

Transfer Learning for Detecting Question Similarity between Stack Overflow Posts

Victor Kwan

c++ - What is the difference between \n and std::endl [duplicate]

Secure | https://stackoverflow.com/questions/19245748/what-is-the-difference-between-n-and-stdendl?noredirect=1&lq=1

stackoverflow

Questions Developer Jobs Tags Users

Search...

Log In Sign Up

What is the difference between \n and std::endl [duplicate]

Ask Question

Get personalized job matches now

Get started

asked 4 years, 6 months ago

viewed 4,346 times

active 4 years, 6 months ago

1

This question already has an answer here:

"\n" or '\n' or std::endl to std::cout? [duplicate] 6 answers

My answer, today, was edited and std::endl was changed to \n.

Does \n have any advantages?

c++ newline std

share improve this question

edited May 23 '17 at 12:12

Community 1 + 1

asked Oct 8 '13 at 10:52

khajvah 1,828 7 26 47

marked as duplicate by Some programmer dude, R. Martinho Fernandes, us2012, Cassio Neri, Jesse Good Oct 8 '13 at 10:57

This question has been asked before and already has an answer. If those answers do not fully address your question, please ask a new question.

Weird, I searched but couldn't find – khajvah Oct 8 '13 at 10:57

Anyway I would mark that edit as invalid. If someone thinks that you should modify the code in your answer, they should have left a comment for you to consider. Modifying your code risks making your answer mistaken, if they are wrong. – SJuan76 Oct 8 '13 at 11:10

@khajvah Edit it back. If you wanted std::endl, then no one should change it by editing to something you didn't want. (In general, when answering questions here, std::endl is probably better, since it's what is more usually wanted.) – James Kanze Oct 8 '13 at 11:15

@SJuan76 I agree. This is absolutely abuse of editing. I've had to edit my own back once or twice for similar reasons. – James Kanze Oct 8 '13 at 11:16

@khajvah In the context of your answer, std::endl is better. In general, unless there is a specific reason not to, std::endl is to be preferred, especially for output which will probably end up on a terminal. About the only time you want to use "\n" is when outputting large quantities of data in a block. – James Kanze Oct 8 '13 at 11:16

BLOG

Stack Overflow Isn't Very Welcoming. It's Time for That to Change.

Want to work remotely?

Find remote jobs based on technology

Get started

Looking for a job?

Senior Software Engineer (Java & Server Side)

MB00 Limited 9 Kowloon, Hong Kong

java sql-server

Core Java Engineer - Elasticsearch Distributed Systems Internals - Anywhere

Elastic 9 No office location

REMOTE

java concurrency

Join Stack Overflow to learn, share knowledge, and build your career.

Email Sign Up

OR SIGN IN WITH

Google

Facebook

Many duplicate questions arise from lack of problem knowledge, non-matching titles, and lack of knowledge of relevant terminology and buzzwords.


(Ahasanuzzaman et al. 2016)

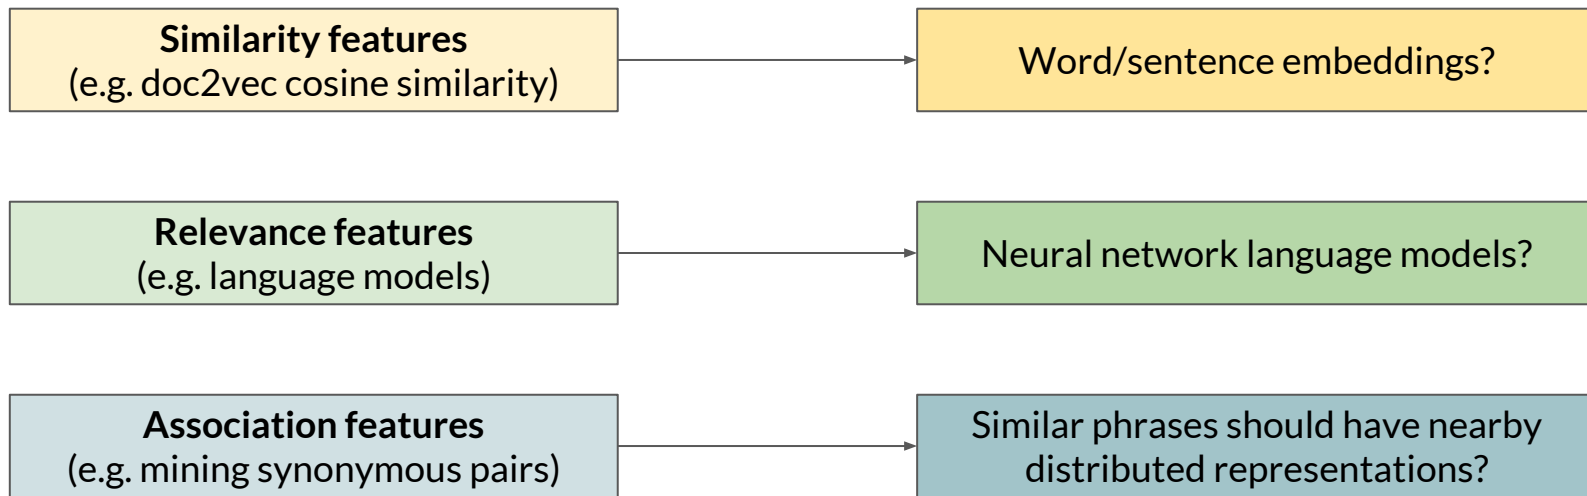
Check connection is
active in ASP.NET?

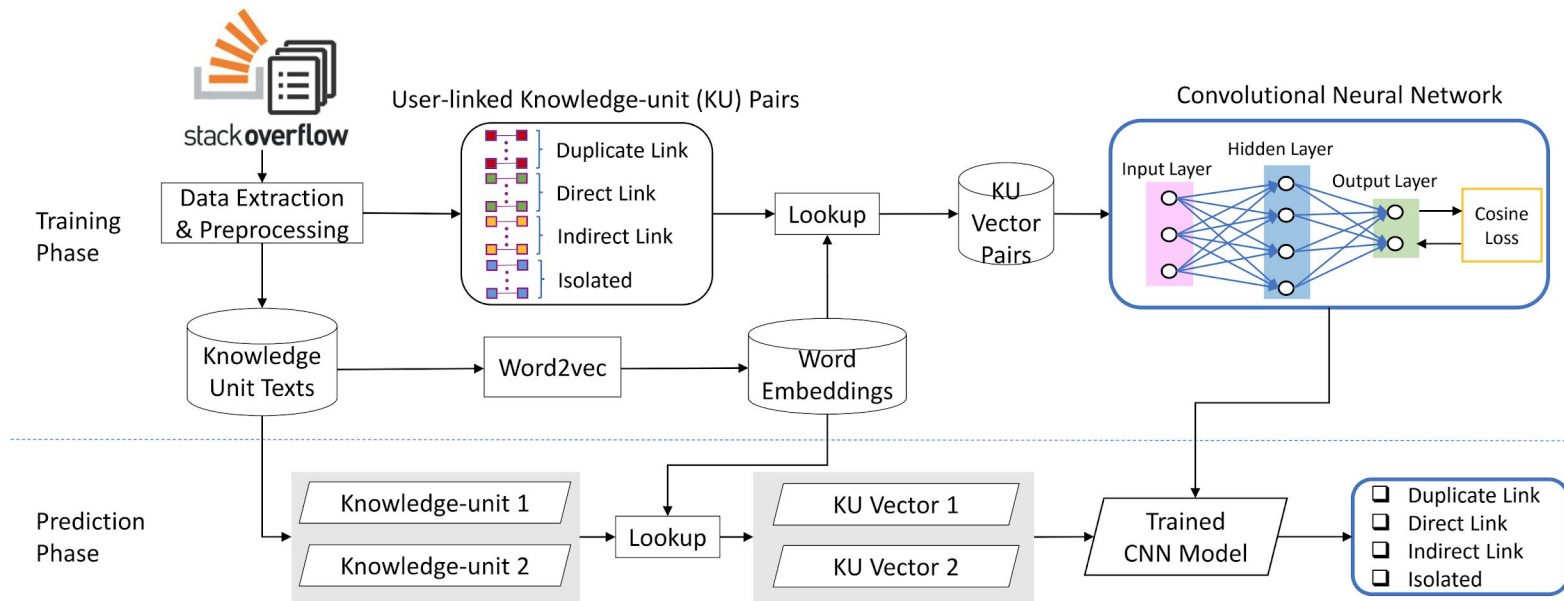
(QID: 761376)

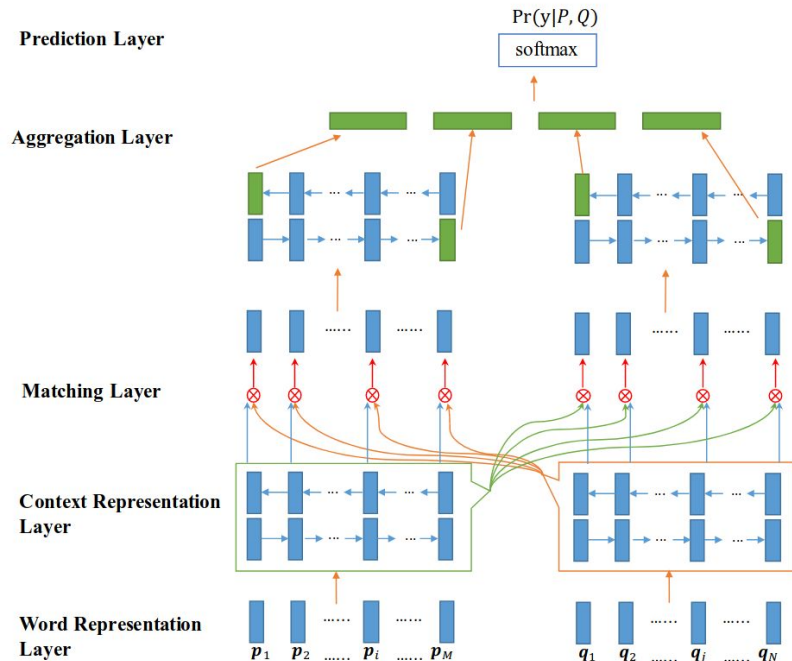
How to check if user
is still present?

(QID: 3374120)











(Wang et al. 2017)



(Parikh et al. 2016)



3. The sentence representations, their absolute difference, their element-wise product, and their tags are concatenated for classification.

2. Each timestep of the GRU output is compared against the context vector.

1. Concatenated word and character-based embeddings are passed as input into the GRUs.

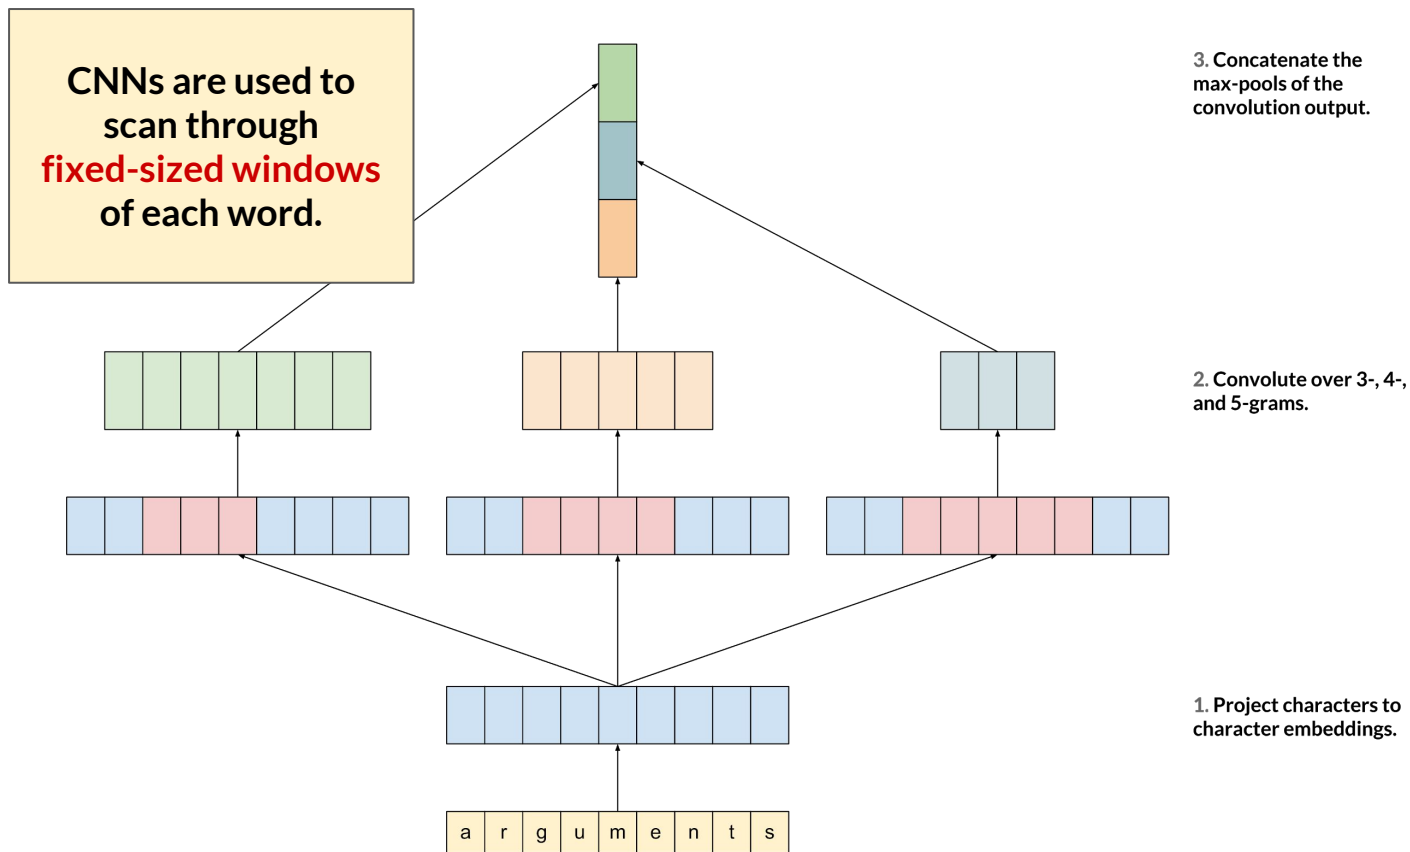
Classifier: (Mou et al. 2016)
Attention: (Yang et al. 2016)
GRU: (Cho et al. 2014)

UIAlertView

Many words on Stack Overflow
that aren't in our dictionary

UI / *Alert* / *View*

Look at **how** the words are spelt!
("subword information")



An illustration of how the model works
*Note that we only deal with question titles;
the paragraph is produced as an example.*

Did you hear? The genome of the banana has been sequenced, an important development in scientist's efforts to produce better bananas. A look at that genome has revealed curious things, said Pat Heslop-Harrison, a plant geneticist at the University of Leicester in England who was a coauthor of the report published this week in the journal Nature. For example, there are regions of the banana genome that don't seem to be involved in making proteins but are shared by many different species of plants, far beyond bananas.



Article: (Mestel, 2012)

Image: <http://pngplay.com/image/1810>

Did you hear? The genome of the banana has been sequenced, an important development in scientist's efforts to produce better bananas. A look at that genome has revealed curious things, said Pat Heslop-Harrison, a plant geneticist at the University of Leicester in England who was a coauthor of the report published this week in the journal Nature. For example, there are regions of the banana genome that don't seem to be involved in making proteins but are shared by many different species of plants, far beyond bananas.



To develop better bananas, scientists have sequenced the banana genome. Pat Heslop-Harrison says in Nature that the genome reveals curious things. Parts of the genome are shared by many other plants.



To develop better bananas, scientists have sequenced the banana genome. Pat Heslop-Harrison says in Nature that the genome reveals curious things. Parts of the genome are shared by many other plants.



Scientists have sequenced the banana genome and have found that parts are shared by other plants.




Scientists have sequenced the **banana** genome and have found that parts are shared by other plants.



Scientists have sequenced the **strawberry** genome and have found that parts are shared by other plants.



Non-duplicate!



3. The sentence representations, their absolute difference, their element-wise product, and their tags are concatenated for classification.



2. Each timestep of the GRU output is compared against the context vector.

1. Concatenated word and character-based embeddings are passed as input into the GRUs.

Classifier: (Mou et al. 2016)
Attention: (Yang et al. 2016)
GRU: (Cho et al. 2014)

Marked as **duplicate**
... but where do we get
non-duplicates?

c++ - What is the difference between `\n` and `std::endl` [duplicate]

Get personalized job matches now

This question already has an answer here:
"`\n`" or "`\n`" or `std::endl` to `std::cout`? [duplicate] 6 answers

My answer, today, was edited and `std::endl` was changed to `\n`.

Does `\n` have any advantages?

asked Oct 8 '13 at 10:52
1,828 7 26 47

marked as duplicate by Some programmer dude, R. Martinho Fernandes, us2012, Cassio Neri, Jesse Good Oct 8 '13 at 10:57

This question has been asked before and already has an answer. If those answers do not fully address your question, please ask a new question.

Weird, I searched but couldn't find - khajvah Oct 8 '13 at 10:57

Anyway I would mark that edit as invalid. If someone thinks that you should modify the code in your answer, they should have left a comment for you to consider. Modifying your code risks making your answer mistaken, if they are wrong. - SJuan76 Oct 8 '13 at 11:10

1 @khajvah Edit it back. If you wanted `std::endl`, then no one should change it by editing to something you didn't want. (In general, when answering questions here, `std::endl` is probably better, since it's what is more usually wanted.) - James Kanze Oct 8 '13 at 11:15

1 @SJuan76 I agree. This is absolutely abuse of editing. I've had to edit my own back once or twice for similar reasons. - James Kanze Oct 8 '13 at 11:16

1 @khajvah In the context of your answer, `std::endl` is better. In general, unless there is a specific reason not to, `std::endl` is to be preferred, especially for output which will probably end up on a terminal. About the only time you want to use `\n` is when outputting large quantities of data in a block. - James Kanze Oct 8 '13 at 11:16

asked 4 years, 6 months ago
viewed 4,346 times
active 4 years, 6 months ago

BLOG
Stack Overflow Isn't Very Welcoming. It's Time for That to Change.

Want to work remotely?
Find remote jobs based on technology

Looking for a job?

Senior Software Engineer (Java & Server Side)
MB00 Limited 9 Kowloon, Hong Kong
java sql-server

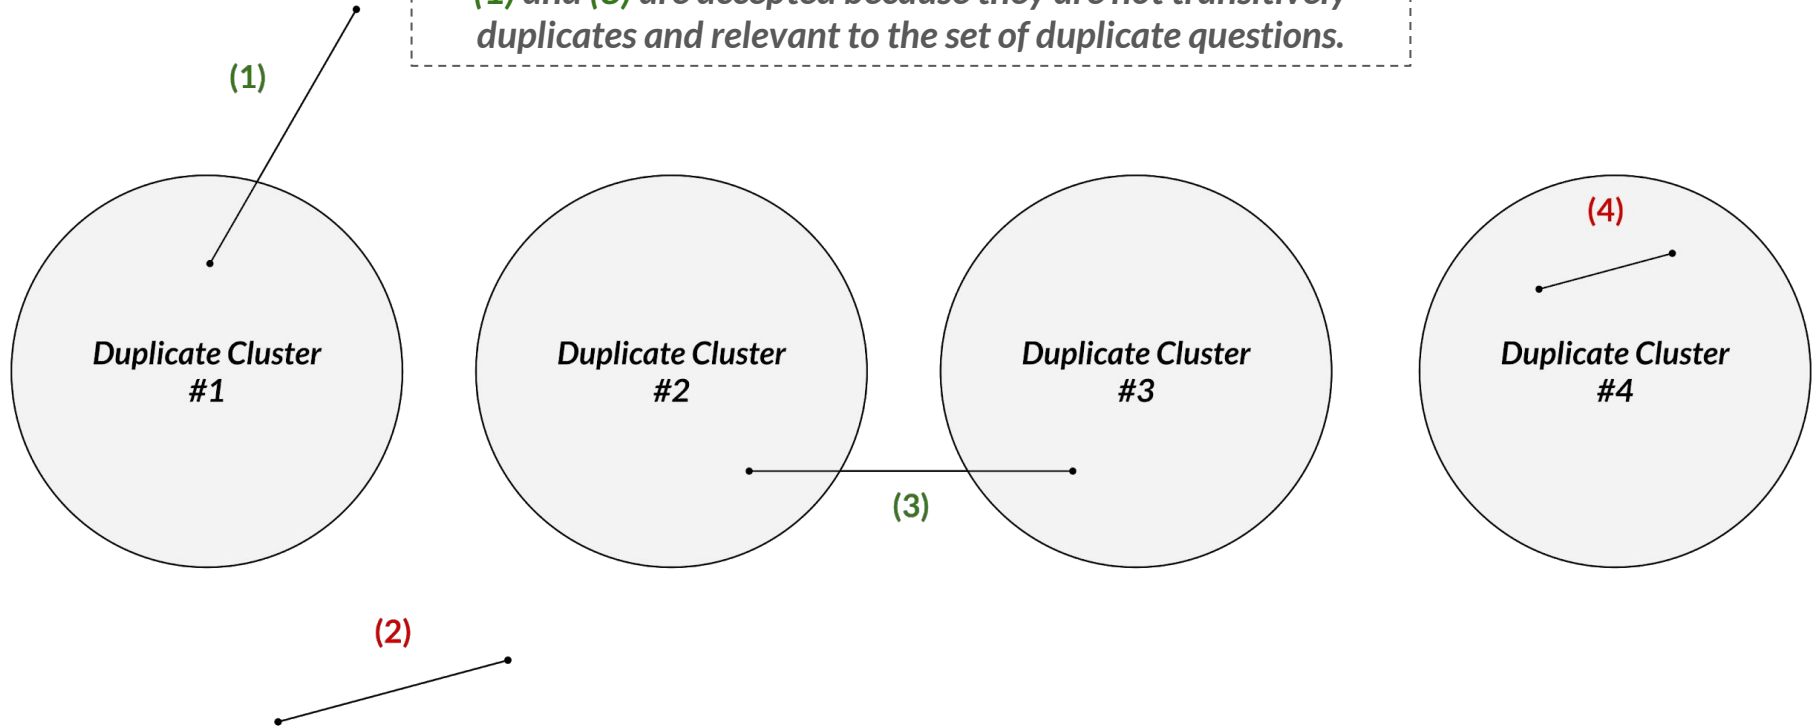
Core Java Engineer - Elasticsearch Distributed Systems Internals - Anywhere
Elastic 9 No office location
REMOTE
java concurrency

Join Stack Overflow to learn, share knowledge, and build your career.

Email Sign Up OR SIGN IN WITH Google Facebook

Question filtering process

(1) and (3) are accepted because they are not transitively duplicates and relevant to the set of duplicate questions.



Quora

Develop some notions about questions on the Quora dataset

149,306 positive examples
255,045 negative examples



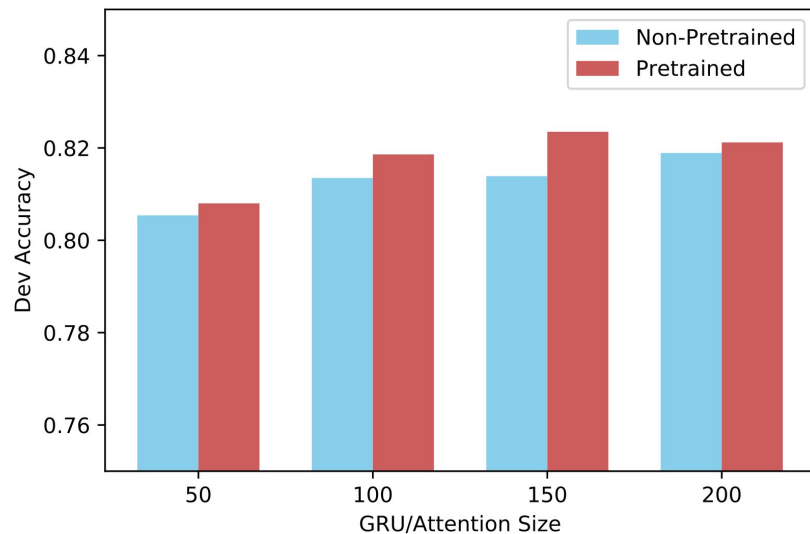
Further *tune* these notions by training against the Stack Overflow dataset?
(Now less likely to fit to noise.)

137,793 positive examples
139,282 negative examples


Configuration	Dev Accuracy	Test Accuracy	Test Macro F-Score
50, non-pretrained	0.805	0.796	0.797
50, pretrained	0.808	0.795	0.795
100, non-pretrained	0.814	0.803	0.803
100, pretrained	0.819	0.807	0.807
150, non-pretrained	0.814	0.801	0.803
150, pretrained	0.824	0.814	0.814
200, non-pretrained	0.819	0.809	0.809
200, pretrained	0.821	0.809	0.809

Hyperparameter tuning results

Pretrained models are pretrained against the Quora dataset then tuned against the SO dataset for 50 epochs. The classifier is reinitialized when tuning.



Model performance on dev set.
We select the model at the epoch with highest dev accuracy for evaluation.



Model performance on test set.
This performance measure tells us the generalizability of the model.

Pretraining as a form of regularization

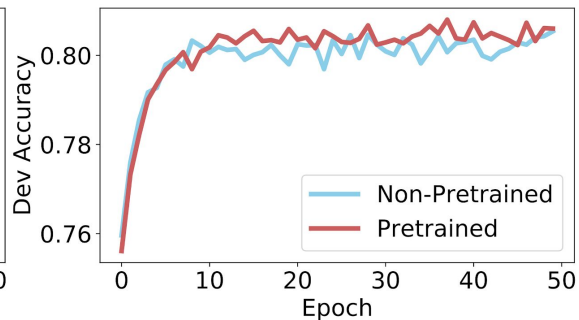
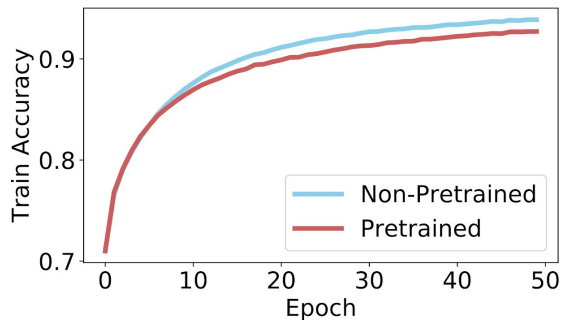


Scientists have sequenced
the banana genome and
have found that parts are
shared by other plants.


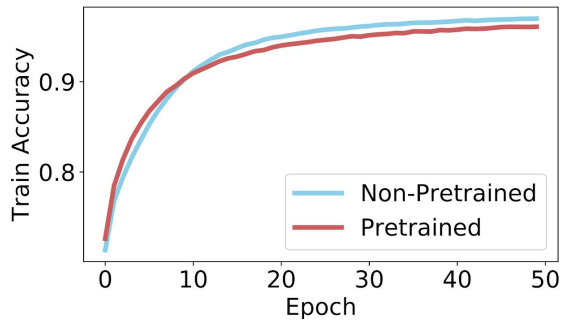


Scientists have sequenced
the banana genome and
have found that parts are
shared by other plants.



size=50



size=150



size=200




How effective are our components?


Ablated Component	Dev Accuracy	Test Accuracy	Test Macro F-Score
None	0.824	0.814	0.814
Character-Based Embeddings	0.818	0.811	0.812
Attention Mechanism	0.814	0.803	0.804
Fully-Connected Classifier	0.659	0.652	0.652
Tags	0.812	0.805	0.806

Ablation study results

The "ablated" component is removed from the model. Each model is pretrained for 25 epochs and tuned for 25 epochs.



Full Model
... but only 0.002 improvement to F1 score?



Ablated Character-Based Embeddings
The [unk] are int32 and int64 respectively.

ruby on rails more
elegant way to
authenticate that
users can edit only
their own content

(QID: 3183210)

how to only user to
edit / update / delete
their own questions ?

(QID: 21632236)

Our full model outperforms our tag
ablated model by **0.008 F1 score**.

ruby on rails more
elegant way to
authenticate that
users can edit only
their own content

`<ruby-on-rails>`


(QID: 3183210)

how to only user to
edit / update / delete
their own questions ?

`<ruby-on-rails>`

(QID: 21632236)

Our full model outperforms our tag
ablated model by **0.008 F1 score**.



Ablated Fully-Connected Classifier


*This model variant uses the Manhattan distance to compare sentence representations. The tag-ablated model outperforms this model by **0.154 F1 score**.*

**What's transferred between the
Quora and Stack Overflow datasets?**


Unfrozen Components	Dev Accuracy	Test Accuracy	Test Macro F-Score
All	0.824	0.814	0.814
>= GRU	0.822	0.807	0.807
>= Attention Mechanism	0.785	0.769	0.770
>= Classifier	0.778	0.760	0.761
None		0.522	0.560

>= indicates that all components including and after the specified component are unfrozen.

Unfrozen components study results
Components that are frozen are not tuned. Tuning is carried out over 25 epochs, with the classifier re-initialized.

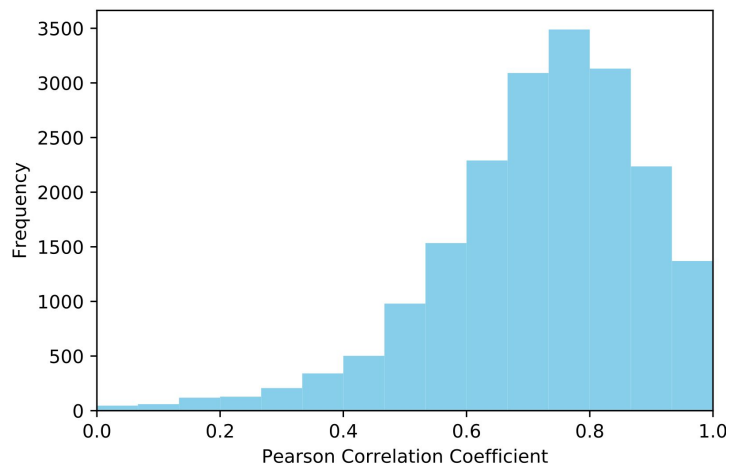


Full model

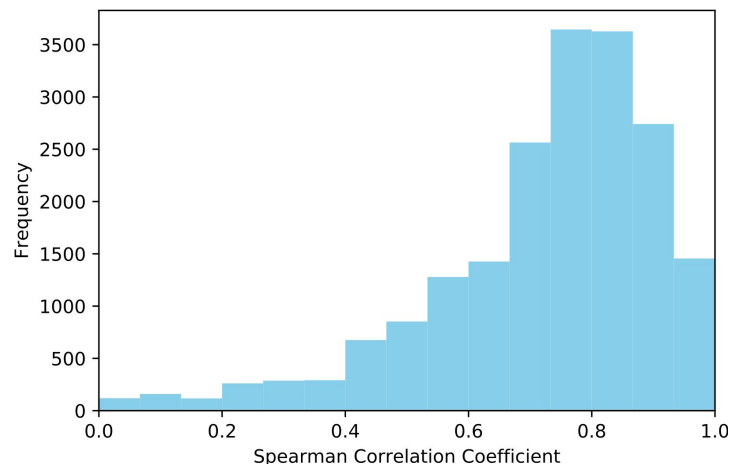


Quora-trained model

Correlation between attention weights produced by Quora-trained and full models



Pearson Correlation
Takes direction and magnitude into account.



Spearman Rank-Order Correlation
Only takes direction into account.

Attention is consistent	✓
Sentence representations contain relevant information	✓
<i>... what else could the difference be?</i>	

		Predicted	
		Non-duplicate	Duplicate
Label	Non-duplicate	4826	174
	Duplicate	4601	399

Quora-trained model confusion matrix

*The model has exceptionally poor recall on positive examples
(0.080) but comparatively better precision (0.696).*

**difference between
static class and
singleton pattern ?**

(QID: 519520)

**what is the main
difference between
static class &
singleton class**

(QID: 23566750)


now donald trump is
the president of us ,
what is the **impact** on
india ?

(QID: 732446)


what will be the
impact of trump
presidency **on india** ?

(QID: 732447)

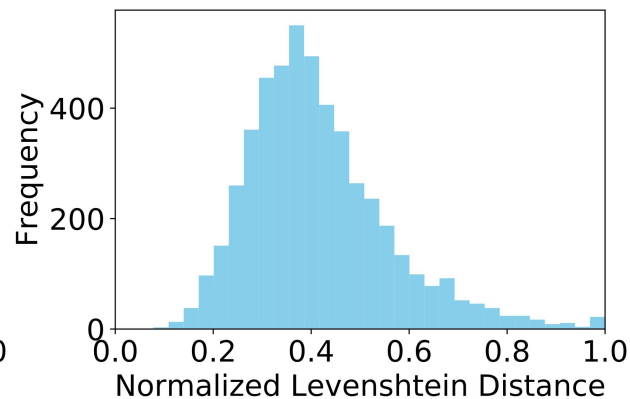
Character-based Normalized Levenshtein Distance between question pairs
This distance is computed as the LCS overlap, normalized by the longer of the two questions.



**Positive SO pairs
(correctly identified by
Quora-trained model)**




Positive Quora pairs




Positive SO pairs

Token-based Normalized Levenshtein Distance between question pairs


This distance is computed as the LCS overlap, normalized by the longer of the two questions.



**Positive SO pairs
(correctly identified by
Quora-trained model)**



Positive Quora pairs



Positive SO pairs

Unfreezing the classifier
(+0.201 F1)

Classifier is less reliant on monotonic alignment between question pairs; better handling of technical vocabulary.

Unfreezing the GRU
(+0.037 F1)

GRU shifts towards better utilizing the character-based embeddings.

get size of file on disk

(QID: 3750590)

file size - actual
number of bytes used
on disk (not
allocated)

(QID: 13959228)

Questions should be clearly written

- Quora requires correct spelling, punctuation, grammar and capitalization to be used in questions. A question should consist of complete sentence(s) and end in a question mark.
- Questions should contain enough information for meaningful answers to be posted, and make clear how answers can meet their expectations.
- Questions should be concise. As a general guideline, questions should generally be no longer than 2 sentences, and only one question should be asked at a time.

Pretraining

Through pretraining, we see a over one point improvement in both **test accuracy** and **f-score**.

We see that pretraining can be an effective form of regularization.

Model Components

Each of our model components contribute positively to our model performance.

However, we note that character-based embeddings may be prone to overfitting.

Transferable Representations

We find that our representations are transferable.

Structural differences between Quora and SO questions may motivate using position-invariant models.