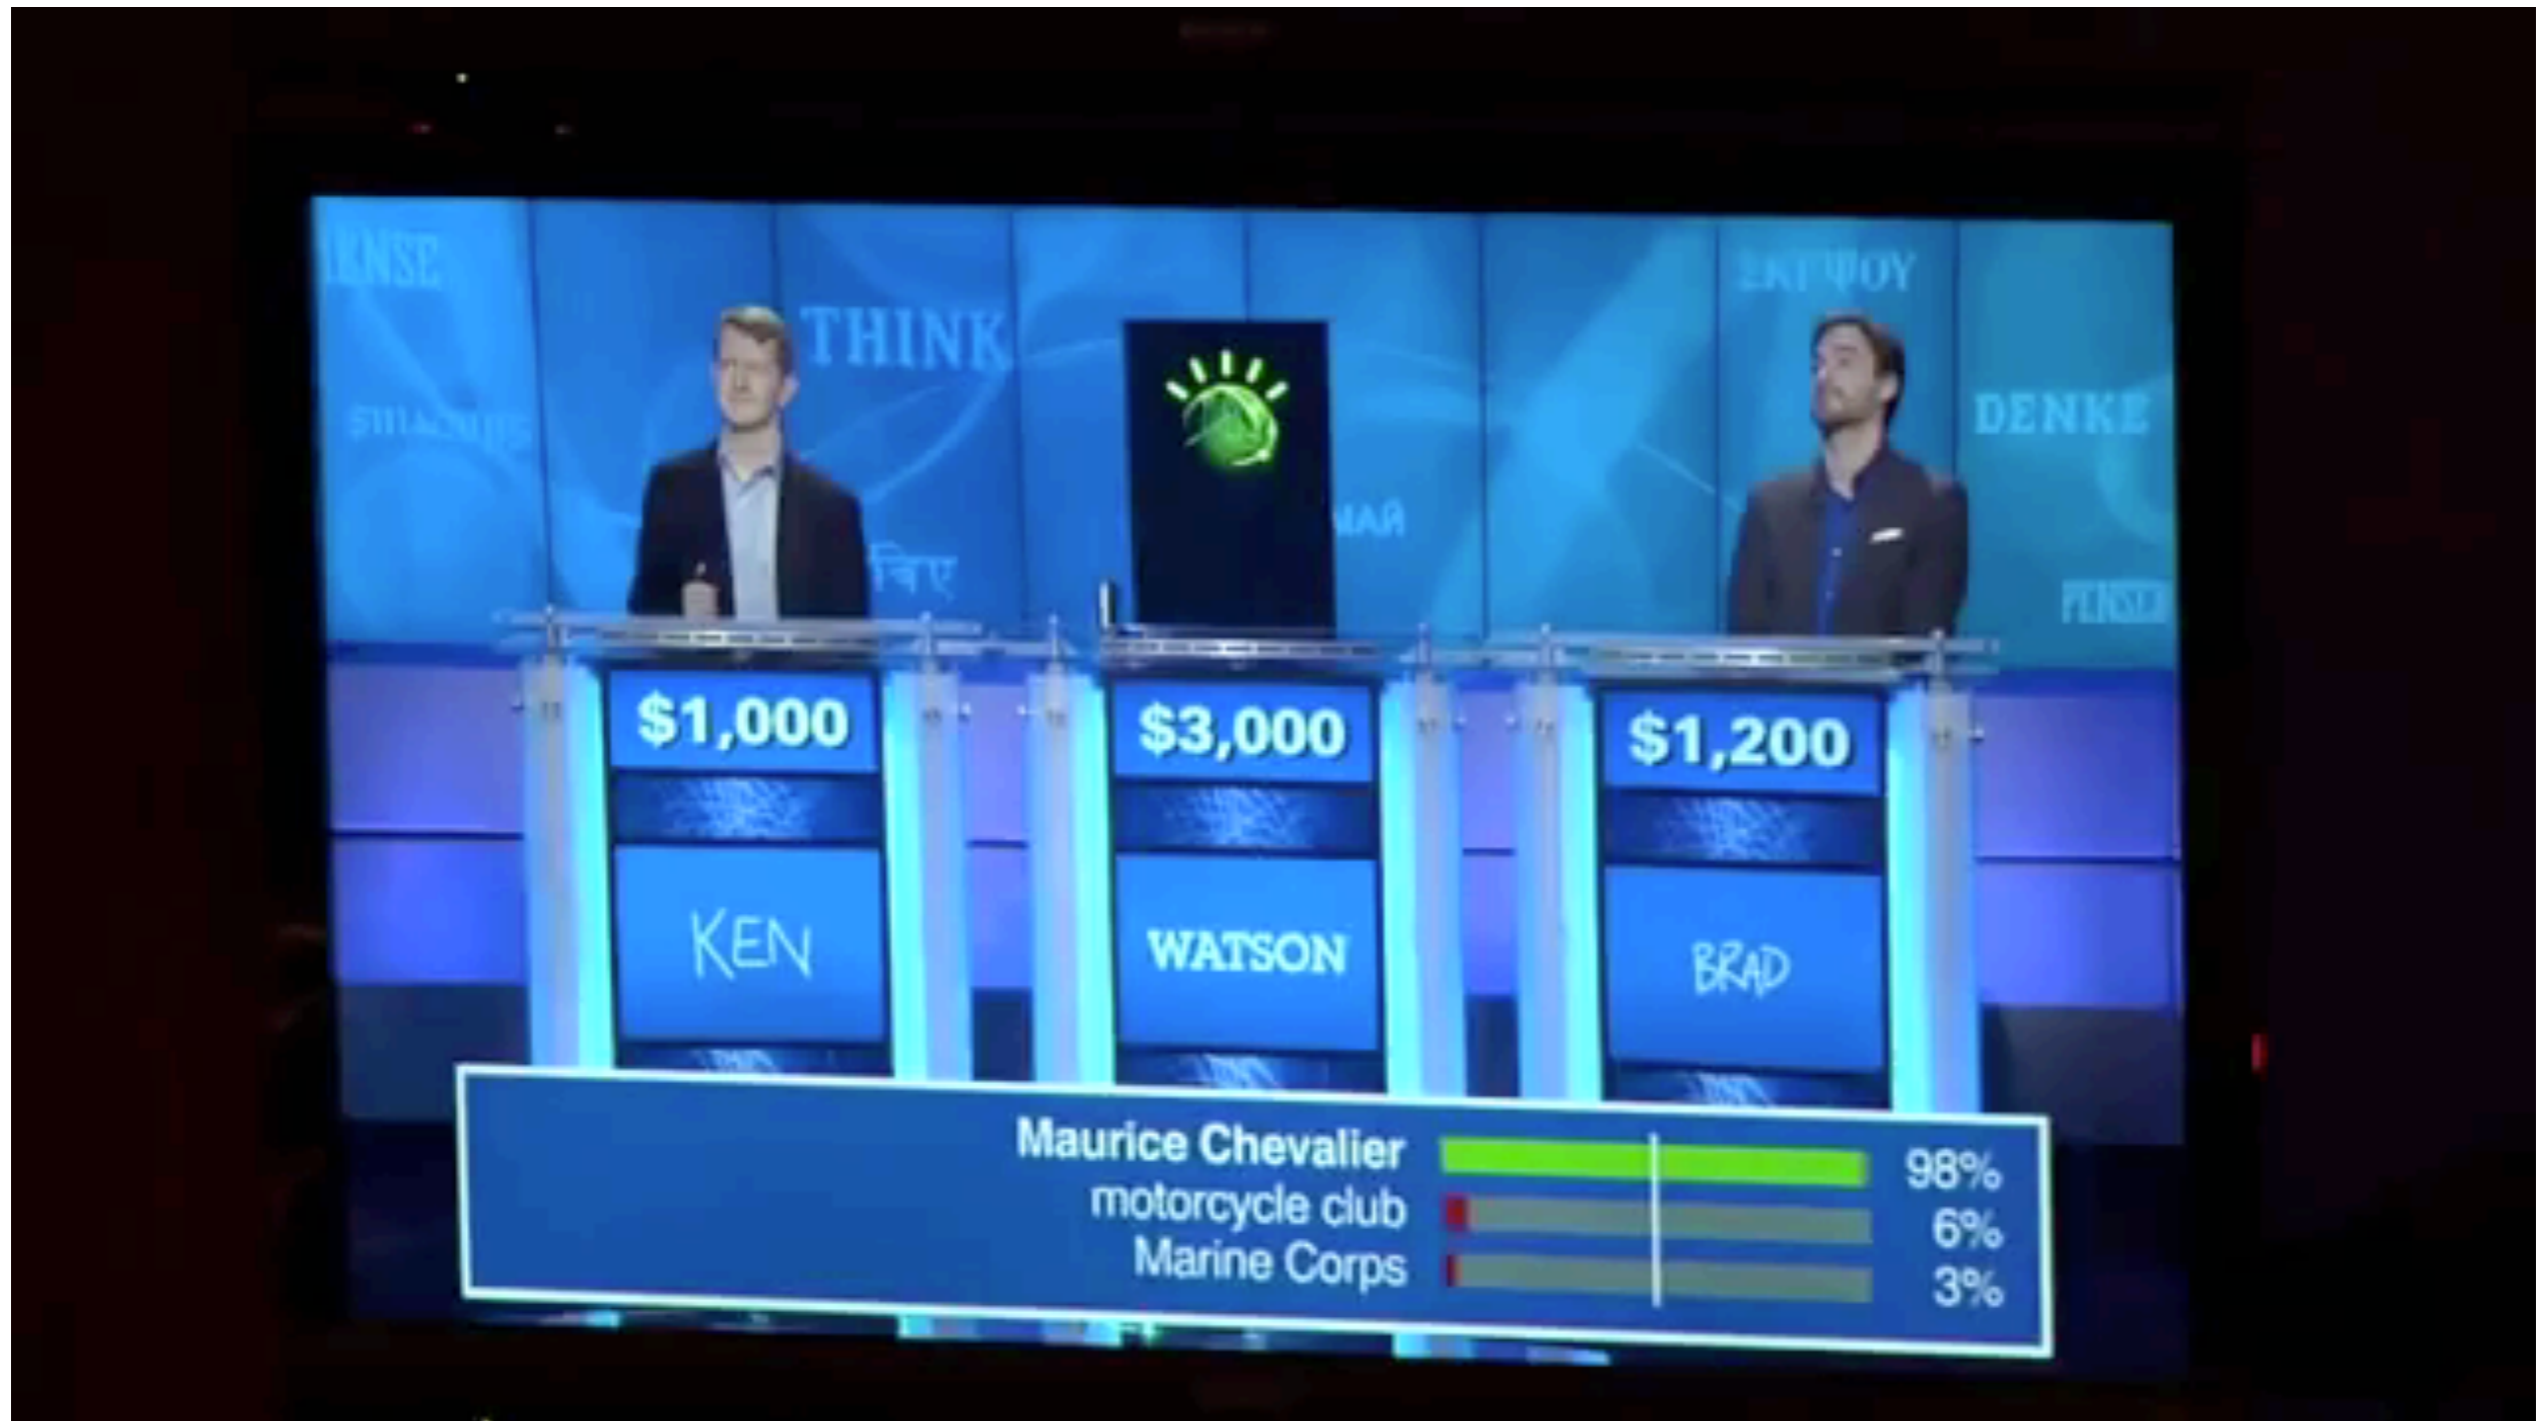
- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients for spontaneous and stimulated emission, the subject of one of his multiple groundbreaking 1905 papers. He further developed the model of statistics sent to him by Bose to describe particles with integer spin. For 10 points, who is this German physicist best known for formulating the special and general theories of relativity?
- Albert Einstein



# Quiz Bowl vs Jeopardy? IBM Watson vs QANTA?



# IBM Watson





# **QANTA: Question Answering is Not a Trivial Activity**

---

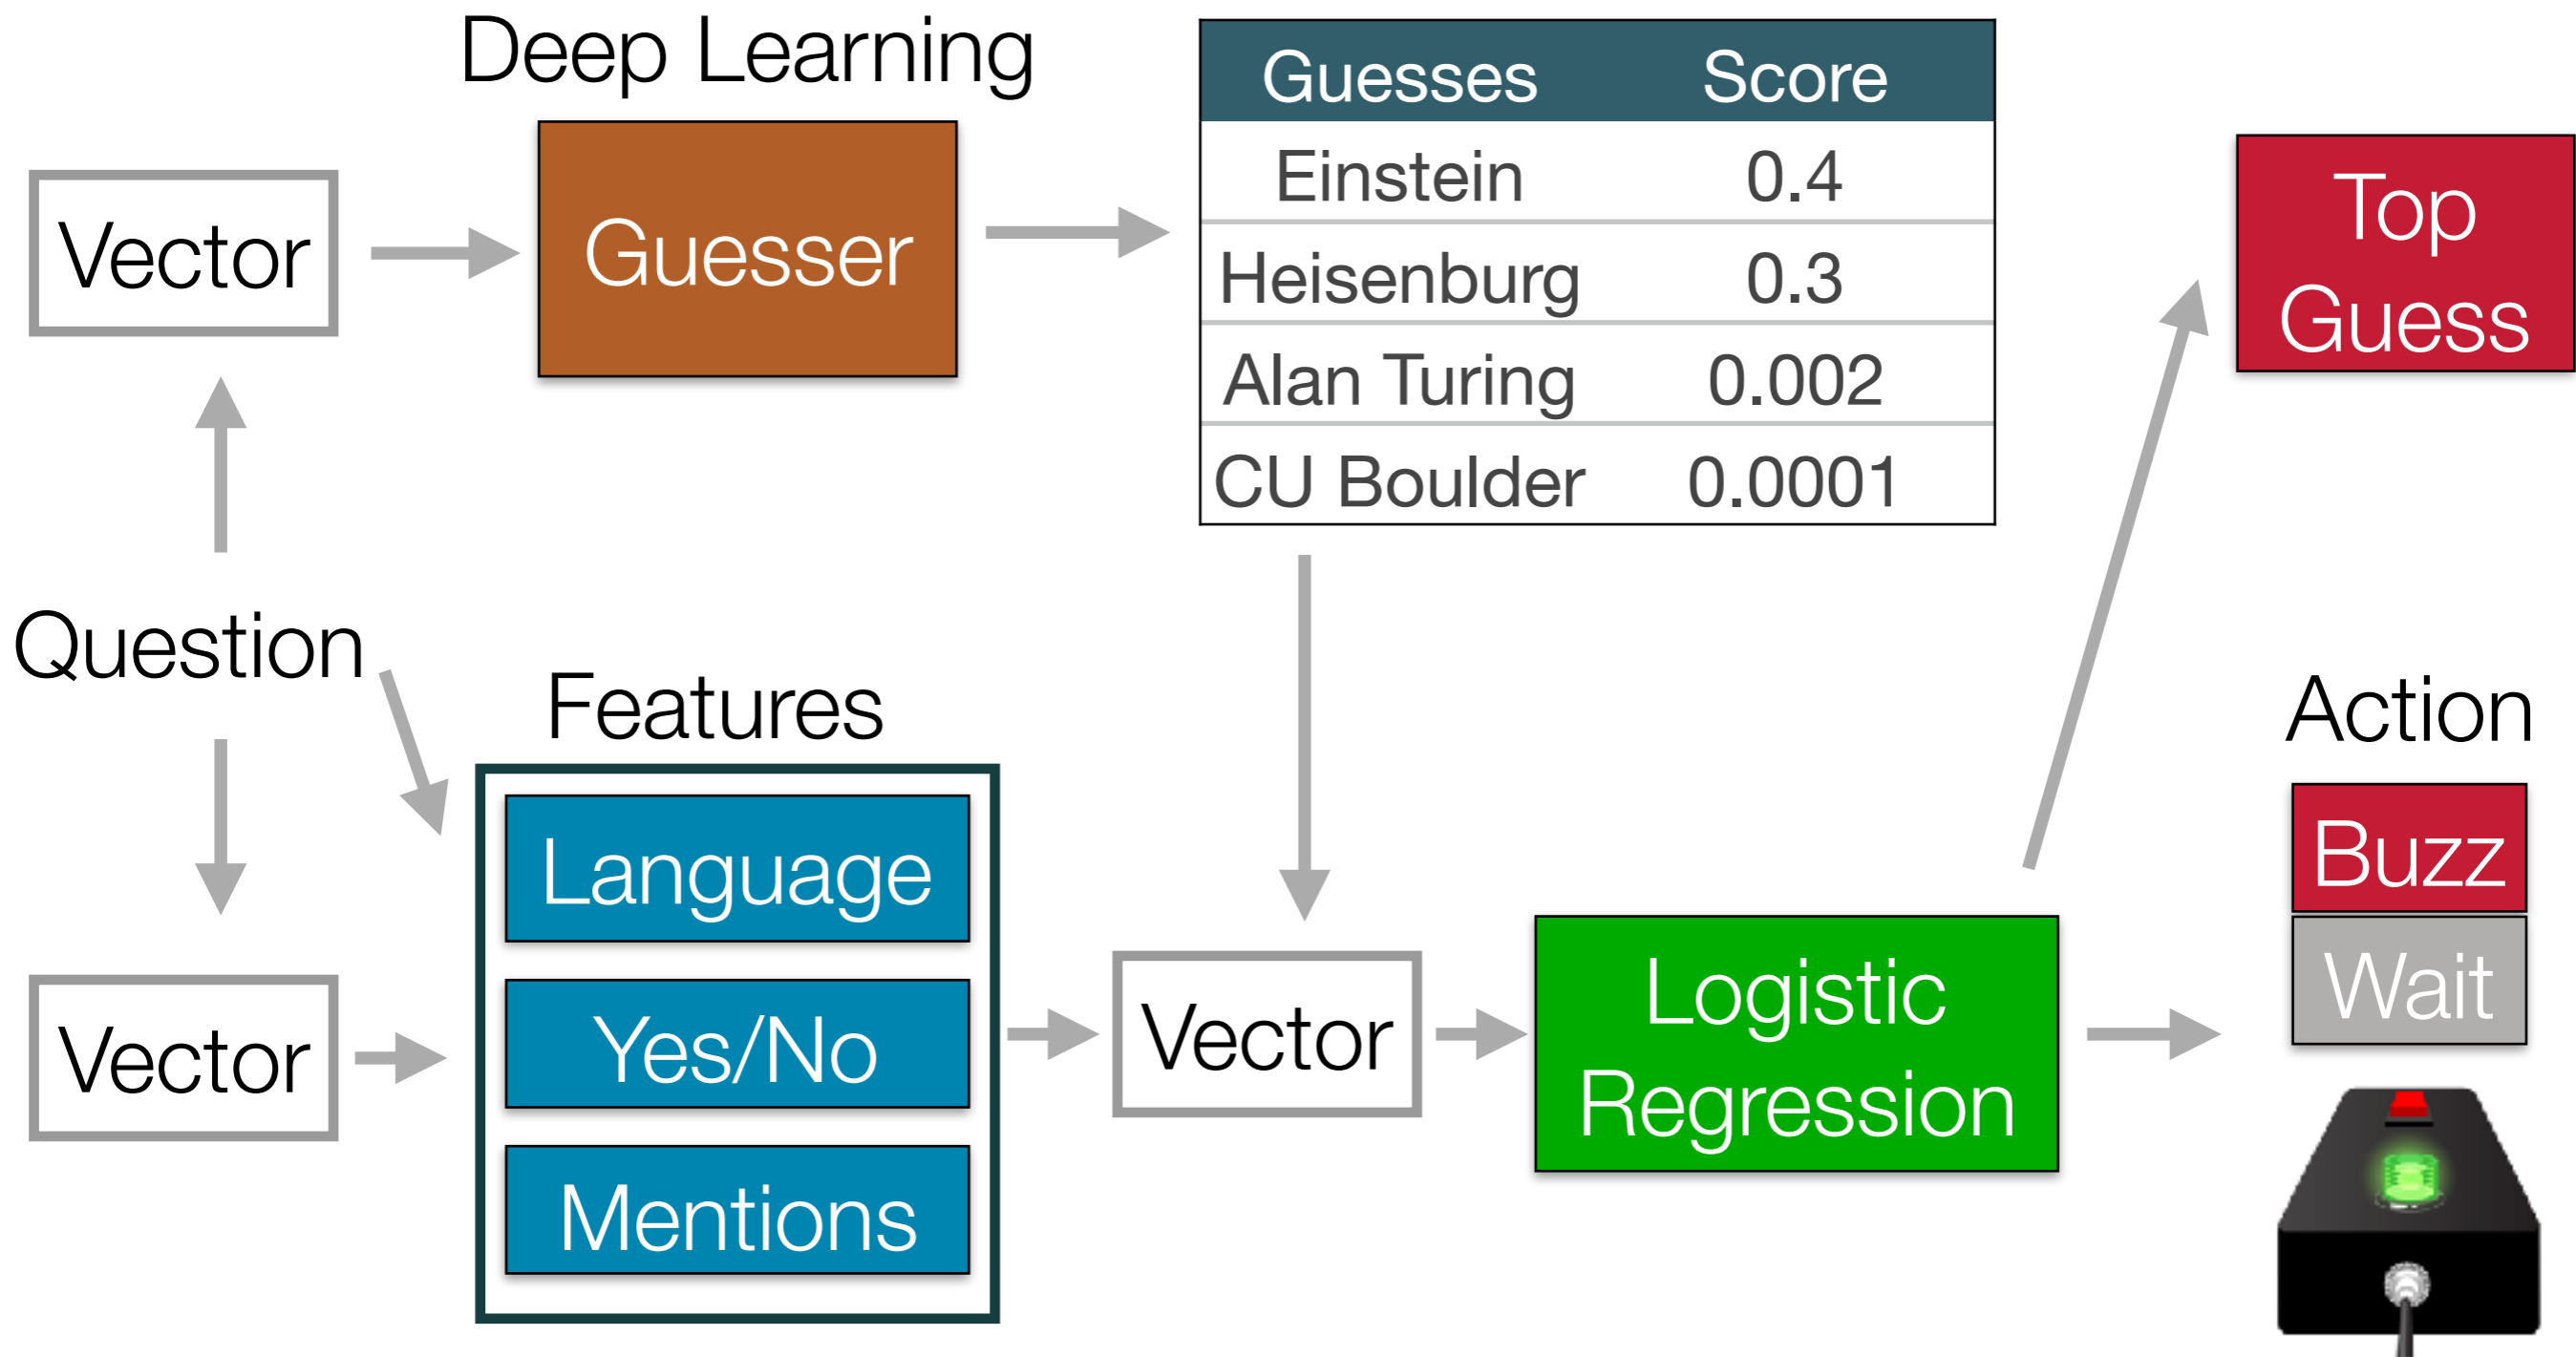


# Similarities and Differences

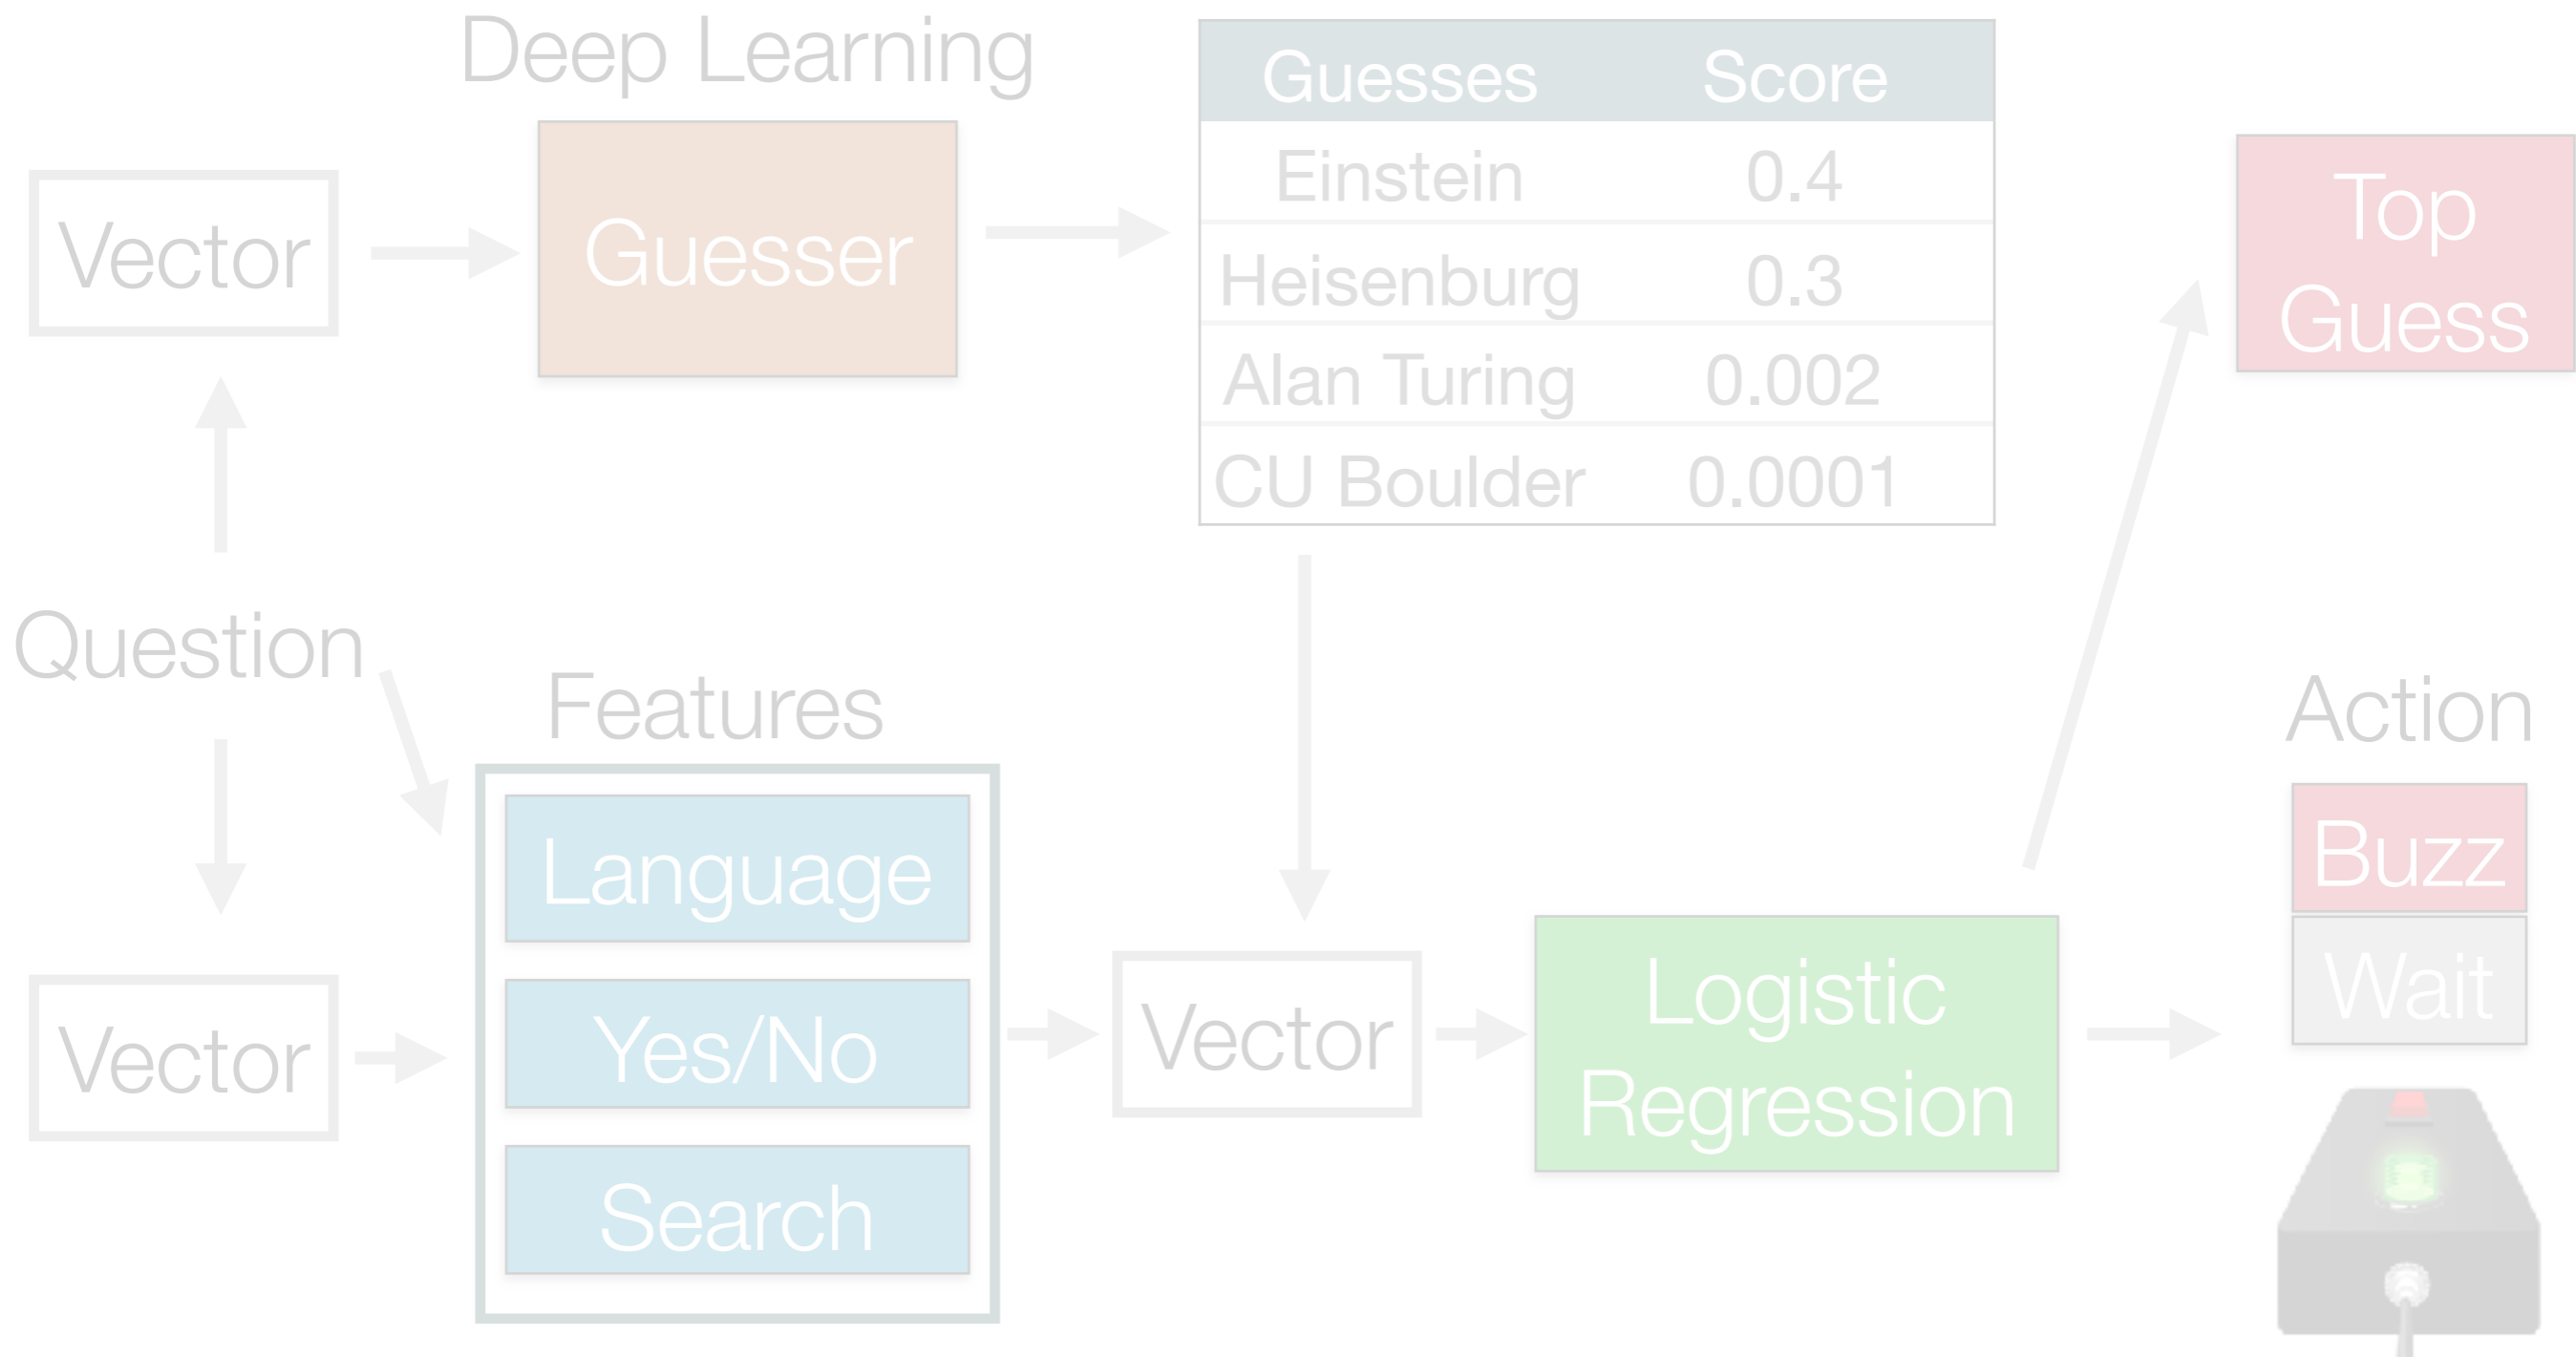
---

- Differences
  - Jeopardy: answer questions only at the end
  - Quiz Bowl: decide after each word
  - Quiz Bowl is pyramidal
  - Humans think more like QANTA than Watson
    - **Why?**

# QANTA Overview




# QANTA Overview



# Datasets

---

- Quiz Bowl Questions
  - ~200,000 questions
  - ~10,000 unique answers
  - ~2,000 answers have 5 or more questions
- Wikipedia
  - 50GB Decompressed (text)
  - Network/Graph Structure



**WIKIPEDIA**  
The Free Encyclopedia

# Incremental Learning

---

- Ordinary algorithms assume **all** of **X** is **known** at **once**
- How does Quiz Bowl differ?
  - Receive input  $X$  “incrementally”, one word at a time
  - Quiz Bowl Task doesn't cleanly fit into single prediction
  - How do we deal with that?
  - What other tasks are like this?

# Quiz Bowl Task

---


- Overall: return the correct answer as soon as possible
- Break the problem down
  - What should we answer with?
  - When should we answer the question?
- Treat as Markov Decision Process



# Markov Decision Process

---

- Framework for modeling decision-making in situations where outcomes are partly random and partly under control of decision-maker



# Quiz Bowl Markov Decision Process

---

- We have control of when we can buzz
- We have control of what we answer
- Don't have control over when opponent answers
- Don't have control over quality of next "clue"
- Do I risk letting opponent answer for more information?

# Two Steps to Answering Questions


---

- Given a question
  - Generate a set of **guesses** (deep learning + reranking)
  - **Buzz** if confident (reinforcement learning + model)
- Deep learning?
- Reranking?
- Reinforcement Learning?

# How can we represent text numerically?

---

- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients for spontaneous and stimulated emission, the subject of one of his multiple groundbreaking 1905 papers. He further developed the model of statistics sent to him by Bose to describe particles with integer spin. For 10 points, who is this German physicist best known for formulating the special and general theories of relativity?



# One Hot Encoding Vectors

---

Text Vector

$$[0, \dots, 1, \dots, 0]$$

- Represent text as a  $|V|$  size vector (vocabulary size)
- 1s mark word presence, 0s absence
- Problems?
  - Sparse, wasteful representation, especially for n-grams
  - No notion of similarity of words with each other

# Context Matters!

---

- From Wikipedia (with some editing)
  - “Albert **Einstein** developed the **general theory** of **relativity**”
  - “This led to the development of Einstein’s **special theory** of **relativity**”
- For 10 points, who is this German physicist best known for formulating the **special** and **general theories** of **relativity**?
- The wiki page on Einstein is pretty similar to the question so the “distance” between the vectors should be small
- n-grams alone can’t capture the level of similarity desired

# Context Dependent Vectors

---

- “You shall know a word by the company it keeps” -J. R. Firth
- Track co-occurrence of words
- Problem: expensive to compute
- Recent solution: word2vec

government debt problems turning into banking crises as has happened in  
saying that Europe needs unified banking regulation to replace the hodgepodge

↖ These words will represent *banking* ↗

# Word2Vec (Mikolov 2013)

---

- Idea: sliding “context” window around current center word
- Maximize probability of any context given center word

$v_w$  = Vector representation of word  $w$

$v_c$  = Vector representation of context  $c$

$\theta$  = Matrix of vector representations

$C(w)$  = set of context words of  $w$

$D$  = Set of all words and context pairs



## Word2Vec: Objectives

---

$$\arg \max_{\theta} \prod_{(w,c) \in D} p(c|w; \theta)$$

$$\arg \max_{\theta} \prod_{w \in \text{text}} \left[ \prod_{c \in C(w)} p(c|w; \theta) \right]$$

## Word2Vec: Model

---

$$p(c|w; \theta) = \frac{e^{v_c \cdot v_w}}{\sum_{c' \in C} e^{v_{c'} \cdot v_w}}$$



$$\arg \max_{\theta} \sum_{(w,c) \in D} \log p(c|w)$$

$$\arg \max_{\theta} (\log e^{v_c \cdot v_w} - \log \sum_{c'} e^{v_{c'} \cdot v + w})$$


# Word2Vec (Mikolov 2013)

---

$$\text{vec}(\text{"king"}) - \text{vec}(\text{"man"}) + \text{vec}(\text{"woman"}) \approx \text{vec}(\text{"queen"})$$

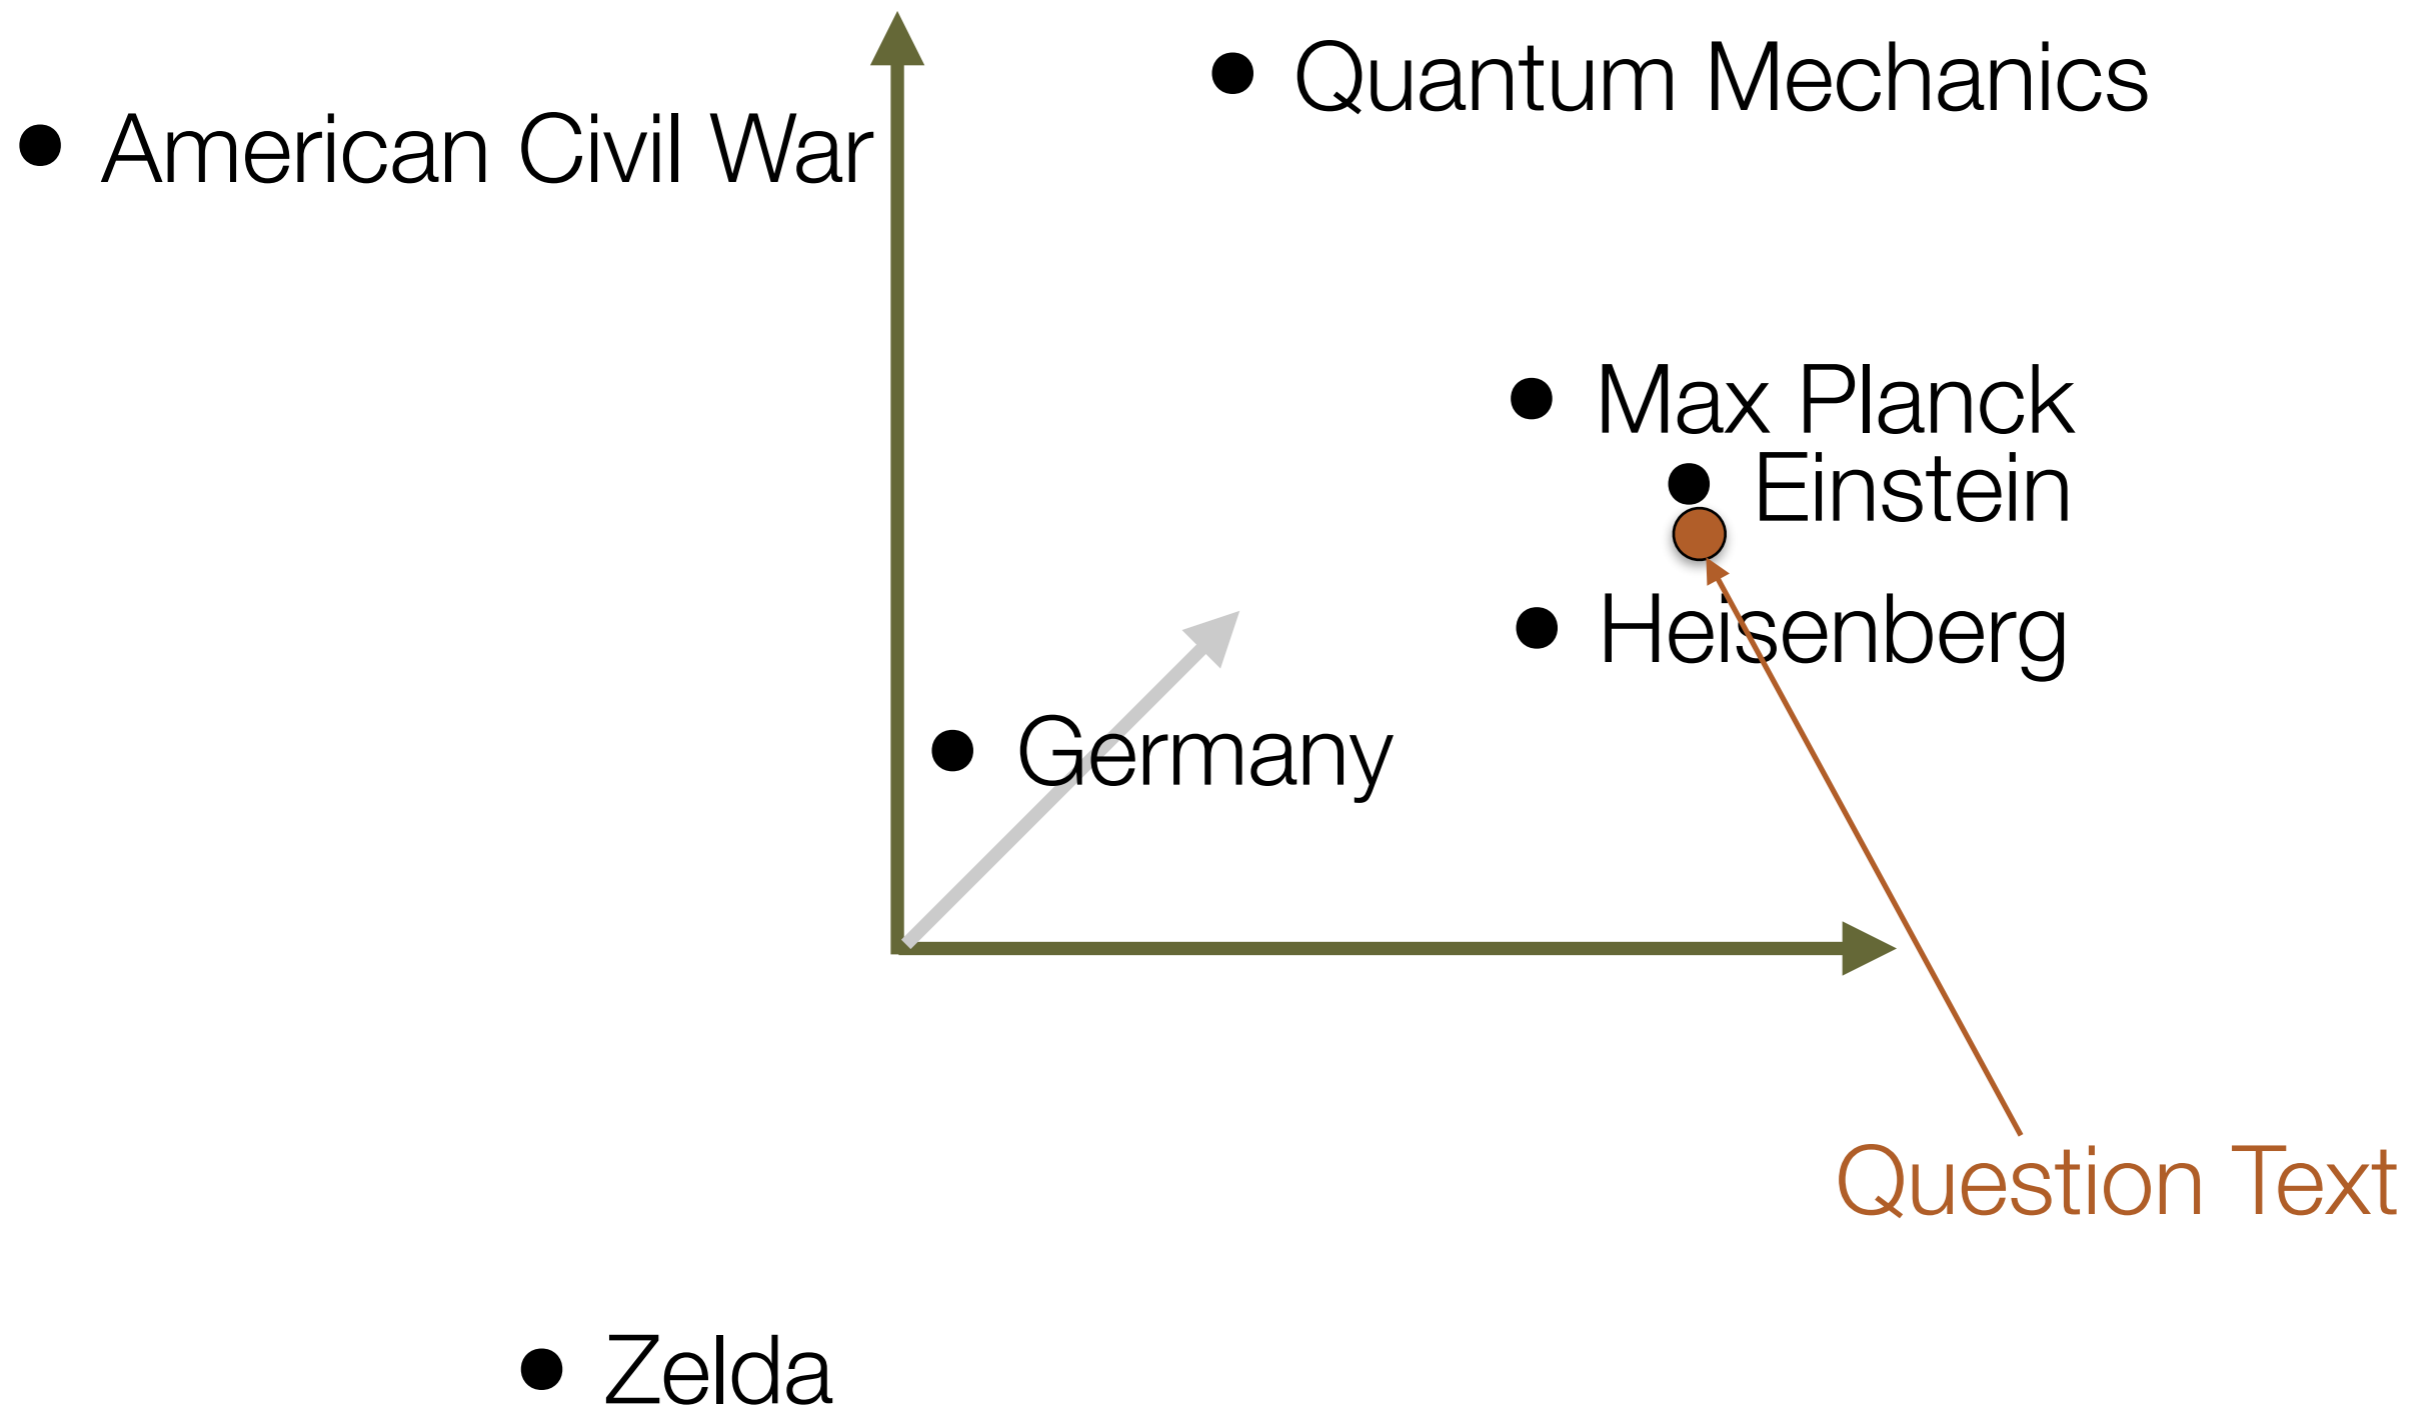


# Word2Vec



# QANTA Embedding

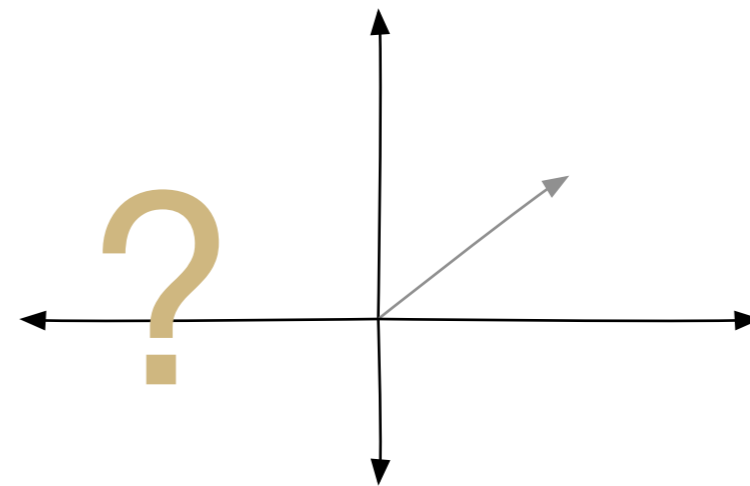

---



# One Embedding: Dependency Parse

---


This Chinese admiral attacked kotte in ceylon



# One Embedding: Dependency Parse

---

This Chinese admiral attacked kotte in ceylon

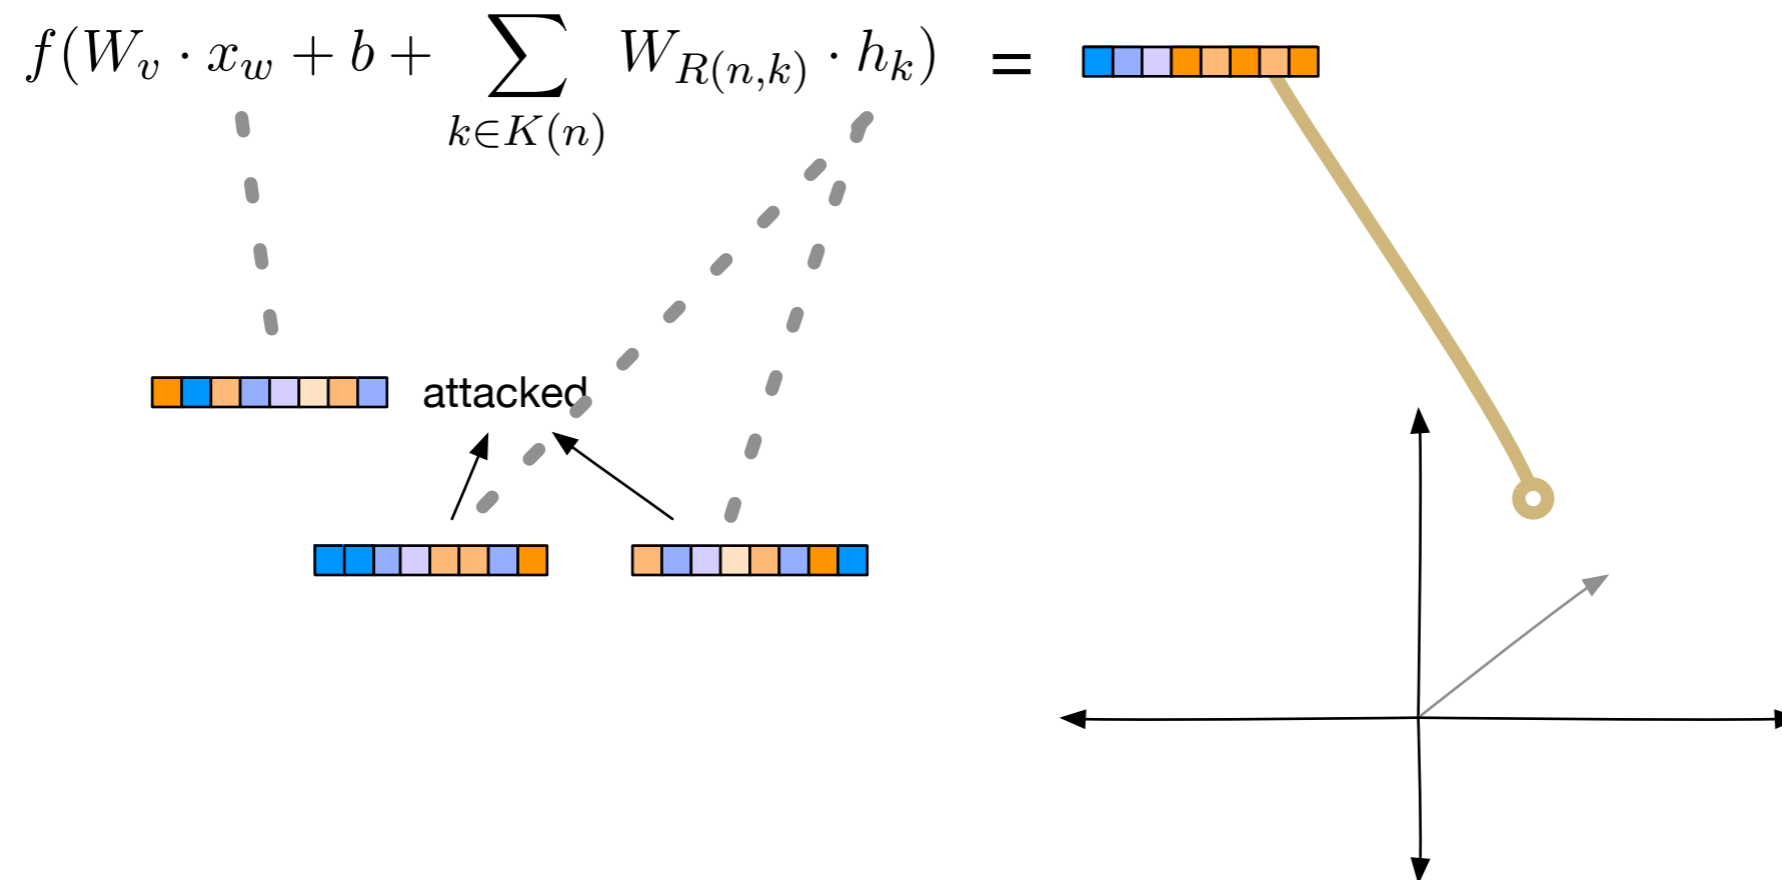




# One Embedding: Dependency Parse

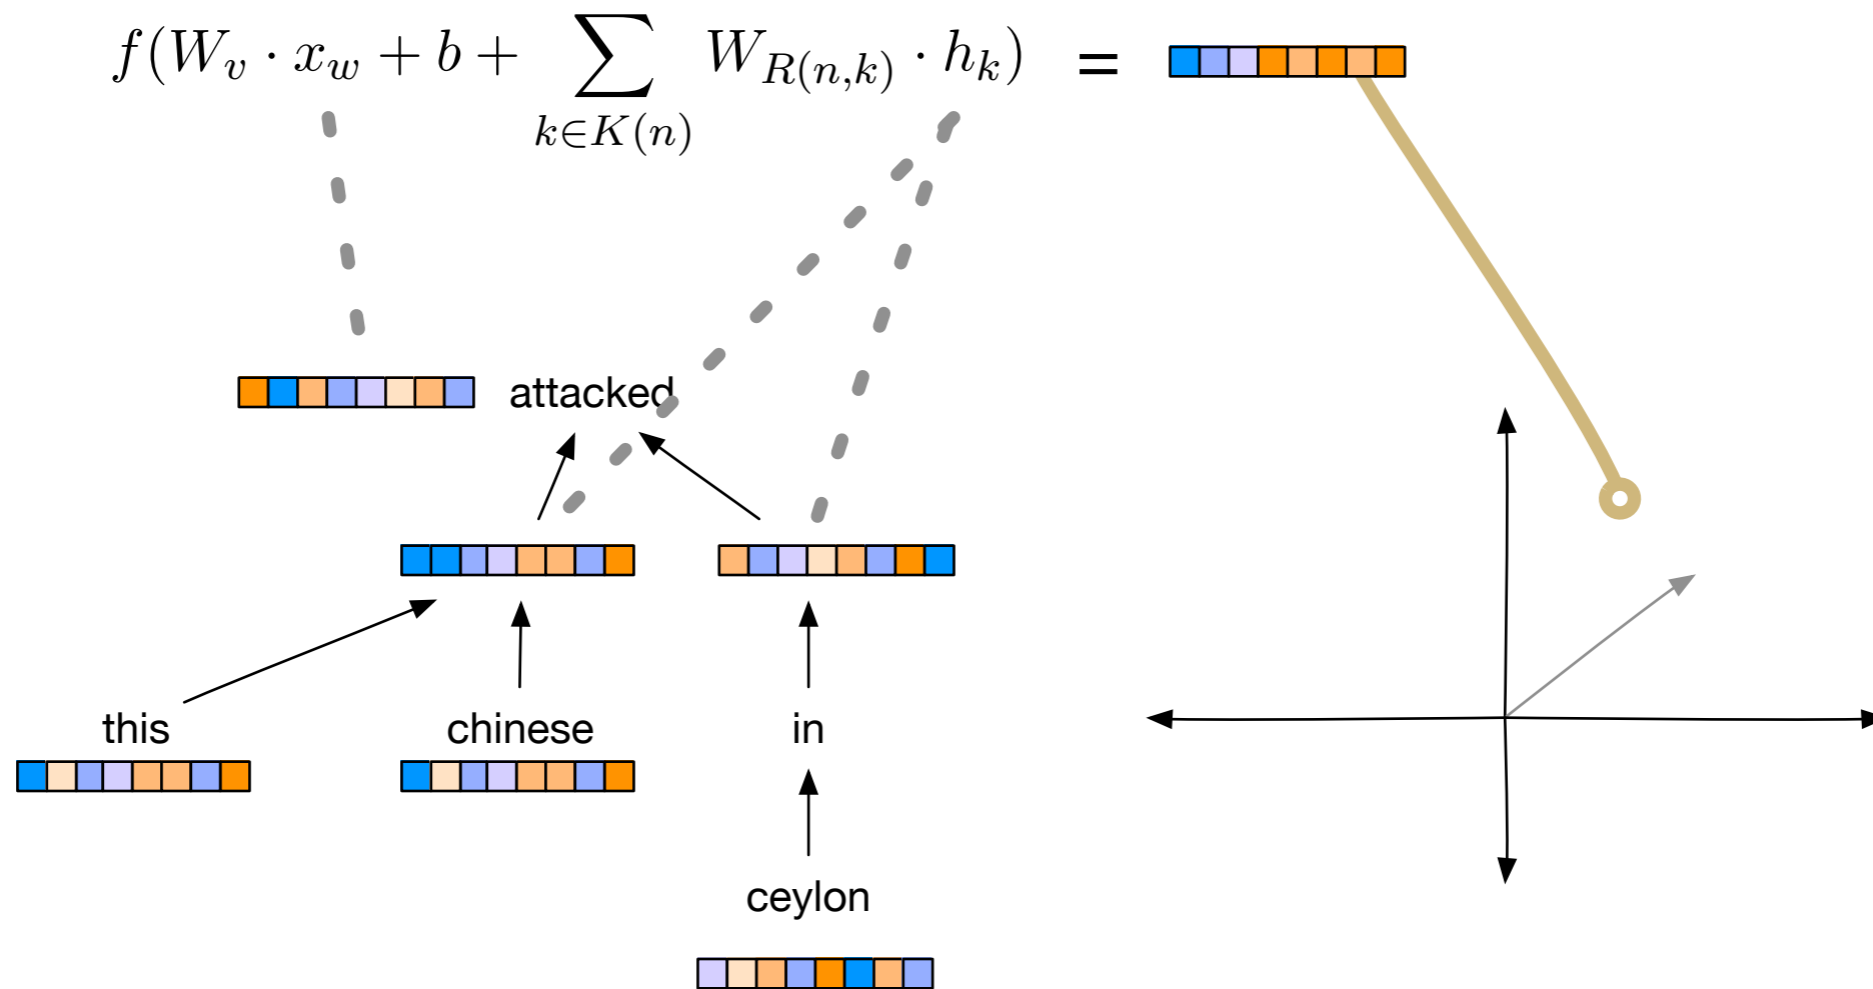
---

This Chinese admiral attacked motte in ceylon



# One Embedding: Dependency Parse


This Chinese admiral attacked motte in ceylon



# QANTA Vector Space

TSNE-2

TSNE-1



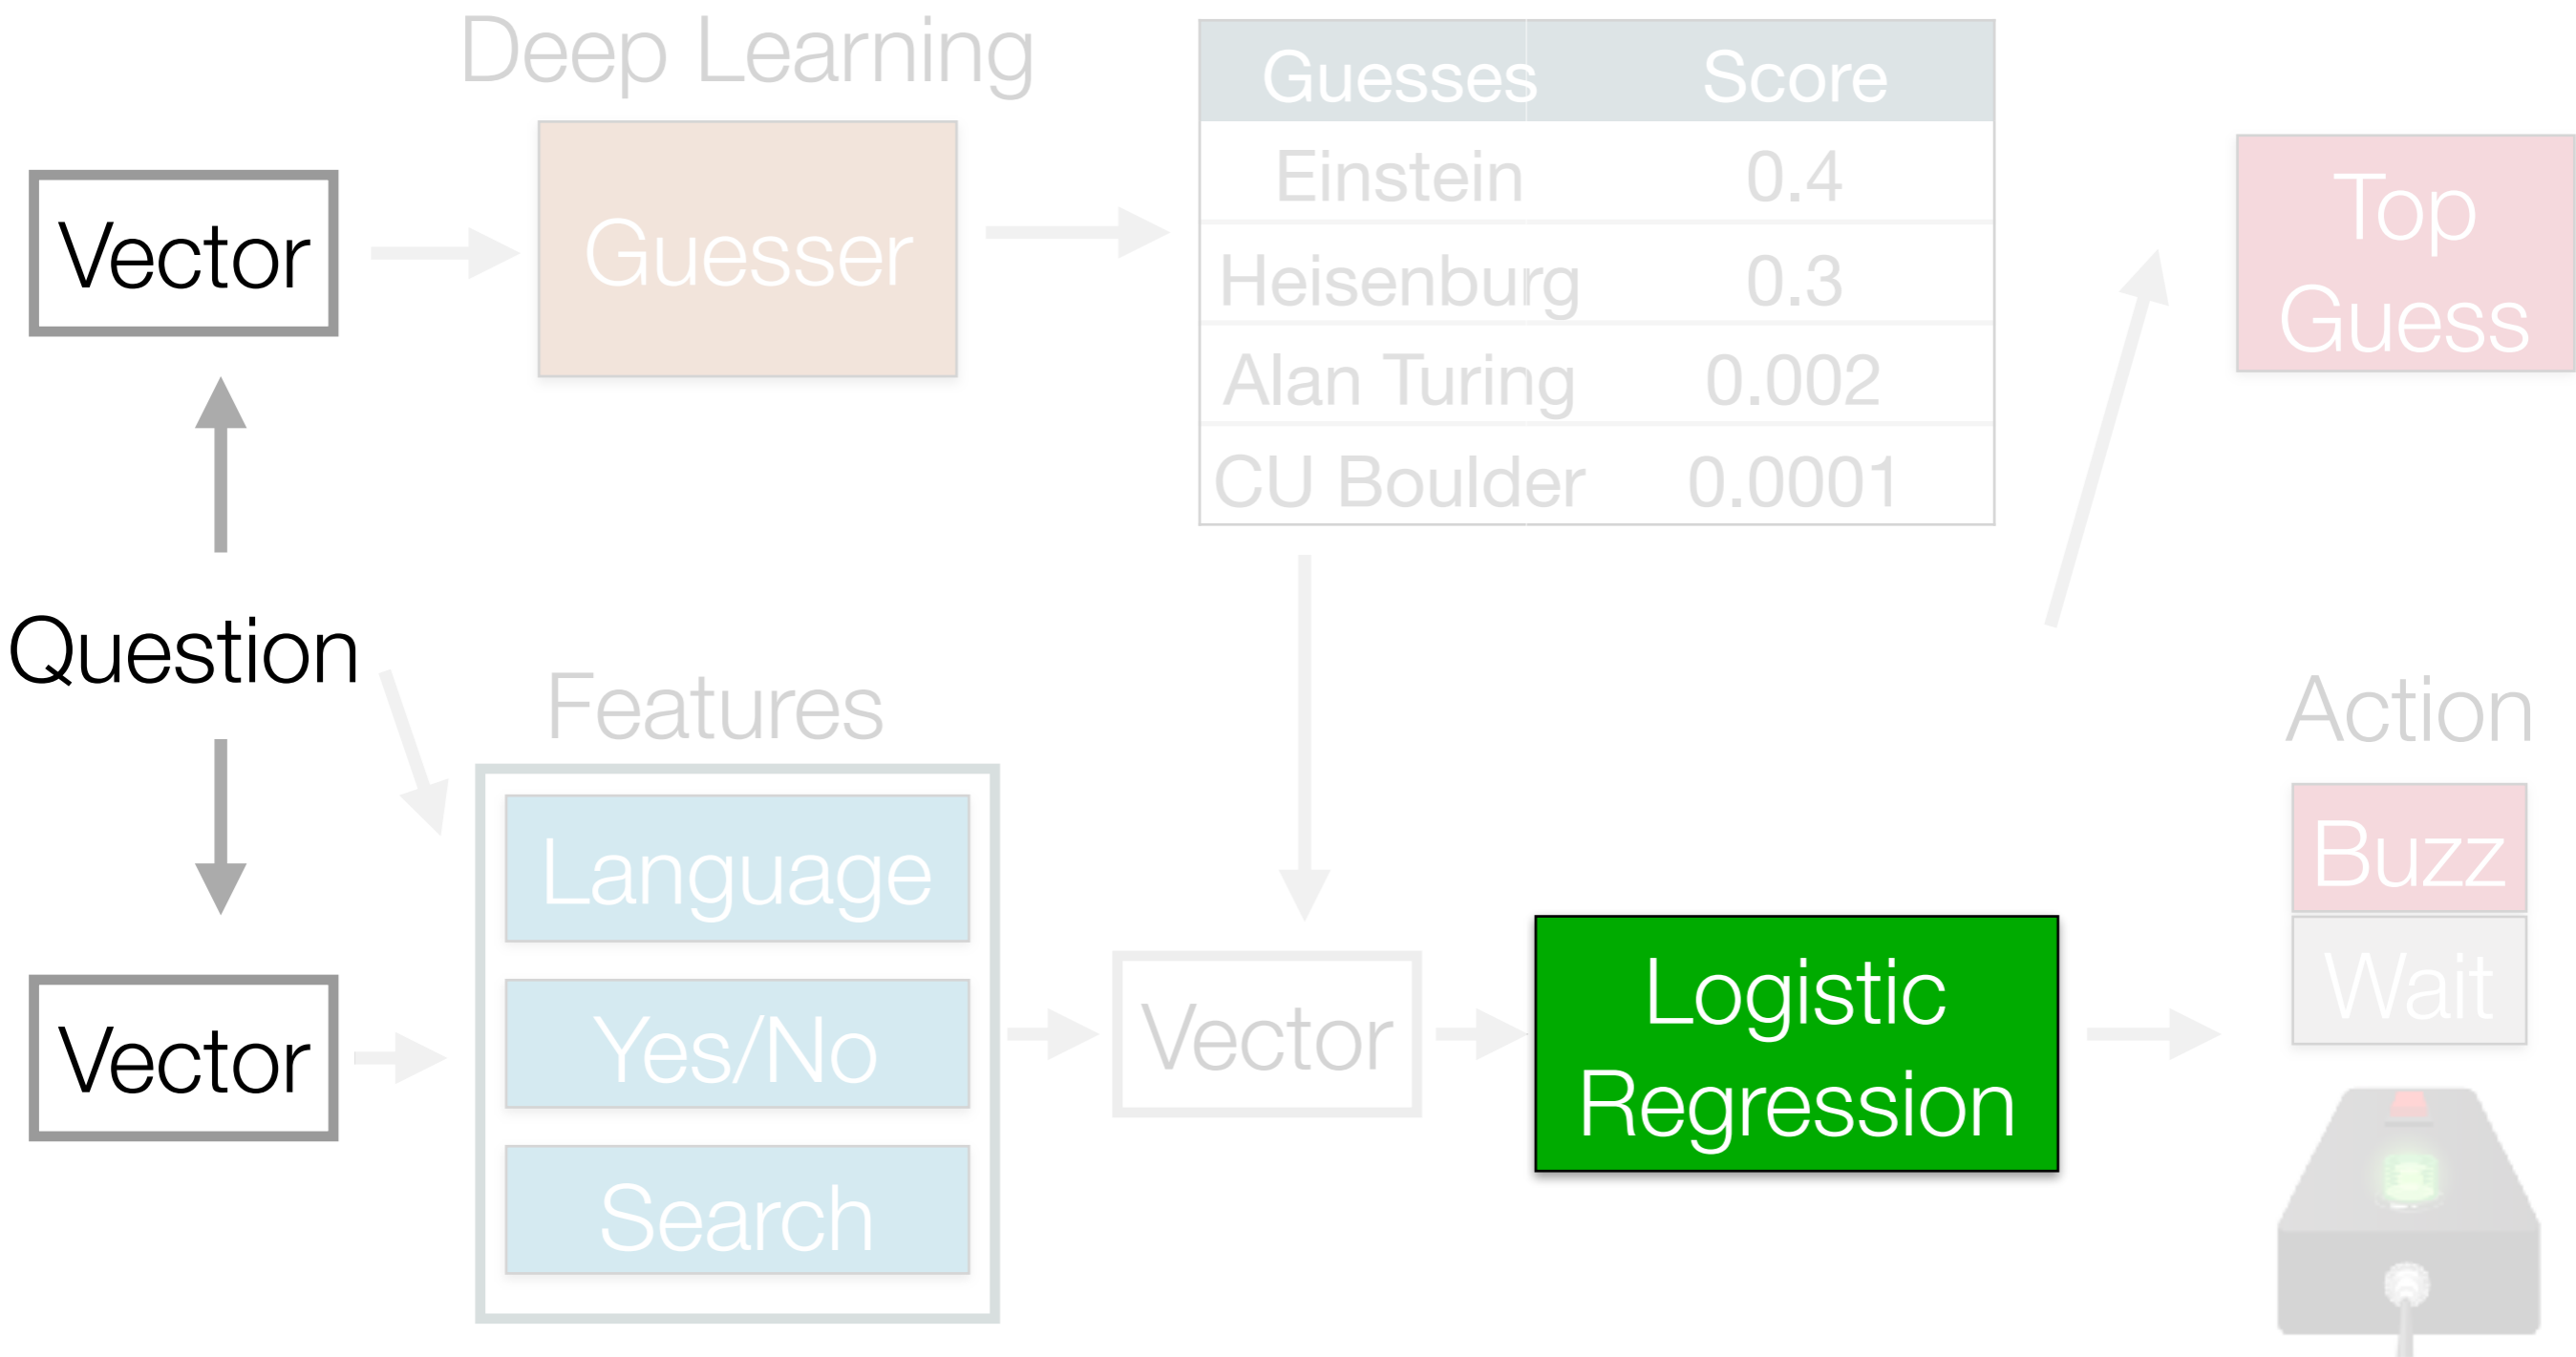
- Wars, rebellions, and battles
- U.S. presidents
- Prime ministers
- Explorers & emperors
- Policies
- Other

# Another approach: Neural Bag of Words



---

- Turns out a much simpler model works almost as well
- Neural bag of words: sum vectors together
- Train using a deep learning method

# QANTA Overview





# Neural Networks and Deep Learning



# Perceptron (like Logistic Regression)

---


$$P(y = 1|x) = h_{\Theta} = \sigma(\Theta^T x)$$
$$P(y = 0|x) = 1 - P(y = 1|x)$$





# Neural Network

---




# Neural Networks


---

- Related to Perceptron
- Instead of outputting 0 or 1, feed into another perceptron
- Stack these “deep” and its called deep learning
- Next:
  - Binary and Multiclass Classification
  - Training Neural Networks

# Neural Networks: Compose many perceptrons



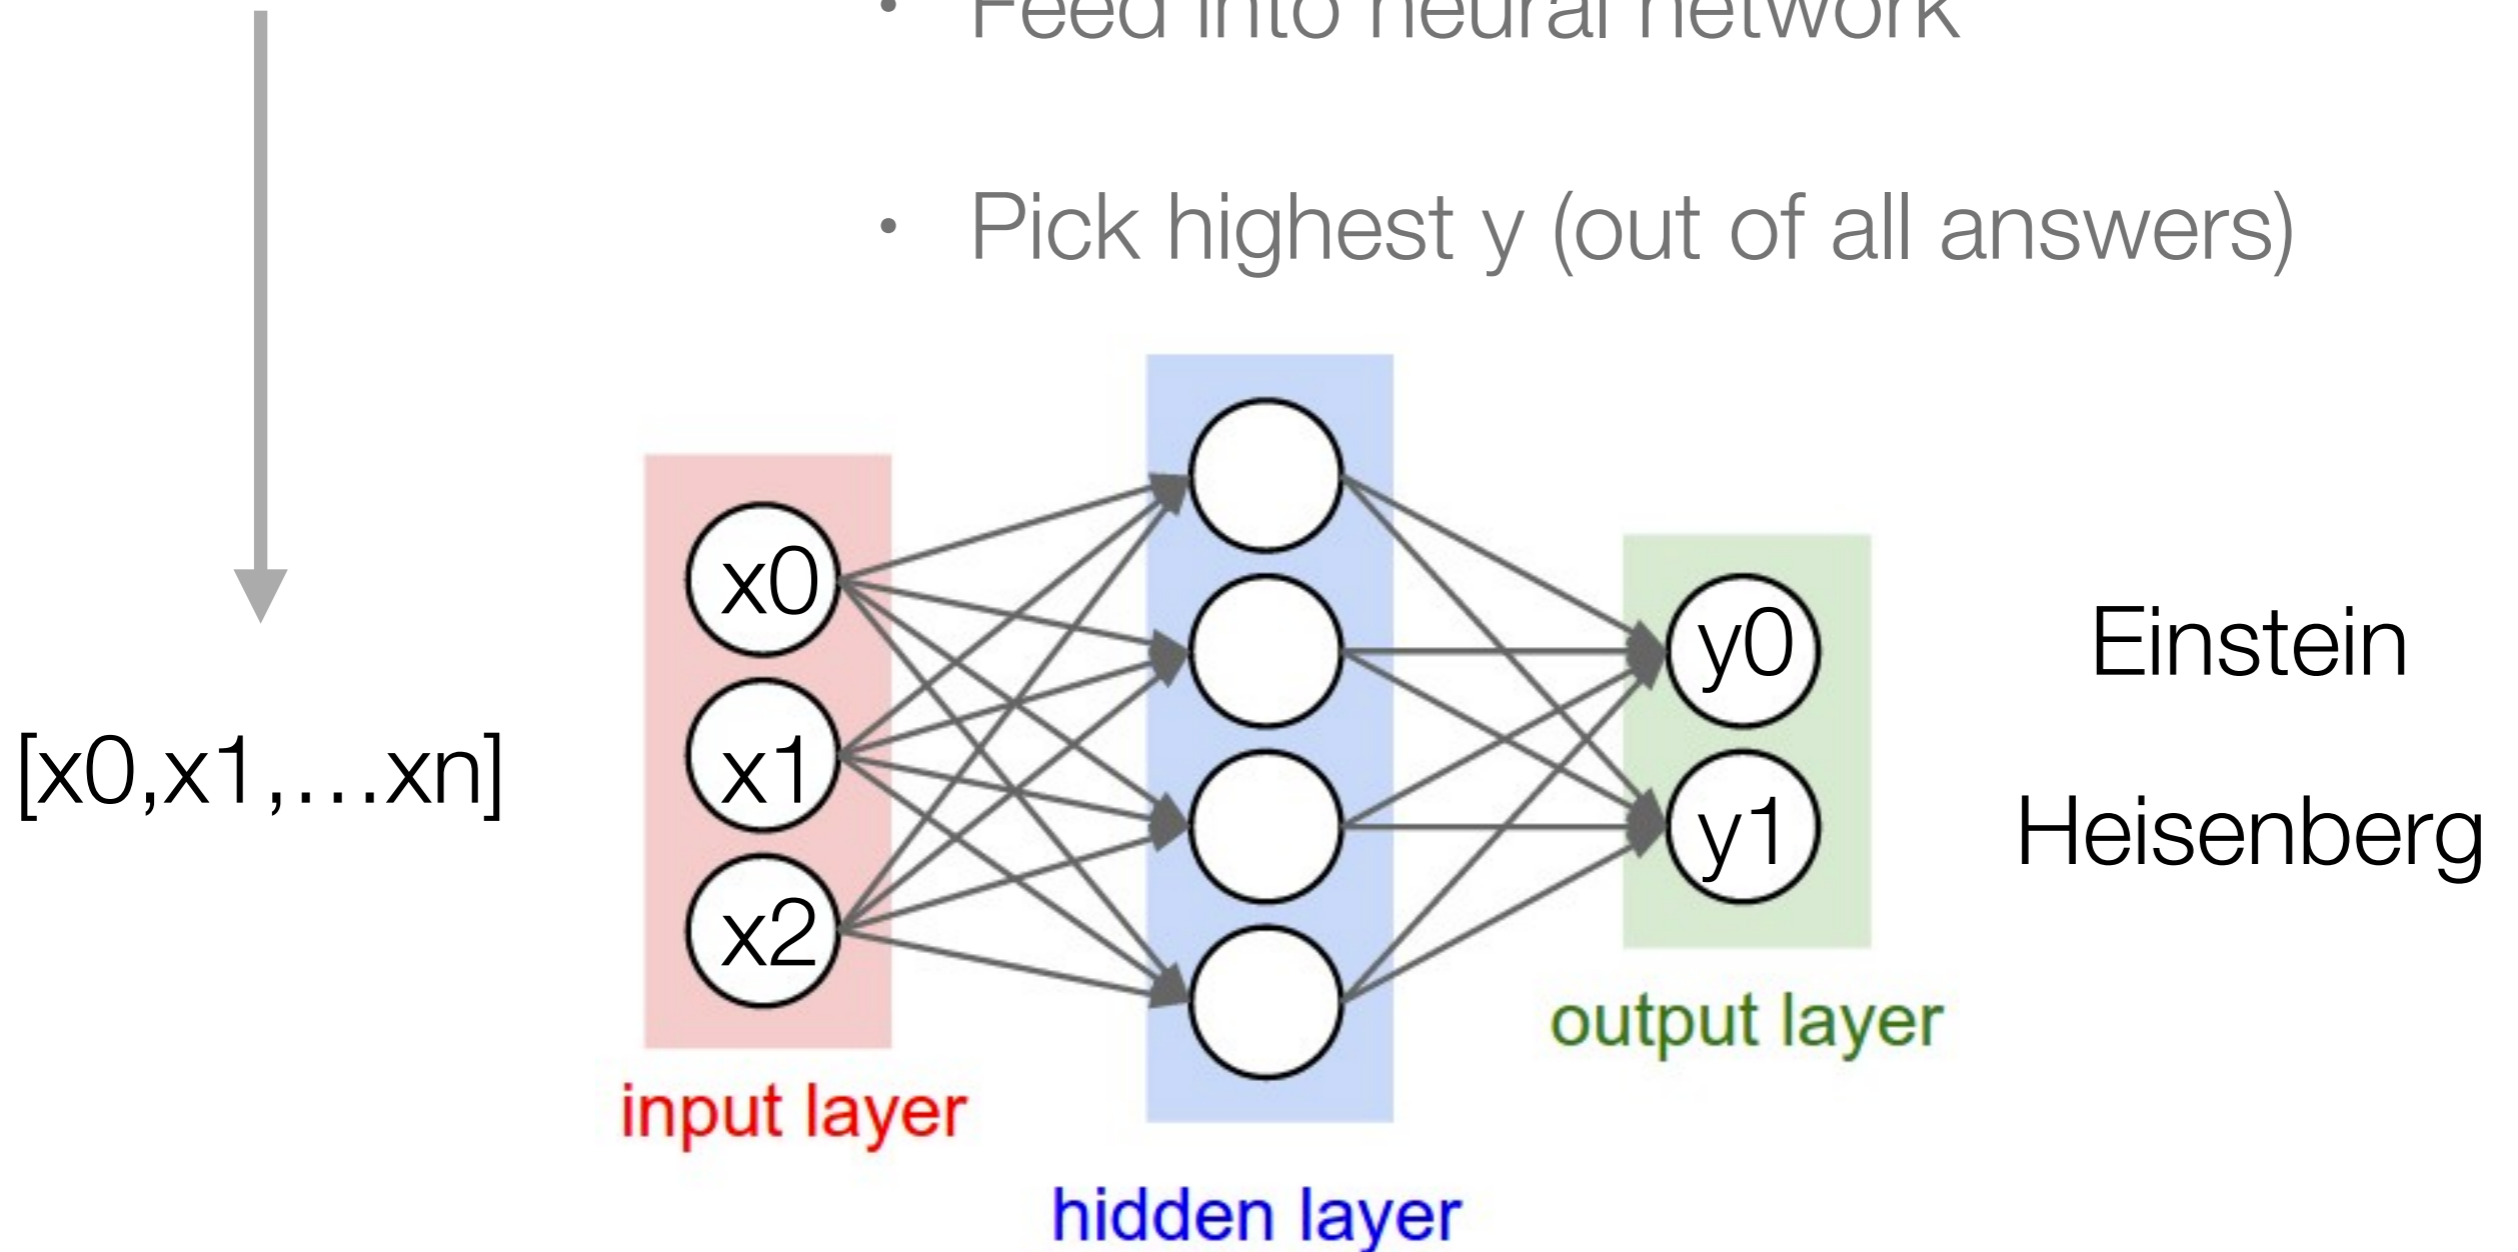
# Recurrent Neural Network



# Generating Guesses

Einstein Question

- Vectorize question using average
- Feed into neural network
- Pick highest  $y$  (out of all answers)



# Softmax

---

- LR is binary
- QB has ~2,000 to ~10,000 answers
- Softmax is a generalization given K classes

$$\sigma(\mathbf{z})_j = \frac{e^{z_j}}{\sum_{k=1}^K e^{z_k}}$$

# Training Neural Networks

---

- First lets write down the neural network math

$x$  = d dimensional input to network

$z_i$  = hidden layer i

$a_i$  = output of layer i

$W_i$  = weights for layer i

$b_i$  = bias for layer i

$z_1 = xW_1 + b_1$

$a_1 = \tanh(z_1)$

$z_2 = a_1W_2 + b_2$


$a_2 = \hat{y} = \sigma(z_2) = \text{softmax}(z_2)$



# Cross Entropy Loss

---

- Need a way to grade our algorithm's performance
- C classes, N data, y is true label and y-hat is prediction
- When is this minimized?



$$L(y, \hat{y}) = -\frac{1}{N} \sum_{n \in N} \sum_{i \in C} y_{n,i} \log \hat{y}_{n,i}$$

Training: How can we minimize the loss?

---

$$z_1 = xW_1 + b_1$$

$$a_1 = \tanh(z_1)$$

$$z_2 = a_1W_2 + b_2$$

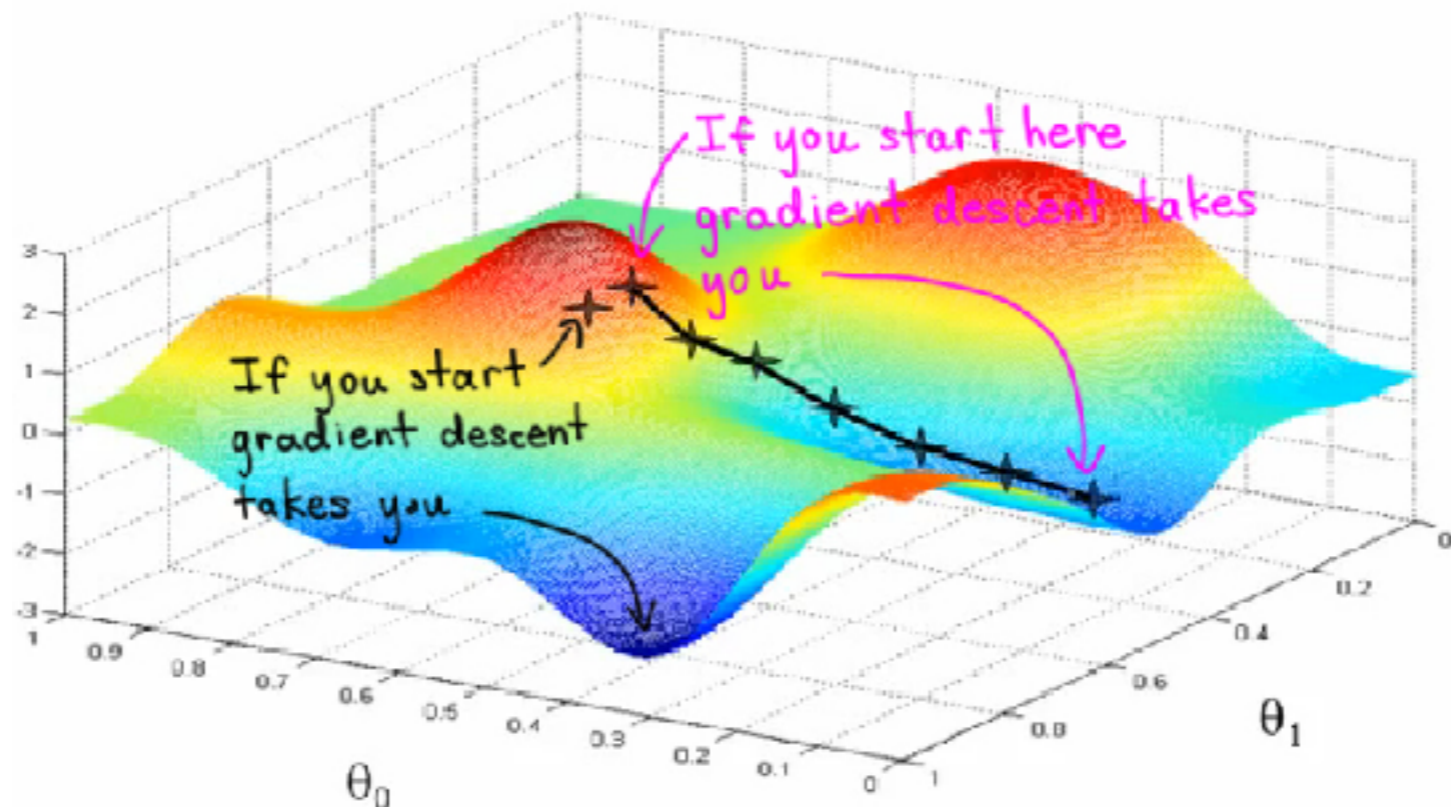
$$a_2 = \hat{y} = \sigma(z_2) = \textit{softmax}(z_2)$$

$$L(y, \hat{y}) = -\frac{1}{N} \sum_{n \in N} \sum_{i \in C} y_{n,i} \log \hat{y}_{n,i}$$

# Training

---

- Evaluate data  $x$  and get a loss from  $L$
- Calculate the gradient of the loss function (backprop)
- Update the weights using the gradient and loss

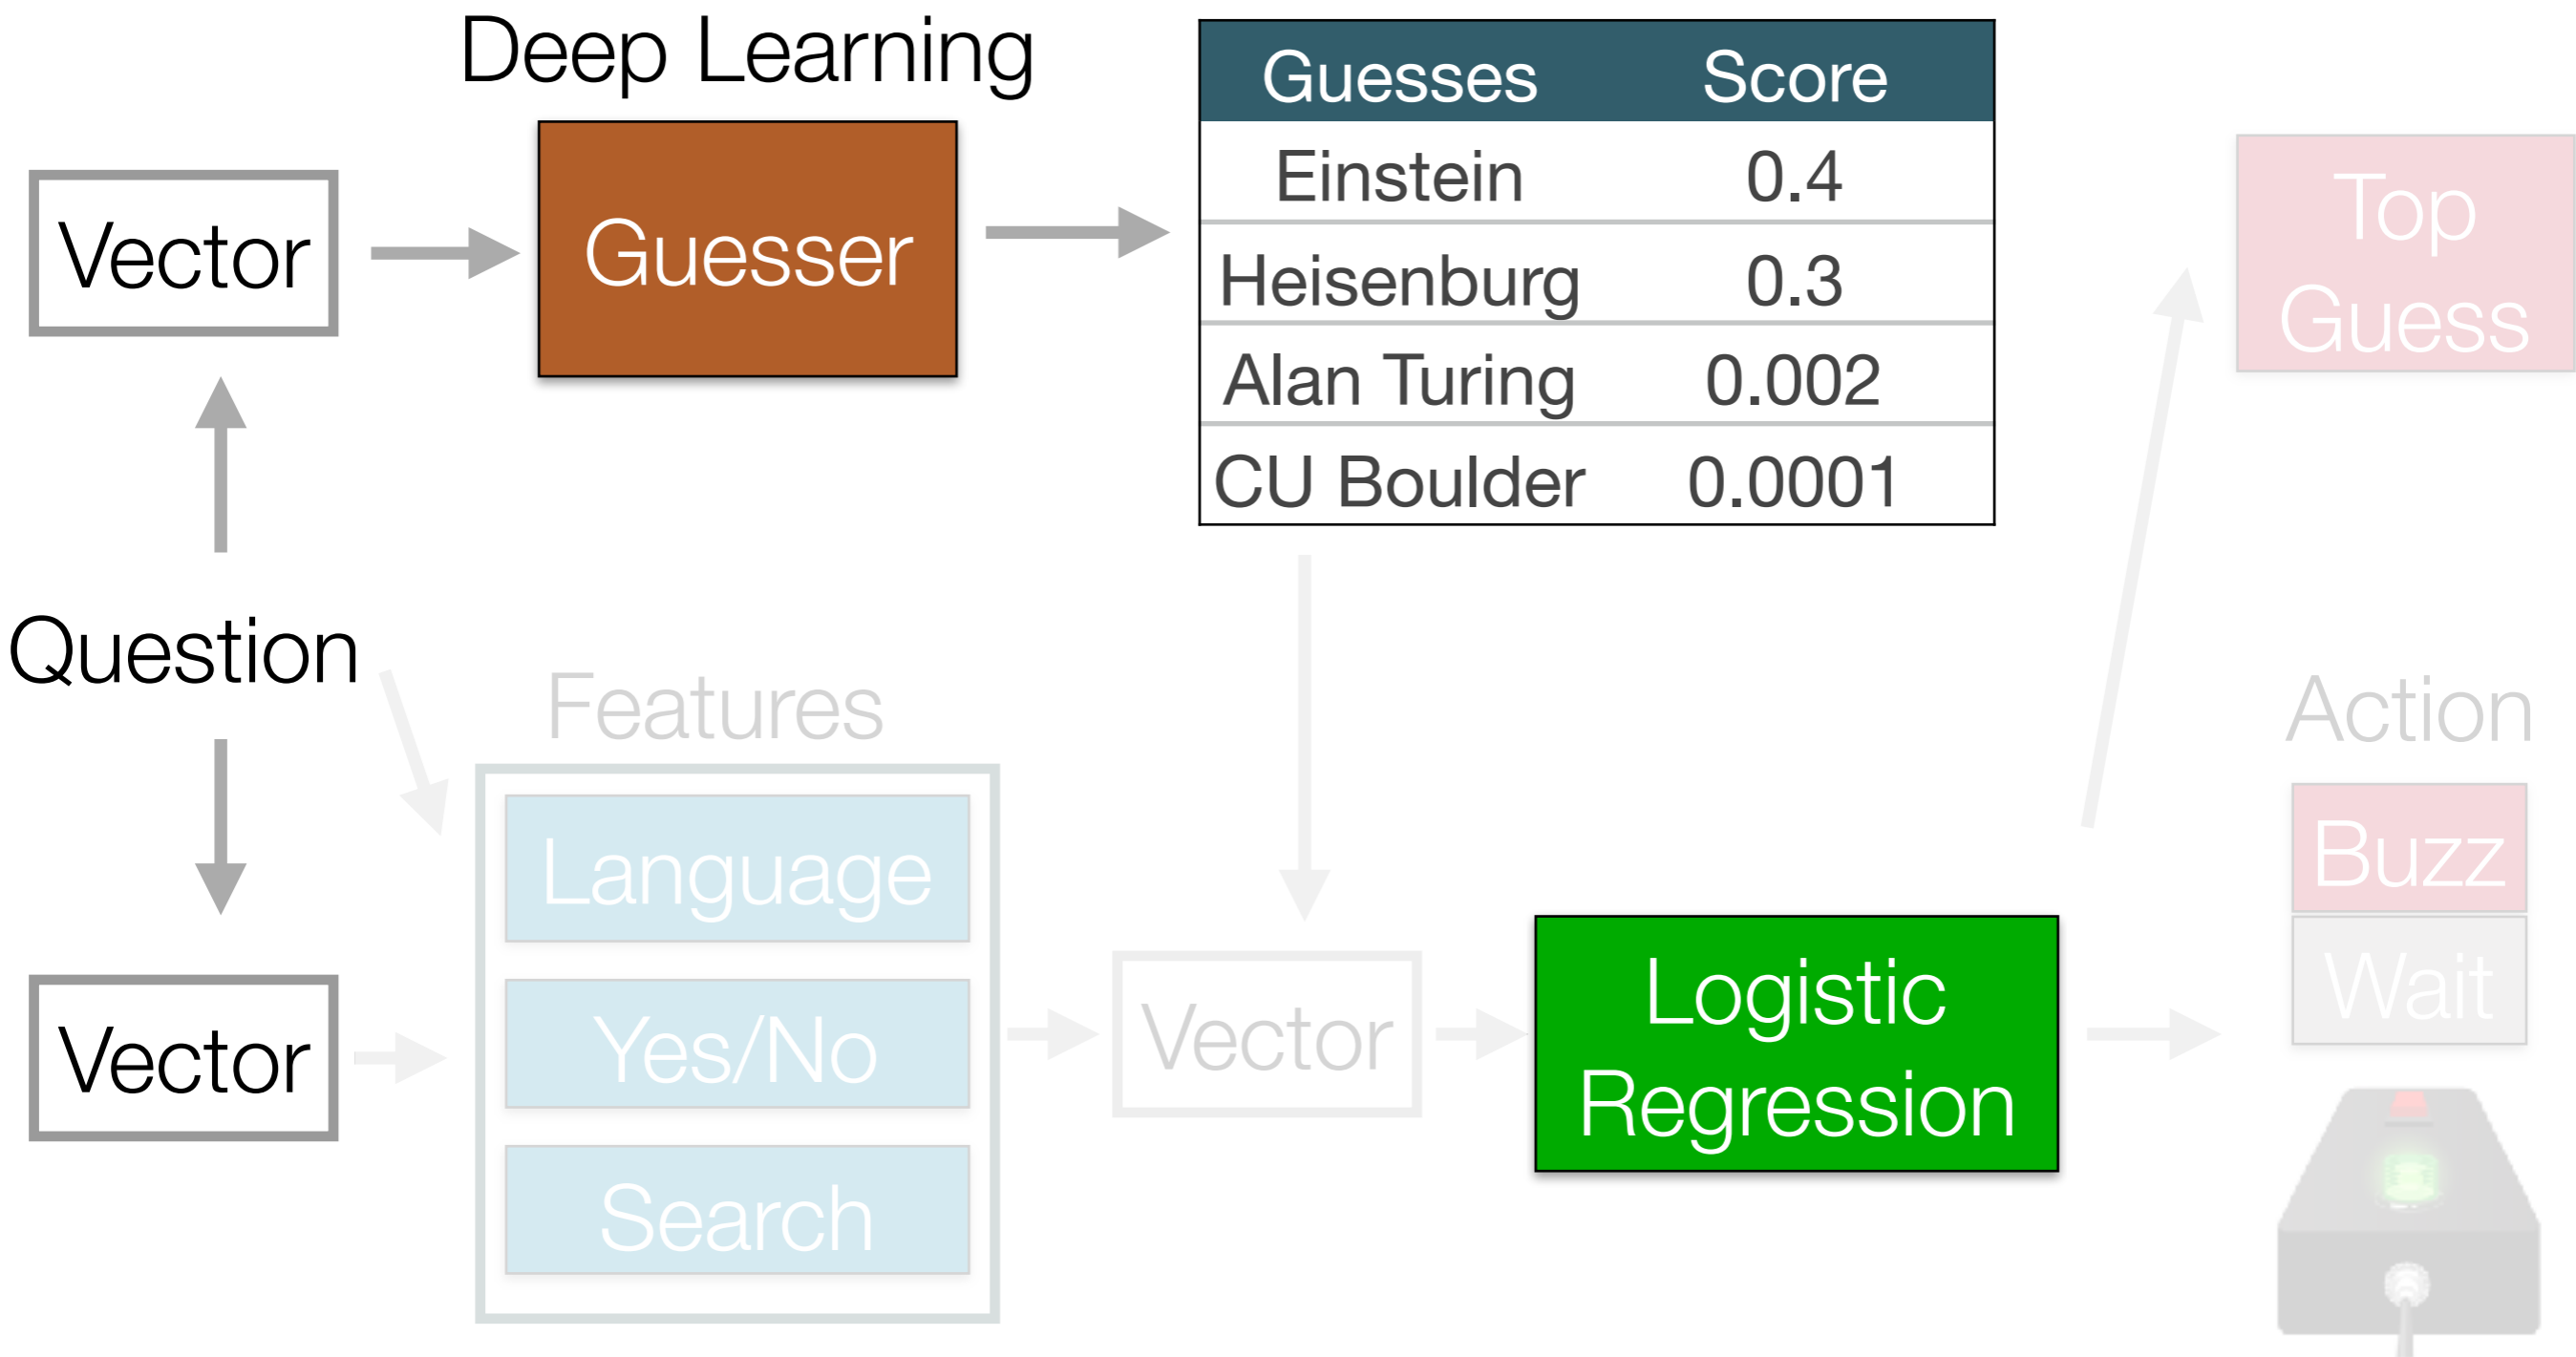


# Summary

---

- Deep Learning used to go from embedding to answer
- Softmax generalizes binary loss
- Cross Entropy Loss
- There is a lot more going on here...

# QANTA Overview



Reranking Guesses with Other Features

# QANTA Features

---

- Language Model
- Binary Features: gender, answer present,...
- Wikipedia “lookup”
- Category Features

# Language Model

---

- Probabilistic Model
- How likely is the sequence of words?
- Assuming a bag of words model (order doesn't matter)
- Use Markov Property, why?

Unigram

$$P(w_0, \dots, w_n) = P(w_0) \cdot \dots \cdot P(w_n)$$



# Language Model

---

- How does this help quiz bowl?
- Condition on what guess is being considered

$$P(w_0, \dots, w_n, guess) = P(w_0, \dots, w_n | guess)P(guess)$$

$$\max_i P(w_0, \dots, w_n | guess_i)$$

↑  
Score

# Classification Features

---

- Is the answer present in the question?
- Is the question about a male or female?
- What is the question category?

# Wikipedia

---

- Idea: “Wikify” the text and look for potential answer
- With Leo Szilard, he invented a doubly-eponymous refrigerator with no moving parts. He did not take interaction with neighbors into account when formulating his theory of heat capacity, so Debye adjusted the theory for low temperatures. His summation convention automatically sums repeated indices in tensor products. His name is attached to the A and B coefficients for spontaneous and stimulated emission, the subject of one of his multiple groundbreaking 1905 papers. He further developed the model of statistics sent to him by Bose to describe particles with integer spin. For 10 points, who is this German physicist best known for formulating the special and general theories of relativity?
- Albert Einstein


# When to Guess?

- Given  $X$  predict  $Y$ 
  - $X$ : all data about the question
  - $Y$ : guess or wait
- Example: when should I sell my stock?
- Quiz Bowl: when should we buzz?



# Human Guesses

---




# What does this say about factors to balance?

---

- How accurate are you?
- How much will your accuracy increase?
- How accurate is your opponent?
- How much will your opponent's accuracy increase?
- How aggressive is your opponent?


We could mimic humans?

---



We could learn from experience like humans

---





# What is important?

---

- Need to explore possible options
- Need to exploit good sequences of actions
- Need to develop a strategy for playing the game (policy)
- How can we learn this?
  - Replay many games
  - Learn by trial and error

# Google DeepMind playing Atari

---



# Reinforcement Learning

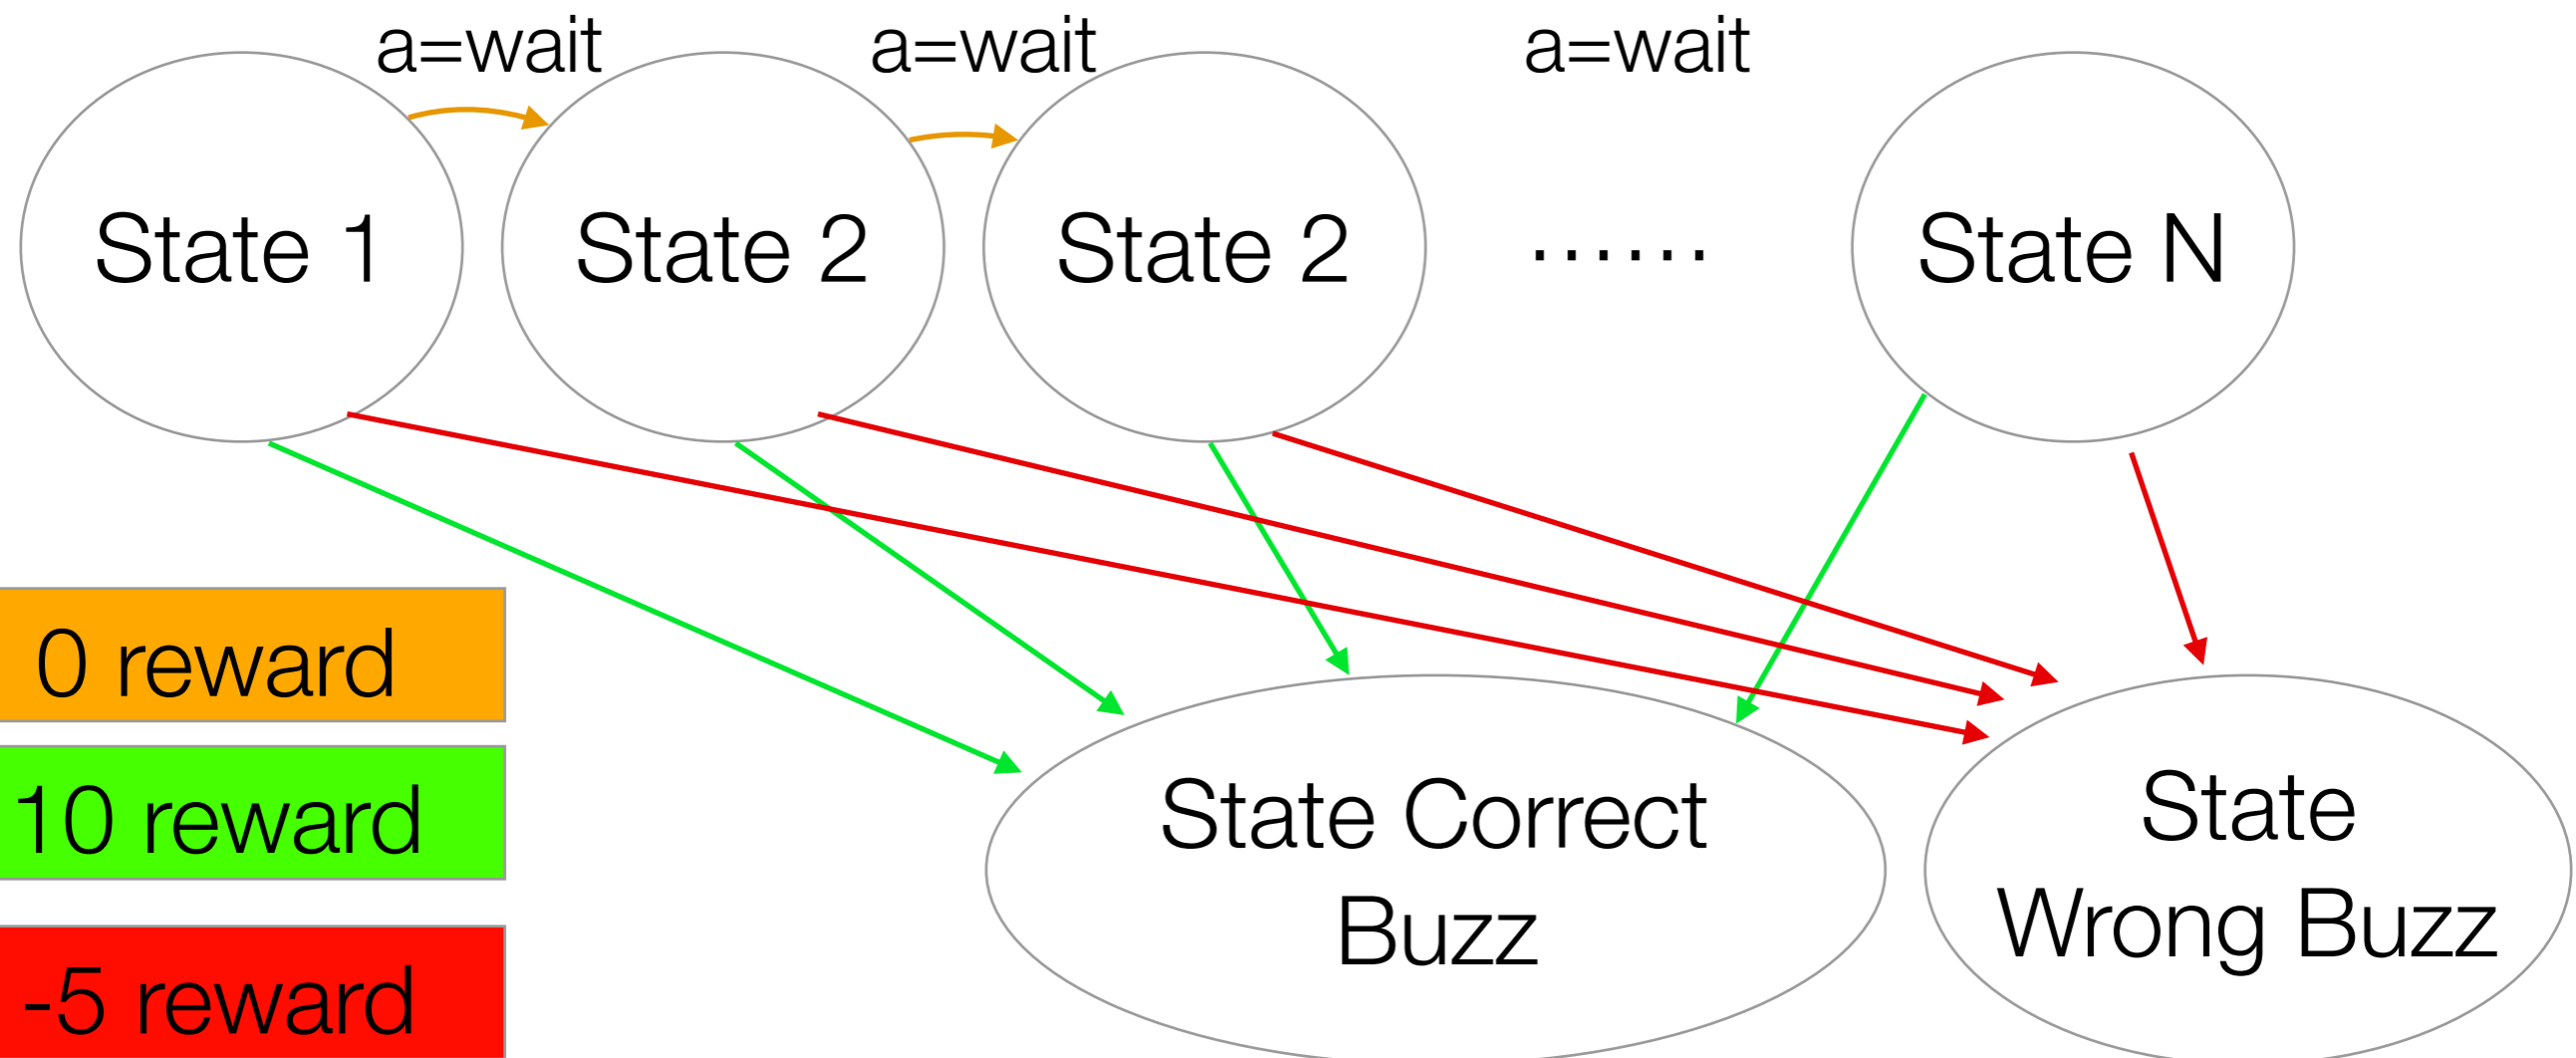
---

- S: set of environmental states
  - sentence/word position, information about question
- A: set of actions -> Buzz or Wait
- Rules determine rewards from going one state to another
  - Buzzing incorrectly loses points
  - Waiting too long loses points
- Examples: chess, checkers, ...

# RL for Buzzing

---

- N word sentence, State=all we know about question
- Is this really all the states? How do you assign credit for reward?



# Q Learning

---

- Method for evaluating actions given current state
- Iterative algorithm, Q represents our model
- Lets unpack this!

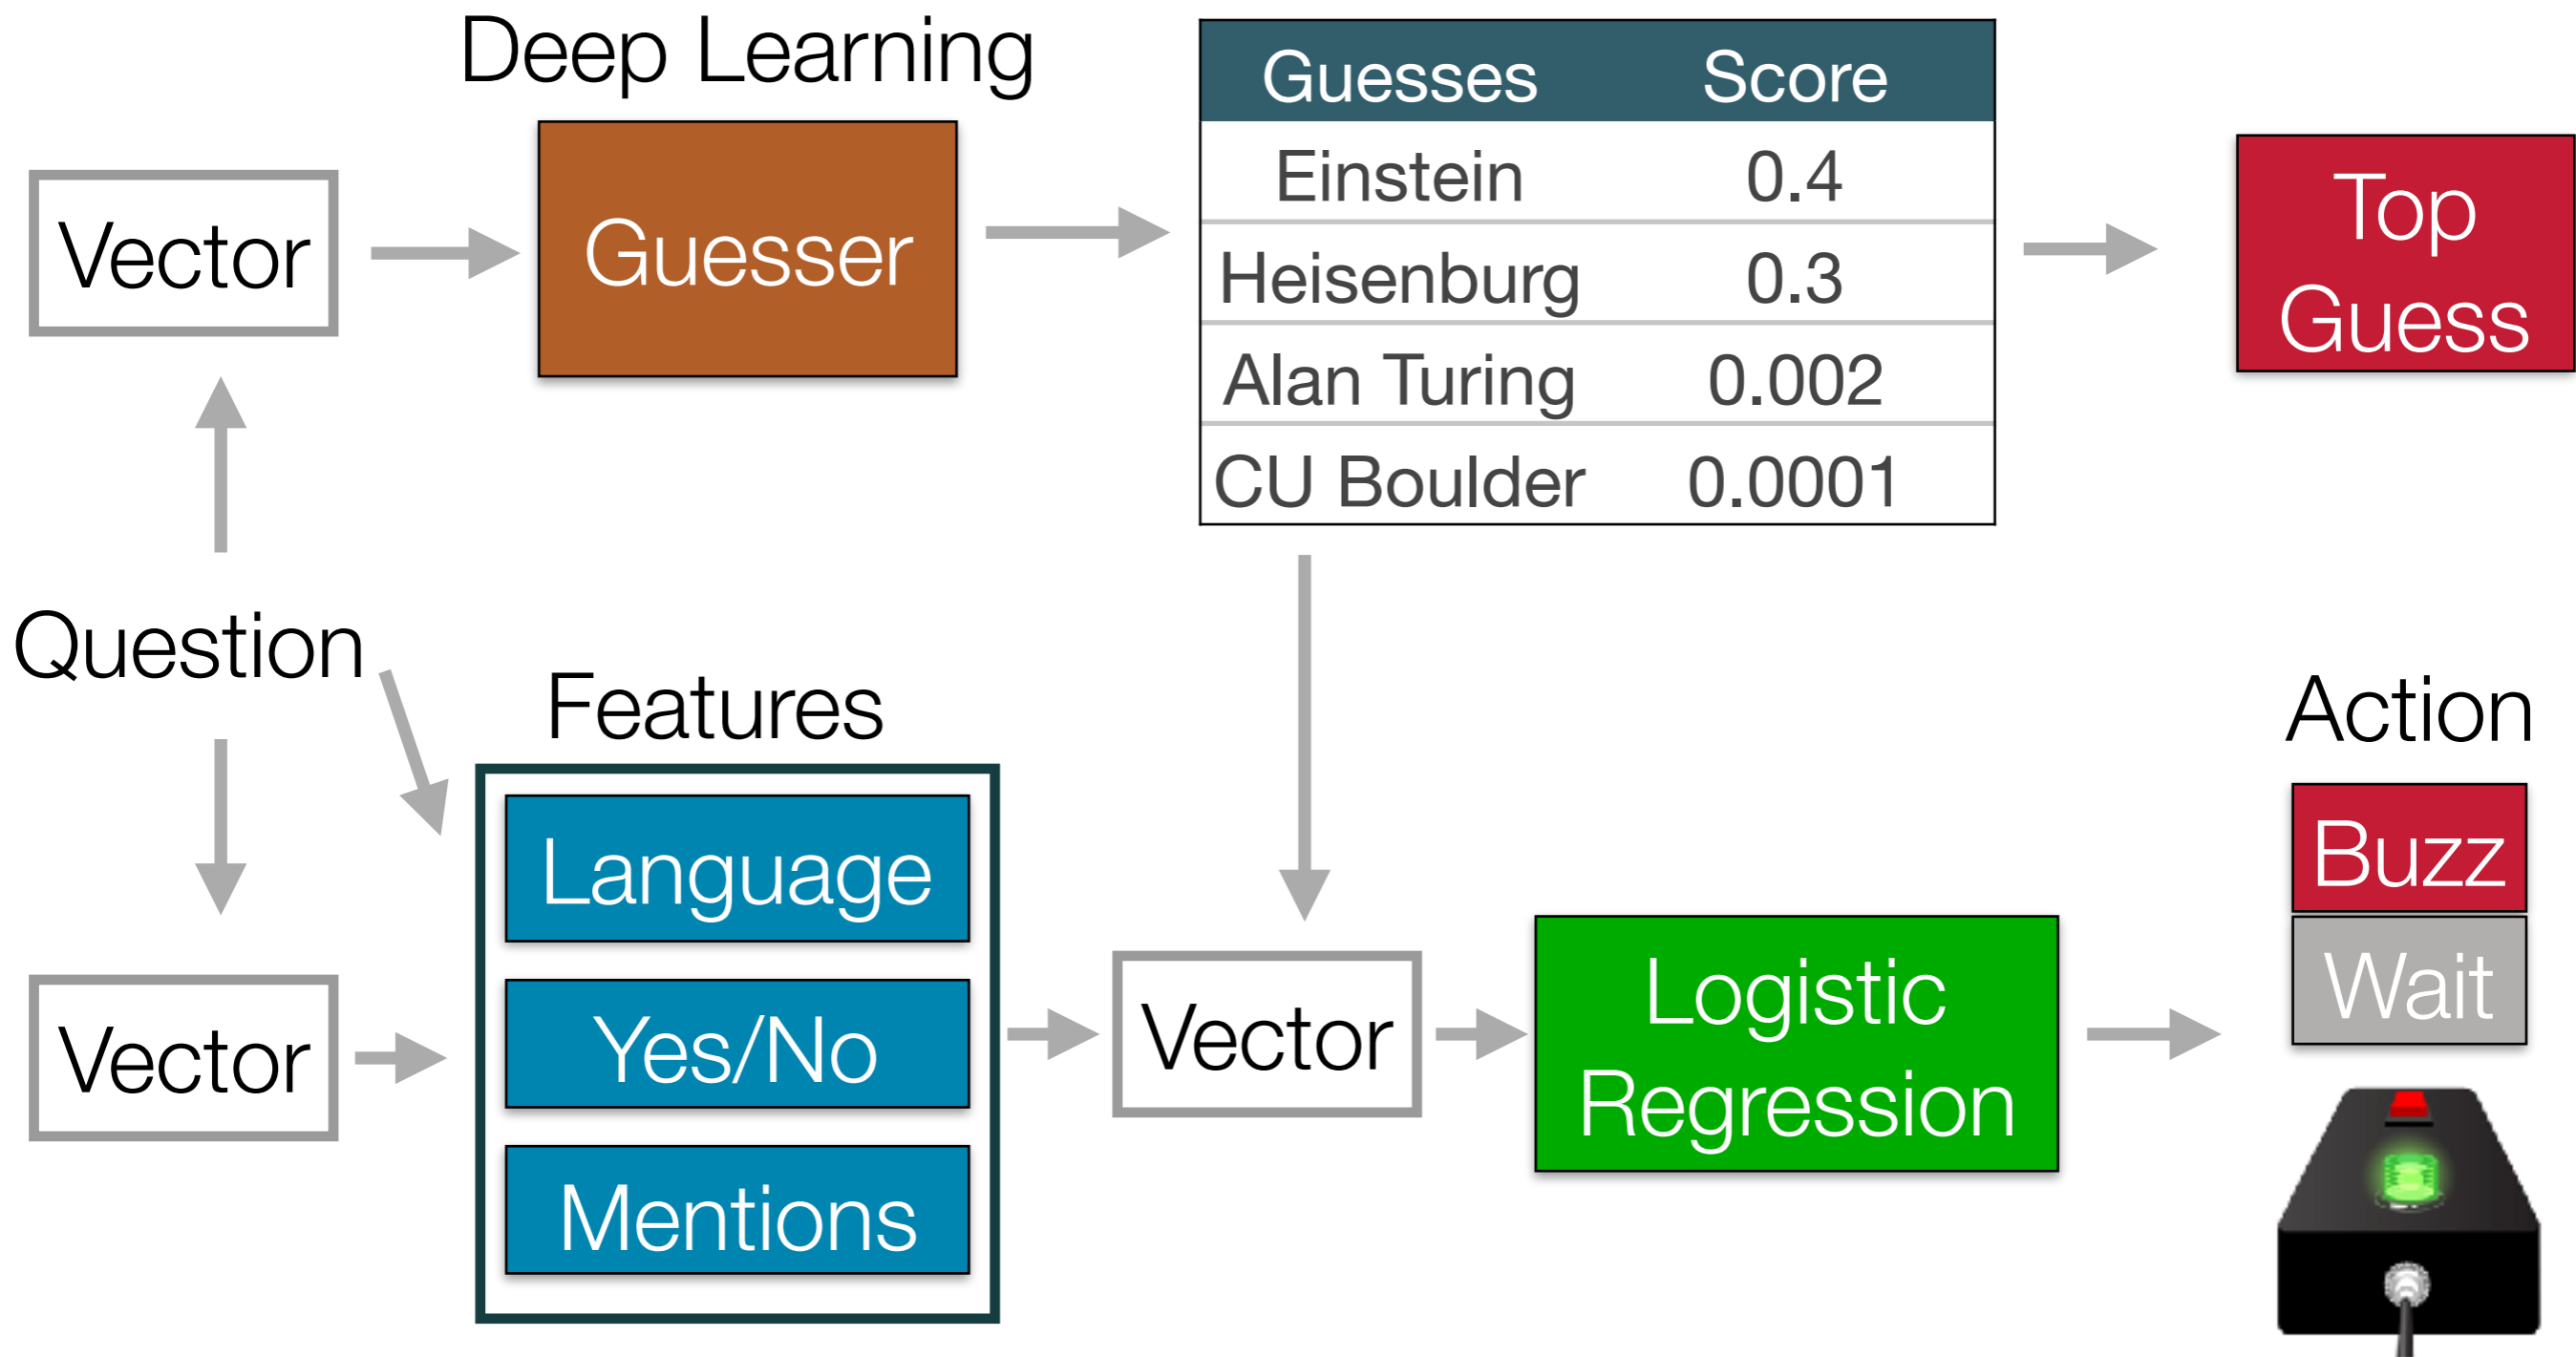
$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left( r_{t+1} + \gamma \cdot \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \right)$$

# Deep Q Learning

---

- Q is a lookup table, which means its  $S \times A$
- What if S is extremely large like in Atari and QB?
  - $84 \times 84$ , 256 gray levels, consider last 4 frames
  - $256^{(84 \times 84 \times 4)} = 10^{67970}$  states.... thats big
- Instead, let Q be a machine learned model
- Deep Q Learning: model is deep neural network

# QANTA Overview



# Future Work

---


- QANTA
  - Wikipedia network properties for guessing
  - Knowledge bases and relational queries
  - Reinforcement Learning and Opponent Modeling
- Alexa Prize: conversational AI



# Thanks!

---

- [NSF: Bayesian Thinking on Your Feet](#)
- [github.com/EntilZha](#)
- QANTA: [github.com/Pinafore/qb](#)
- [UC Berkeley CS188 Course Materials](#)
- About Me: [pedrorodriguez.io](#)
- Contact: [p.rodriguez@colorado.edu](mailto:p.rodriguez@colorado.edu)



# References

---

- Besting the Quiz Master: Crowdsourcing Incremental Classification Games, EMNLP 2012
- A Neural Network for Factoid Question Answering over Paragraphs, EMNLP 2014
- Deep Unordered Composition Rivals Syntactic Methods for Text Classification, ACL 2015
- Removing the training wheels: A coreference dataset that entertains humans and challenges computers, ACL 2015