

# Machine Learning for Question and Answer Understanding at Quora

Nikhil Dandekar

Quora: Nikhil-Dandekar

Twitter: @nikhilbd

**Quora**

# What is Quora?

Quora is a platform to ask questions, get useful answers, and share what you know with the world.

---



Question asked · Yosemite National Park · 52m

## What is it like to hike Half Dome in Yosemite?

12 Answers


[Answer](#)

[Pass](#)

[Follow](#) 39 [Downvote](#) ⋮

 Ian Goodfellow wrote this · Deep Learning · Jun 14

### I want to pursue PhD in deep learning in Canada. Will I be able to pursue it if I have a wife and 2+ year old kid? I'm asking money wise.

 Ian Goodfellow, Lead author of the Deep Learning textbook:  
<http://www.deeplearningbook.org>  
Written Jun 14

It depends a lot on which university / which city, but in general, yes, it should be possible in terms of money. General advice:

- Apply for lots of external fellowships. These can usually boost your i...

[\(more\)](#)

[Upvote](#) 12


[Downvote](#) [Comment](#)



 Loring Chien wrote this · Sun

### Why can't Apple just make a phone with a battery that lasts one day?

 Loring Chien, Electrical engineer  
Written Sun



[More Stories](#)

 Marc Bodnick a... 2h

### What major institutional investors (outside of VCs) are active users on Twitter?

[Answer](#)

[Pass](#)

[Follow](#) 2 ⋮

Answer written · Computer Science · 22m

### What does the notion of state and memory mean in the realm of computer science?

 Daniel Dilger  
Written 22m ago

Consider "computer science" to be the study of algorithms. This broad definition would also include the theory of what problems are computable and the spac...

[Upvote](#)

[Downvote](#) [Read More](#) ⋮

 Abhinav Sharma upvoted this answer from 2015 · Sat

### Flying: What things should I look out for as a passenger on an airline flight?



Read



Answer



Notifications



You



AI Takeover   Existential Risks   Threats   Disruptive Technology   Future of Technology

Artificial General Intelligence   Artificial Intelligence   The Human Race and Condition

## Is A.I. an existential threat to humanity?

Promoted by interviewing.io

### You're good at machine learning, but what about interviewing?

Free interview practice with top engineers. If you make the cut, you get fast-tracked at top companies!

Learn more at interviewing.io [↗](#)

### 100+ Answers



Yoshua Bengio, Head of Montreal Institute for Learning Algorithms, Professor @ U. Montreal

Answered Jan 19, 2016 · Upvoted by Zeeshan Zia, PhD in Computer Vision and Machine Learning and Alberto Bietti, PhD student in machine learning

Originally Answered: Yoshua Bengio: Is AI an existential threat to humanity?

There is a lot of uncertainty about the kind of AI we may one day reach that would achieve human-level intelligence or possibly more. Even though I am an optimist, because of that uncertainty, I cannot deny that things we do not want to happen could happen. That is why I have signed the Future of Life letter for a robust and beneficial AI (see [FLI - Future of Life Institute open letter](#)). [↗](#)

However, keep in mind that such potential dangers are quite far from us right now. There has been a tendency in the press to blow this up out of proportion to attract attention. It could hurt AI research. Instead, we should encourage

Recurrent Neural Networks (RNNs)   Natural Language Processing

## How does an attention mechanism work in deep learning?

Promoted by Interana

### Top analytics skills you need to run your data team.

Though marketing is somewhere between an art and a science, analytics can still give you an edge.

Read more



### 6 ANSWERS



Nikhil Dandekar, leads engineers working on NLP and ML problems at Quora

Answered May 10

At a high-level, an attention mechanism enables your neural network to focus on relevant parts of your input more than the irrelevant parts when doing a prediction task.

Attention aims to capture something we do as humans. E.g. if you are manually translating a long sentence from

# Talk outline

- Data at Quora
- Lifecycle of a question
- Deep dive: Automatic question correction
- Other question and answer understanding examples

# Data at Quora

# Questions and Answers

## How related are artificial intelligence and data science?

Last followed 4m ago

Follow 4 Downvote



## How does one start learning artificial intelligence from scratch?

Last followed 5m ago · 4 Answers

Follow 12 Downvote



## How can I do data preprocessing during image classification with artificial neural networks when images have different sizes?

Last followed 1h ago · 3 Answers

Follow 2 Downvote



## Why have the social sciences basically ignored algorithmic probability's applicability to unified models of society?

<https://herox.com/SocialCausalityPrize/>

Last followed 1h ago · 1 Answer

Follow 4 Downvote



## Which courses are best to learn about hardware design (silicon) for artificial intelligence and deep learning?

Last followed 15m ago · 1 Answer

Follow 23 Downvote



## What are the advantages of logistic regression over decision trees? Are there any cases where it's better to use logistic regression instead of decision trees?



Claudia Perlich, Chief Scientist Dstillery, Adjunct Professor at NYU

Answered Jun 5 · Featured on HuffPost and 3 more · Upvoted by Ayush Sharma, Masters Digital Signal Processing & Machine Learning, Georgia Institute of Technology (2015) and Naran Bayanbat, MSCS with focus in machine learning

First of, you need to be clear what exactly you mean by advantages. People have argued the relative benefits of trees vs. logistic regression in the context of [interpretability](#), robustness, etc.

But let's assume for now that all you care about is out of sample predictive performance. Again, you may need to specify what kind of predictive performance you need: accuracy, ranking, probability estimation. In short: all things equal, trees might have a leg up on accuracy whereas logistic might be better at ranking and probability estimation.



Kah Seng Tay, Jay Wacker, Alberto Bietti, and 21 more upvoted this



Andrew Ng, Co-founder of Coursera; Adjunct Professor of Stanford

Answered Jan 29, 2016 · Featured on HuffPost and 2 more · Upvoted by Alexandru Rosianu, M.Eng. Computer Science & Artificial Intelligence, University of Southampton and Rob Cohn, Ph. D. Artificial Intelligence, University of Michigan (2016)

Originally Answered: Andrew Ng: Is AI an existential threat to humanity?

Worrying about AI evil superintelligence today is like worrying about overpopulation on the planet Mars. We haven't even landed on the planet yet!

AI has made tremendous progress, and I'm wildly optimistic about building a better society that is embedded up and down with machine intelligence. But AI today is still very limited. Almost all the economic and social value of deep learning is still through supervised learning, which is limited by the amount of suitably formatted (i.e., labeled) data. Even though AI is helping hundreds of millions of people already, and is well poised to help hundreds of millions more, I don't see any realistic path to AI threatening humanity.

# Topics



## Recurrent Neural Networks (RNNs)

[Read](#) [Answer](#) [Topic FAQ](#) [Most Viewed Writers](#)

Following Topic | 23.4k | ...

### Feed

Question asked · Recurrent Neural Networks

#### How do you supply input to a recurrent neural network?

Last followed Jun 1 · 2 Answers

[Answer](#) [Pass](#) [Follow](#) 8 [Downvote](#)



Answer · Recurrent Neural Networks

#### Why doesn't CNN suffer from the vanishing gradient problems of RNN?



Angli Liu, works at University of Washington  
Answered Aug 6

They do. RNNs unfolded are deep in common application scenarios, thus prone to severer vanishing gradient problems. For example, when used in language modeling, RNN depth can go as long as the longe... [\(more\)](#)

[Upvote](#) 2 [Downvote](#)



Yair Livne and Xavier Amatriain upvoted this · Recurrent Neural Networks · 2016

#### How similar are the internal parameters of two structurally identical RNN's trained on the same data?



Ian Goodfellow, AI Research Scientist  
Answered Nov 14, 2016

This depends a lot on:

- The way you measure similarity between parameter vectors.
- The training algorithm

### About

Describe the [Recurrent Neural Networks \(RNNs\)](#) topic

579  
Questions

23.4k  
Followers

31  
Edits

### Related Topics



[Long Short-term Memory](#)  
12.6k Followers



[Convolutional Neural Networks](#)  
30.7k Followers



[Artificial Neural Networks](#)  
198.4k Followers



[Neural Machine Translation](#)  
1.2k Followers



[Deep Learning](#)  
217.8k Followers





## Daphne Koller

President and co-founder of Coursera

Following 11.1k Turn On Notifications Ask Question

### Feeds

Answers 7

Questions 0

Activity

Posts 0

Blogs 0

Followers 11,124

Following 0

Topics 0

Edits 7

### 7 Answers

Most Recent / 30-Day Views

#### How valuable are Coursera certificates and would they help me find jobs easier if I put them on my resume?

Daphne Koller, President and co-founder of Coursera  
Answered Apr 27, 2016 · Upvoted by Mark Pan, worked at Coursera

A Coursera Certificate is certainly a valuable addition to a resume. In the Coursera Learner Outcomes Survey [that we published last year](#), we saw that learners were more likely to report tangible ca... [\(more\)](#)

Upvote 390 Downvote [f](#) [t](#) [e](#) [...](#)

#### Do you have any data regarding the success obtained by Coursera students (in regards of salary) after finishing a course?

Daphne Koller, President and co-founder of Coursera  
Answered Apr 27, 2016 · Upvoted by Abhinav Sharma, I used to work at Coursera

We don't have precise quantitative data on salary changes, but we do have strong evidence that many learners who complete courses see a salary benefit. Results from last year's Coursera Learner Outcomes Survey [... \(more\)](#)

Upvote 90 Downvote [f](#) [t](#) [e](#) [...](#)

#### When will Coursera be offering degrees?

Daphne Koller, President and co-founder of Coursera  
Answered Apr 27, 2016 · Upvoted by Abhinav Sharma, I used to work at Coursera

We do now! We've partnered with the University of Illinois at Urbana-Champaign [to offer two all-online, fully-accredited degrees: an MBA program](#) [, which launched in January](#), and a Master of Computer ... [\(more\)](#)

Upvote 277 Downvote [f](#) [t](#) [e](#) [...](#)

#### What benefit do universities get who put courses on Coursera?

Daphne Koller, President and co-founder of Coursera  
Answered Apr 27, 2016 · Featured on HuffPost · Upvoted by Abhinav Sharma, I used to work at Coursera and Mark Pan, worked at Coursera

### Credentials & Highlights

More

#### 18 Followers You Know

Jay Wacker, Tom Dong (董玉冰), Xavier Amatriain, Wael Ghonim, and 14 others you follow

3.4m answer views  
5.5k this month

Session Host  
Apr 2016

Published Writer  
HuffPost, Quora's Twitter, and Quora Sessions's Twitter



## Ian Goodfellow

AI Research Scientist

I'm a staff research scientist at Google Brain, where I lead a research group studying adversarial techniques in AI (adversarial examples, generative adversarial networks, etc.). Together with Yosh... [\(more\)](#)

Following 20.3k Turn On Notifications Ask Question

### Feeds

Answers 235

Questions 4

Activity

Posts 0

Blogs 0

Followers 20,387

Following 113

Topics 77

Edits 355

### 235 Answers

Most Recent / 30-Day Views

#### What majors are recommended from B.S to M.S to model something like Generative Adversarial Network or Black-Scholes Merton Model as goal?

Ian Goodfellow, I invented generative adversarial networks  
Answered Sep 6

I invented generative adversarial networks and I majored in computer science for both my B.S. and M.S. For the M.S., I had a concentration in AI.

2.7k Views · 82 Upvotes · Answer requested by Sean Kim

Upvote 82 Downvote [f](#) [t](#) [e](#) [...](#)

Add a comment... Recommended All

#### What are some good references for adversarial attack and defence?

Ian Goodfellow, AI Research Scientist  
Answered Aug 11

Here is a reading list that Alexey Kurakin and I put together for the adversarial example contest: [NIPS 2017: Non-targeted Adversarial Attack](#)

3.3k Views · 34 Upvotes · Answer requested by Yilun (Tom) Zhang

Upvote 34 Downvote [f](#) [t](#) [e](#) [...](#)

Add a comment... Recommended All

#### Does solving for generating targeted adversarial attacks also solve for generating non-targeted adversarial attacks?

Ian Goodfellow, Lead author of the Deep Learning textbook:  
<http://www.deeplearningbook.org>  
Answered Jul 17 · Upvoted by Sadid Hasan, Ph.D Computational Linguistics & Machine Learning, University of Lethbridge (2013)

### Credentials & Highlights

More

29 Followers You Know  
Adrien Lucas Ecoffet, Jay Wacker, Kapil Dhawan, Shreyes Seshasai, and 25 others you follow

Staff Research Scientist at Google Brain  
2017-present

Studied at Stanford University

Lives in Mountain View, CA

1.6m answer views  
82.1k this month

Session Host  
Aug 2016 and Jul 2017

Published Writer  
HuffPost, Forbes, and 3 more

### Knows About

Deep Learning  
113 answers

Artificial Intelligence  
62 answers

Generative Adversarial Networks  
43 answers

Generative Models (statistics)  
15 answers

OpenAI (company)  
11 answers

View More

# Other data



**Alberto Bietti**, PhD student in machine learning. Former ML engineer

Sep 24 · 6 upvotes including Eric Jang and Nikhil Garg

FYI, two recent papers have gotten much closer to justifying the results on random vs non-random generalization results, by finding relevant “complexity” measures of the learned neural nets and proving generalization bounds based on those measures. They show empirically that those bounds are rele... [\(more\)](#)

[Reply](#) · [Upvote](#) · [Downvote](#) · [Report](#)



**Brando Miranda**

Sep 24 · 1 upvote

Hey Eric, sorry if its easy, but I was wondering if you minded expanding what you meant with:

“But if a deep net is *also* capable of memorizing random data, it suggests that the generalization capabilities are not entirely explained by the inductive biases of the model (e.g. the convolutional/pooli... [\(more\)](#)

[Reply](#) · [Upvote](#) · [Downvote](#) · [Report](#)



**Eric Jang**, Research Engineer at Google Brain

Sep 24 · 3 upvotes including Brando Miranda

> isn't it expected given such overparametrization model despite the inductive biases used?

Yes, I agree that the result is somewhat expected. I don't think the paper made any arguments that radically opposed to this intuition, but it was important that they demonstrated the experiments showing th... [\(more\)](#)

[Reply](#) · [Upvote](#) · [Downvote](#) · [Report](#)

**Brando Miranda**: After thinking a little bit more about the paragraph you ...

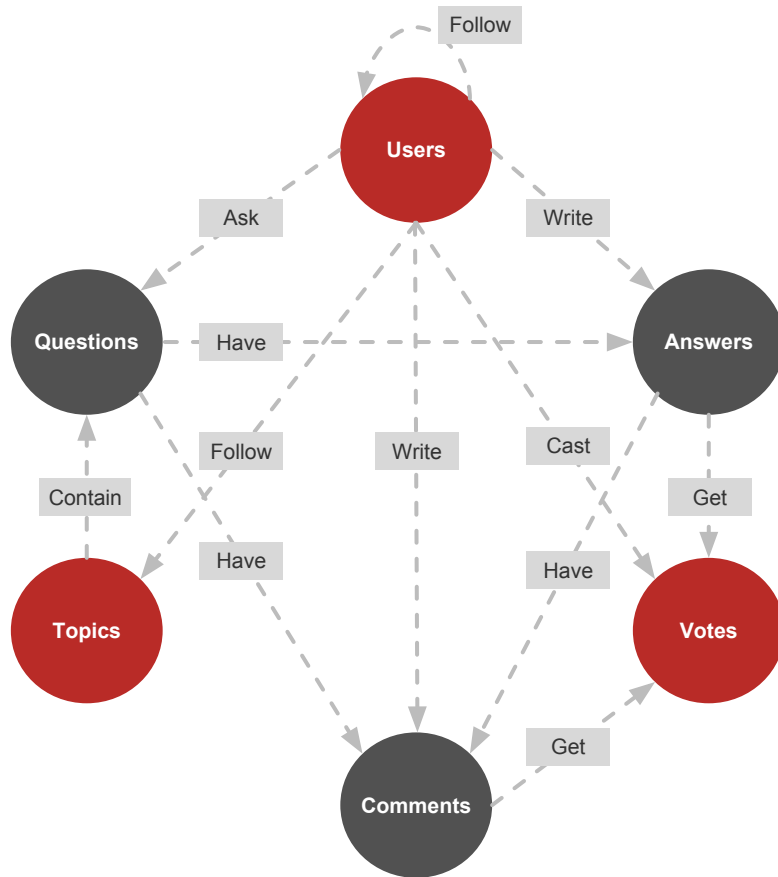
## Answer Wiki

[Marketing's](#) role is one of demand management.

"Create demand" ([Shrikant Latkar](#)), "demand generation" ([Daniel Woodard](#)) and "customer acquisition" ([Anon User](#)) are all terms that get invoked. A couple of answers imagine the role of marketing as a "kickstart" ([Naga Chokkanathan](#) and [Nemo Chu](#)) for a extended range of follow-on activities. But simply put, "Marketing is the management of the customer and their needs" ([Fabio Barca](#)).

The concept of managing demand then gives rise to a range of functions and emphases. Those mentioned include:

- [Targeting, Lead Generation](#) and Sales Support ([Jonathan Rabinowitz](#))
- [Product Development](#) ([Daniel Woodard](#))
- [Brand Building](#) ([Shrikant Latkar](#))
- [Brand Positioning](#) ([Sten Svensson](#))
- [Communications](#) ([Daniel Woodard](#) lists a hierarchy of communication objectives and [Shrikant Latkar](#) a range of communication tools that provide brands with "aircover," including [Socialmedia](#), [Public Relations \(PR\)](#) and [Advertising](#))



**Lots of  
data  
relations**

# Lifecycle of a question

User asks a question

### Question quality

- Adult detection
- Quality classification (high vs low)
- Automatic question correction
- Duplicate question detection and merging
- Spam/abuse detection
- Policy violations
- etc.

Q

Leo C. Stein and 92 more follow this

Magnetism Alternative Energy Physics Edit Topics

### Why do like charges repel and opposite charges attract?

This question previously had details. They are now in a comment.

Answer

Request

Follow 94

...

Why do protons like charges repel each other while attracting electrons and opposite charges attract?

Quality: 4.515

#### Merged Questions

↑ Why do positive and negative charges attract?

Undo Merge

↑ Why is positive charged particles are attracted towards negatively charged particles?

Undo Merge

↑ Why do particles with like charges repel?

Undo Merge

↑ Why do positive and negative charges attracts?

Undo Merge

↑ Why do like charges repel?

Undo Merge

↑ Why do like charges "repel" each other and other way around?

Undo Merge

## Question understanding

- Question-Topic labeling
- Question type classification
- Question locale detection
- Related Questions
- etc.

Software Libraries TensorFlow (software library) Open Source Software  
Machine Learning Google (company) ✎

## Why did Google open source TensorFlow? What's in it for them?

[http://googleresearch.blogspot.com/2015/11/tensorflow-googles-latest-machine\\_9...](http://googleresearch.blogspot.com/2015/11/tensorflow-googles-latest-machine_9...)

### Related Questions

What is Google's motivation for open sourcing the TensorFlow model code for ML / NLP problems such as automatic text summarization?

What is unique about Tensorflow from the other existing deep learning libraries?

How does Google's TensorFlow handle non-SGD based Machine Learning?

What are some great uses of Google TensorFlow?

TensorFlow (open source s/w library): How tensor flow can be used and what are best ways to understand it better?

## Matching questions to writers

- “Request Answers”
- Feed ranking for questions

Big Wall Climbing Aid Climbing Rock Climbing Climbing






### How do multi-pitch climbers carry all the gear?

This question previously had details. They are now in a comment.

Request Follow 1 Comment 1 Downvote



#### Request Answers:

-  **Request From Quora**  
We will distribute this question to writers, and notify you about new answers. +
  -  **Marc Bow**, Journeyman Rock Climber  
8 Answers in Rock Climbing +
  -  **Jared Anwyl**, Climbing since 91  
162 Answers in Rock Climbing +
  -  **Eugene Miya**, Doing it long enough, alpinist.  
206 Answers in Climbing +
-  View More or Search

 Vo Nghi Nguyen asked this · Amazon (company) · Oct 16

### What is it like to buy a house on Amazon?

<https://www.dailydot.com/debug/amazon-house-shipping/>

Last followed Thu

 Answer Follow 3

 Katie Nguyen asked this · Google Apps · Oct 10

### Which companies are currently using Google's recruiting app, Hire?

Last followed Sat · 1 Answer

 Answer Follow 9

 Scott Borden requested answers to this · Military Technology · Thu

### What major military innovations and advances took place in the 17th and 18th centuries?

Last followed Fri · 1 Answer

 Answer Follow 4

Writer writes an answer to a question

### Answer quality

- Answer ranking for questions
- Answer collapsing
- Adult detection
- Spam/abuse detection
- Policy violations
- etc.

### Is AI an existential threat to humanity?

Answer

Request

Follow 615

Comments 5+

Downvote



100+ Answers

Tom Dong (董玉冰), Sumi Kim, Shankar Iyer, and 13 more upvoted this

Yoshua Bengio, Head of Montreal Institute for Learning Algorithms, Professor @ U. Montreal

Answered Jan 19, 2016 · Upvoted by Zeeshan Zia, PhD in Computer Vision and Machine Learning and Alberto Bietti, PhD student in machine learning

### 204 Answers Collapsed (Why?)

JOE  
Answered Aug 22

Contact:-

TROJANLORD100@Gmail.com

Or Text Him Via:-

+1 408 682 0067

For Any Hacking Problems




## Matching answers to readers

- Feed ranking for answers
- Digest emails
- Search ranking
- Visitors coming from Google

 Loring Chien answered this · Water

### I spilled water over my computer keyboard. What should I do?

 Loring Chien, Electrical Engineer for 40 years & IEEE Sr. Life member  
Answered Thu

Unplug the USB cord from the computer. Gently take and turn the keyboard upside down and shake as much water as you can out of it, Place a fan so it blows over the keyboard if possible, or bake gent... [\(more\)](#)

Upvote

Downvote



Answer

### Why are linguists supposed to be descriptive and not prescriptive?

 Daniel Ross, I study Linguistics  
Answered Oct 12 · Upvoted by Steve Rapaport, Linguistics PhD candidate at Edinburgh. Has lived in USA, Sweden, Italy, UK.

The goal of Linguistics is to understand how human language works and

🔍 roger federer

🔍 Search: **roger federer** >

 Topic: **Roger Federer** (tennis player) >

What is **Roger Federer** like in real life? >

Why do most people consider **Roger Federer** to be better than Rafael Nadal? >

Why do most people love **Roger Federer** so much? >

## Other ML applications

- Ads
  - Ads CTR prediction
  - Ads-topic matching
- ML on other content types
  - Comment quality + ranking
  - Answer wiki quality + ranking
- Other recommender systems
  - Users to follow
  - Topics to follow
- Under the hood
  - User understanding signals
  - User-topic affinity
  - User-user affinity
  - User expertise
- ... and more





Promoted by Springboard

**Switch careers to data science. Get a job or your money back.**

Master machine learning, inferential statistics & data storytelling. Work on data projects with a mentor.

Apply Now at [springboard.com](https://springboard.com) 

 Discover new **topics**

 <b>Quora's Technical Infrastructure</b> <a href="#">Follow</a>   4.7k Followed by Alexandr Wang and 14 more	 <b>Network Effects</b> <a href="#">Follow</a>   720 Followed by Nikhil Garg and 5 more	 <b>Restaurants in San Francisco</b> <a href="#">Follow</a>   46.6k Followed by Veni Johanna and 13 more	 <b>Recommendation System</b> <a href="#">Follow</a> Followed by D and 3 more
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------

# Deep dive: Automatic question correction

# “Bad” questions on Quora

- Users often ask questions with grammatical and spelling errors
- Example:
  - Which coin/token is next big thing in crypto currencies? And why?
  - Which coin/token is the next big thing in cryptocurrencies? Why?
- These are well-intentioned questions, but the lack of correct phrasing hurts them
  - Less likely to be answered by experts
  - Harder to catch duplicate questions
  - Can hurt the perception of “quality” of Quora

# “Bad” questions on Quora

- Types of errors in questions
  - Grammatical errors, e.g., “How I can ...”
  - Spelling mistakes
  - Missing preposition or article
  - Wrong/missing punctuation
  - Wrong capitalization
  - etc.
- Can we use Machine Learning to automatically correct these questions?
- Started off as an “offroad” hack-week project
- Since shipped

# Automatic question correction: research

## CoNLL-2014 Shared Task: Grammatical Error Correction

**CoNLL-2014** will continue the CoNLL tradition of having a high profile shared task in natural language processing. This year's shared task will be grammatical error correction, a continuation of the **CoNLL shared task in 2013**. A participating system in this shared task is given short English texts written by non-native speakers of English. The system detects the grammatical errors present in the input texts, and returns the corrected essays. The shared task in 2014 will require a participating system to correct all errors present in an essay (i.e., not restricted to just five error types in 2013). Also, the evaluation metric will be changed to  $F_{0.5}$ , weighting precision twice as much as recall.

## Grammatical Error Correction: Machine Translation and Classifiers

**Alla Rozovskaya**  
Department of Computer Science  
Virginia Tech  
Blacksburg, VA 24060  
alla@vt.edu

**Dan Roth**  
Department of Computer Science  
University of Illinois  
Urbana, IL 61820  
danr@illinois.edu

## System Combination for Grammatical Error Correction

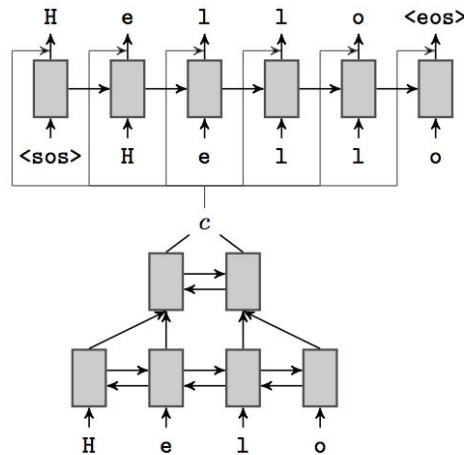
**Raymond Hendy Susanto**    **Peter Phandi**    **Hwee Tou Ng**  
Department of Computer Science  
National University of Singapore  
13 Computing Drive, Singapore 117417  
{raymondhs, peter-p, nght}@comp.nus.edu.sg

## Neural Language Correction with Character-Based Attention

**Ziang Xie, Anand Avati, Naveen Arivazhagan, Dan Jurafsky, Andrew Y. Ng**  
Computer Science Department, Stanford University  
{zxie, avati, naveen67, ang}@cs.stanford.edu, jurafsky@stanford.edu

# Automatic question correction: Model

- We frame this problem similar to the machine translation problem
- Final Model:
  - Multi-level, sequence-to-sequence, character-level GRU with attention





# Stepping back: Neural networks

- At the core: A neuron
- Convert one or more inputs into a single output via this function

$$f\left(\sum_i w_i x_i + b\right)$$


- Objective: Learn the values of weights  $w_i$  given the training data
- Can solve simple ML problems well
- At the core of all the deep learning revolution (and hype)







# Neural networks

- Layers of neurons connecting the inputs to the outputs
- **Training:** Adjust the weights of the network via gradient descent using the backpropagation algorithm
- **Serving:** Given a trained network, predict the output for a new input



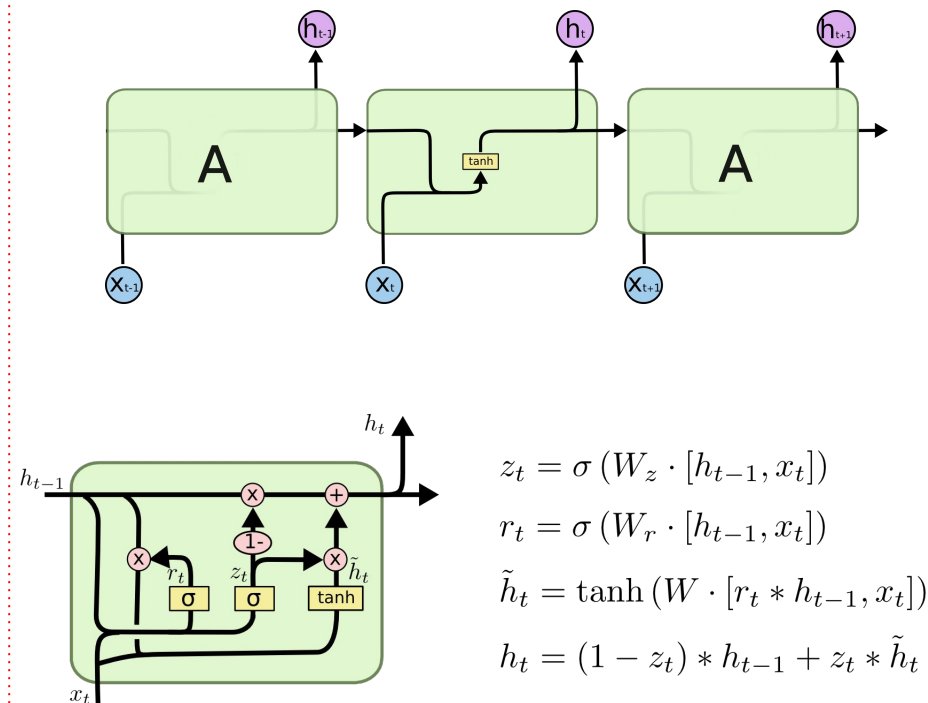
# Recurrent Neural Networks

- Standard NNs
  - Take in all the inputs at once
  - Can't capture sequential dependencies between input data
- Recurrent Neural Networks
- Great for data that is in a sequence form: Text, Videos etc.
- Example tasks: Language modeling (predict the next word in a sentence), language generation, sentiment analysis, video scene labeling etc.




# Gated Recurrent Units

- Standard RNNs
  - Hard to capture long-term dependencies
  - Perform worse on longer sequences
- Modifications to handle long-term dependencies better:
  - Long Short Term Memory (LSTMs)
  - Gated Recurrent Units (GRUs)
- Better than vanilla RNNs for most tasks




# Sequence-to-sequence models

- Takes a sequence as input, predicts a sequence as output. E.g. machine translation
- Also known as the encoder-decoder model
- Ideal when input and output sequences can be of different lengths
- Base case: Input sequence  $\rightarrow s \rightarrow$  output sequence
- Example tasks: Machine translation, speech recognition, sentence correction etc.




# Attention mechanism

- Base sequence-to-sequence model: Hard to capture longer context
- **Attention mechanism:** When predicting a particular output, tells you which part of the input to focus on
- Works really well when the output sequence has a strong 1:1 mapping with the input sequence
- Better than sequence models without attention for most tasks




# Other concepts

- Character-level RNNs
- Bidirectional RNNs
  - Captures dependencies in both directions
- Beam search decoding (vs. greedy decoding)

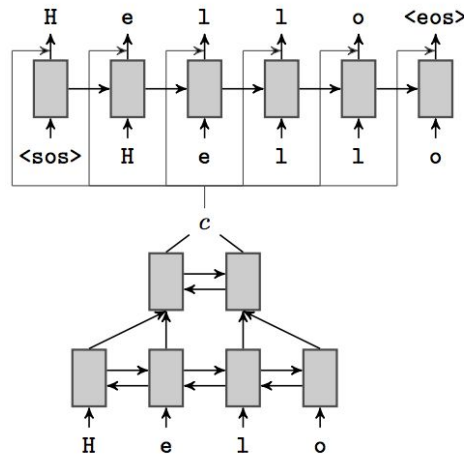


Example of a beam search tree




# Automatic question correction: Model

- Final question correction model:
  - **Multi-level, sequence-to-sequence, character-level GRU with attention**
- Tried solving the subproblems individually, but didn't work as well



# Automatic question correction: System Details

- Training
  - Training data: Pairs of [bad question, corrected question]
  - Training data size:  $O(100,000)$  examples
  - Tensorflow, on a single box with GPUs
  - Training time: 2-3 hours
- Serving:
  - Tensorflow, GPU-based serving
  - Latency: <500 ms p99
- Run on new questions added to Quora





# Automatic question correction: Results

Hey I am studying bsc computer science, what is the duration of MCA course for me?

Hey I am studying bsc for a BSc in computer science, what . What is the duration of the MCA course for me?

Do you thinks students should take CPR before leaving High School?

Do you thinks think students should take CPR before leaving High Schoolhigh school?

How much Uber cabs charge for the distance from mulapur to hotel park plaza Chandigarh?

How much do Uber cabs charge for the distance from mulapur Mulapur to hotel park plaza Hotel Park Plaza, Chandigarh?

Will Windows 10 improve exponentially by July 2018? (speed, stability, etc.)

Will Windows 10 improve exponentially by July 2018? 2018 (speed, stability, etc.)?

How many questions asked per day and per week on Quora?

How many questions are asked per day and per week on Quora?

Which coin/token is next big thing in crypto currencies? And why?

Which coin/token is the next big thing in crypto currenciescryptocurrencies? And whyWhy?

Which is a better course to do ACCA or CFA?

Which is a better course to do do, ACCA or CFA?

# Other question and answer understanding examples

# Question-Topic Labeler

- **Goal:** Given a question, come up with topics that describe it
- Traditional topic labeling: *Lots of text, few topics*
- Question-topic labeling: *Less text, huge topic space*
- Features:
  - Question text
  - Relation to other questions
  - Who asked the question
  - etc.

The Night's Watch (Game of Thrones)

Jon Snow (Game of Thrones character)

Game of Thrones (creative franchise)

A Song of Ice and Fire (book series)

Game of Thrones (TV series)



**What is the "official" reason Jon Snow was able to leave the Night's Watch?**

# Duplicate Question detection

- **Goal:** Single *canonical* question per intent
- Duplicate questions:
  - Make it harder for readers to seek knowledge
  - Make it harder for writers to find questions to answer
- *Semantic* question matching. Not simply a syntactic search problem.

## Why do like charges repel and opposite charges attract?

This question previously had details. They are now in a comment.

### Merged Questions

- ↑ Why do positive and negative charges attract?  
Undo Merge
- ↑ Why is positive charged particles are attracted towards negatively charged particles?  
Undo Merge
- ↑ Why do particles with like charges repel?  
Undo Merge
- ↑ Why do positive and negative charges attracts?  
Undo Merge
- ↑ Why do like charges repel?  
Undo Merge
- ↑ Why do like charges "repel" each other and other way around?  
Undo Merge
- ↑ Why do similar charges repel each other while opposite charges attract?  
Undo Merge
- ↑ Why do same charges repel and opposite charges attract?  
Undo Merge
- ↑ Why are charges attracted or repelled?  
Undo Merge

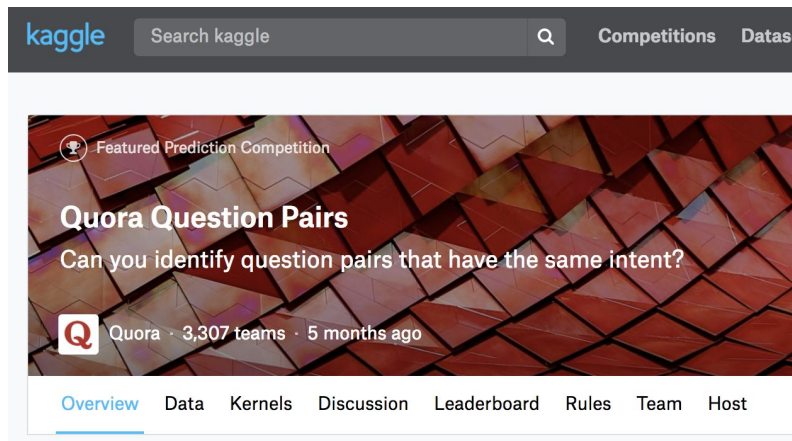
## First Quora Dataset Release: Question Pairs

Kornél Csernai

id	qid1	qid2	question1	question2	is_duplicate
447	895	896	What are natural numbers?	What is a least natural number?	0
1518	3037	3038	Which pizzas are the most popularly ordered pizzas on Domino's menu?	How many calories does a Dominos pizza have?	0
3272	6542	6543	How do you start a bakery?	How can one start a bakery business?	1
3362	6722	6723	Should I learn python or Java first?	If I had to choose between learning Java and Python, what should I choose to learn first?	1

Authors: *Shankar Iyer, Nikhil Dandekar, and Kornél Csernai*

Today, we are excited to announce the first in what we plan to be a series of public dataset releases. Our dataset releases will be oriented around various problems of relevance to Quora and will give researchers in diverse areas such as m...



## Is That a Duplicate Quora Question?

Published on February 27, 2017



Abhishek Thakur [+ Follow](#)

Chief Data Scientist @ Boost ai

[14 articles](#)

 574  39  68

## Deep text-pair classification with Quora's 2017 question dataset

BY MATTHEW HONNIBAL ON FEBRUARY 13, 2017

### Bilateral Multi-Perspective Matching for Natural Language Sentences

Zhiguo Wang, Wael Hamza, Radu Florian  
IBM T.J. Watson Research Center  
{zhigwang, whamza, raduf}@us.ibm.com

### Neural Paraphrase Identification of Questions with Noisy Pretraining


Gaurav Singh Tomar Thyago Duque Oscar Täckström  
Jakob Uszkoreit Dipanjan Das  
Google Inc.

{gtomar, duque, oscar, uszkoreit, dipanjand}@google.com



# BNBR classification

- BNBR = Be Nice, Be Respectful policy
- Binary classifier: Checks for BNBR violations on questions, answers, comments.
- Training data:
  - Positive: Confirmed BNBR violations
  - Negative: False BNBR reports, other good content
- Model: NN with 1 hidden layer (fastText)




# Answer ranking

- **Goal:** Given a question and  $n$  answers, come up with the ideal ranking
- What makes a good answer?
  - Truthful
  - Reusable
  - Well formatted
  - Clear and easy to read
  - ...

## Is AI an existential threat to humanity?

 Answer  Request  Follow  614  Comments  5+  Downvote


100+ Answers

 Tom Dong (董玉冰), Sumi Kim, Shankar Iyer, and 13 more upvoted this

 **Yoshua Bengio**, Head of Montreal Institute for Learning Algorithms, Professor @ U. Montreal

Answered Jan 19, 2016 · Upvoted by Zeeshan Zia, PhD in Computer Vision and Machine Learning and Alberto Bietti, PhD student in machine learning

Originally Answered: Yoshua Bengio: Is AI an existential threat to humanity?

There is a lot of uncertainty about the kind of AI we may one day reach that would achieve human-level intelligence or possibly more. Even though I am an optimist, because of that uncertainty, I cannot deny that things we do not want to happen could happen. That is why I have signed the Future of Life letter for a robust and beneficial AI (see [FLI - Future of Life Institute open letter](#)). 

 Sumi Kim, Kah Seng Tay, Jay Wacker, and 5 more upvoted this

 **Pedro Domingos**, Professor at the University of Washington and author of "The Master Algorithm" 

Answered Feb 2, 2016 · Featured on HuffPost · Upvoted by Alberto Bietti, PhD student in machine learning

Originally Answered: Pedro Domingos: Is AI an existential threat to humanity?

No, provided we stick to a simple rule: don't create AIs with goals of their own. AIs can come up with their own subgoals, but only in service of the goals we set them, and within the constraints we specify. This is how all AIs work today, and as long as they keep doing so, they can be infinitely intelligent without being a threat to us. You don't stay awake at night worrying that your dog will attack you. Why would you worry about your robot, which was evolved to serve you even more finely than your dog?

# Answer ranking

- **Features**

- Answer features: Quality, Formatting etc.
- Interaction features (upvotes/downvotes, clicks, comments...)
  - Network features: Who interacted with the answer?
- User features: Credibility, Expertise
- etc.


 Tom Dong (董玉冰), Sumi Kim, Shankar Iyer, and 13 more upvoted this



**Yoshua Bengio, Head of Montreal Institute for Learning Algorithms, Professor @ U. Montreal**

Answered Jan 19, 2016 · Upvoted by Zeeshan Zia, PhD in Computer Vision and Machine Learning and Alberto Bietti, PhD student in machine learning

Originally Answered: Yoshua Bengio: Is AI an existential threat to humanity?

There is a lot of uncertainty about the kind of AI we may one day reach that would achieve human-level intelligence or possibly more. Even though I am an optimist, because of that uncertainty, I cannot deny that things we do not want to happen could happen. That is why I have signed the Future of Life letter for a robust and beneficial AI (see [FLI - Future of Life Institute open letter](#)). 

-----  
that science and technology give us, the law of the jungle is not appropriate: we need an increase in individual and collective wisdom, for our sake and that of our descendants.

35.2k Views · 333 Upvotes

Upvoted 333 Downvote





# In conclusion

- Machine Learning systems form an important part of what drives Quora
- Lots of interesting Machine Learning problems and solutions all along the question lifecycle
- Machine Learning helps us make Quora more personalized and relevant to you *at scale*

# Thank you!

**Nikhil Dandekar**

Quora: Nikhil-Dandekar

Twitter: @nikhilbd

**Quora**