

Learning to Ask Good Questions: Ranking Clarification Questions using Neural Expected Value of Perfect Information

Sudha Rao¹ and Hal Daumé III^{1,2}



**¹University of Maryland
College Park**



**²Microsoft Research
New York**

CODE + DATA: https://github.com/raosudha89/ranking_clarification_questions

Natural Language Understanding



Natural Language Understanding



How long does it take to get a PhD?

Natural Language Understanding



How long does it take to get a PhD?

Give me a recipe for lasagna



Natural Language Understanding



How long does it take to get a PhD?

Please bring me my coffee mug from the kitchen



Give me a recipe for lasagna



Natural Language Understanding



How long does it take to get a PhD?

Please bring me my coffee mug from the kitchen



Give me a recipe for lasagna



Human Interactions



Human Interactions



Please bring me
my coffee mug
from the kitchen



Human Interactions



Please bring me
my coffee mug
from the kitchen



Human Interactions



Teach Machines to Ask Clarification Questions

Teach Machines to Ask Clarification Questions



How long does it take to get a PhD ?

In which field?

Teach Machines to Ask Clarification Questions



How long does it take to get a PhD ?

In which field?



Teach Machines to Ask Clarification Questions



How long does it take to get a PhD ?

In which field?

Please bring me
my coffee mug
from the kitchen

What color is
your coffee mug?



Give me a recipe
for lasagna

Any dietary
restrictions?



Teach Machines to Ask Clarification Questions



How long does it take to get a PhD ?

In which field?

Please bring me
my coffee mug
from the kitchen

What color is
your coffee mug?



Context-aware questions about
missing information

Give me a recipe
for lasagna

Any dietary
restrictions?



Reading Comprehension Question Generation

My class is going to the movies on a field trip next week.

We have to get permission slips signed before we go.

We are going to see a movie that tells the story from a book we read.

Goal: Assess someone's understanding of the text

Q: What do the students need to do before going to the movies?

- Heilman. "Automatic factual question generation from text" Ph.D. thesis 2011
- Vasile, et al. "The first question generation shared task evaluation challenge." NLG 2010
- Olney, Graesser, and Person. "Question generation from concept maps." Dialogue & Discourse 2012
- Chali and Hasan. "Towards Topic-to-Question Generation." ACL 2015
- Serban, et al. "Generating Factoid Questions With Recurrent Neural Networks" ACL 2016
- Du, Shao & Cardie "Learning to ask: Neural question generation for reading comprehension" ACL 2017
- Tang et al. "Learning to Collaborate for Question Answering and Asking." NAACL 2018
- Mrinmaya and Xing. "Self-Training for Jointly Learning to Ask and Answer Questions." NAACL 2018

Question Generation for Slot Filling

USER: I want to go to **Melbourne** on **July 14**
SYSTEM: What **time** do you want to leave?
USER: I must be in Melbourne by 11 am
SYSTEM: Would you like a **Delta** flight that arrives at 10.15 am?
USER: Sure
SYSTEM: In what **name** should I make the reservation?

SLOTS

<origin city>
<departure city>
<origin time>
<departure time>
<airline>

- Goddeau, et al. "A form-based dialogue manager for spoken language applications." 1996
- Bobrow., et al. "GUS, a frame-driven dialog system." Artificial intelligence 1977
- Lemon, et al. "An ISU dialogue system exhibiting reinforcement learning of dialogue policies: generic slot-filling in the TALK in-car system." EACL 2006
- Williams, et al. "The Dialog State Tracking Challenge" SIGDIAL 2013
- Young, et al. "Pomdp-based statistical spoken dialog systems: A review." IEEE 2013
- Dhingra, et al. "Towards End-to-End Reinforcement Learning of Dialogue Agents for Information Access." ACL 2017
- Bordes, et al. "Learning end-to-end goal-oriented dialog." ICLR 2017

Other types of Question Generation

- Liu, et al. "Automatic question generation for literature review writing support." International Conference on Intelligent Tutoring Systems. 2010
- Penas and Hovy, "Filling knowledge gaps in text for machine reading" International Conference on Computational Linguistics: Posters ACL 2010
- Artzi & Zettlemoyer, "Bootstrapping semantic parsers from conversations" EMNLP 2011
- Labutov, et al. "Deep questions without deep understanding" ACL 2015
- Mostafazadeh et al. "Generating natural questions about an image." ACL 2016
- Mostafazadeh et al. "Multimodal Context for Natural Question and Response Generation." IJCNLP 2017.
- Rothe, Lake and Gureckis. "Question asking as program generation" NIPS 2017.

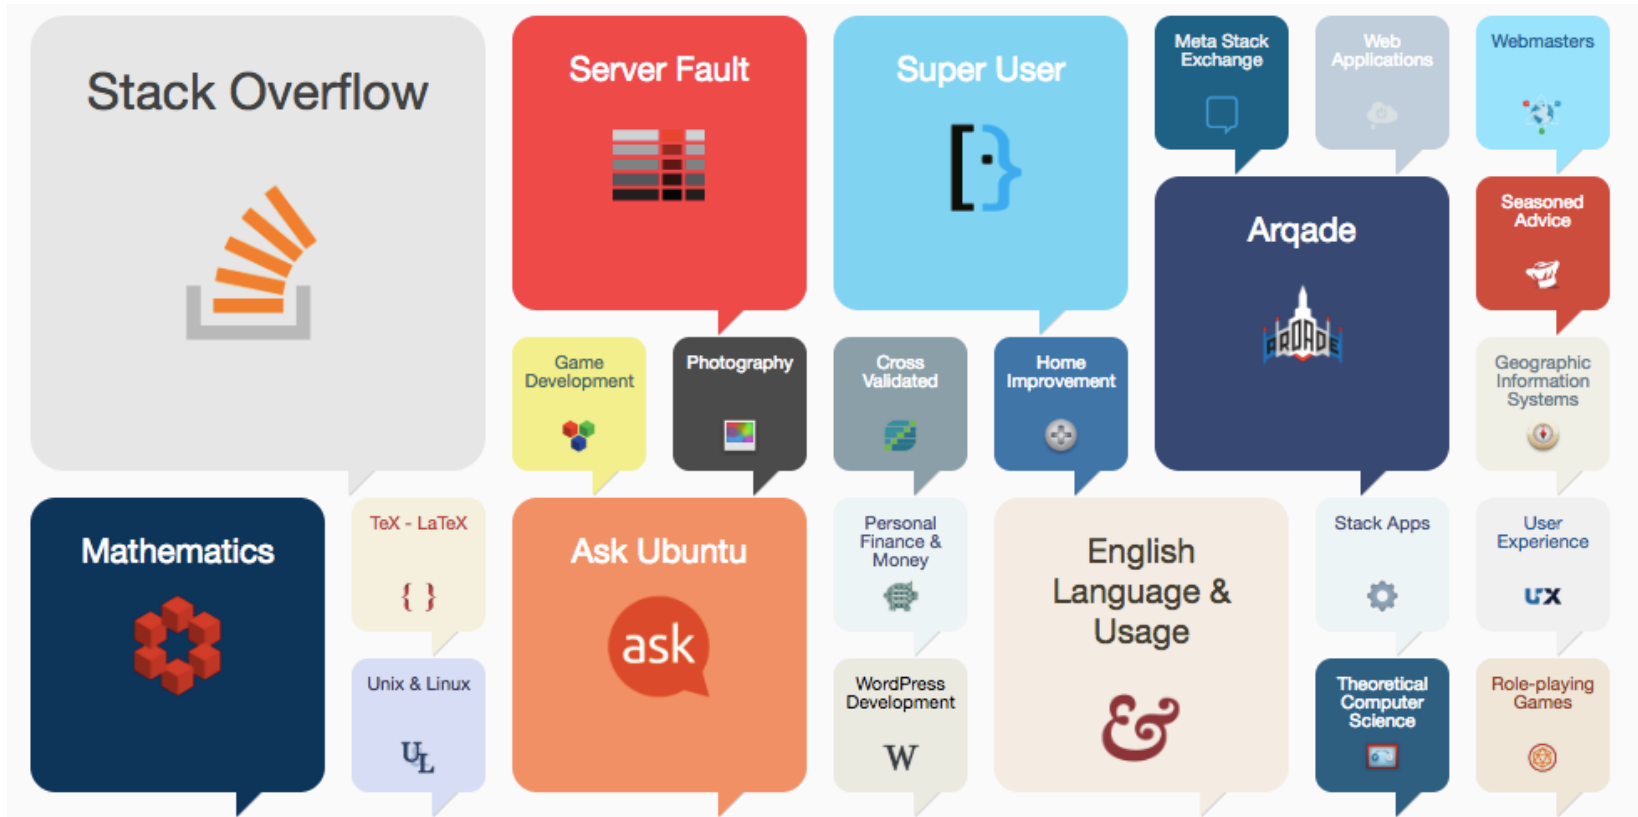
Talk Outline

- Clarification Questions Dataset
- Problem Formulation: Question Ranking
- Expected Value of Perfect Information (EVPI) inspired model
- Evaluation
- Conclusion

Talk Outline

- Clarification Questions Dataset
- Problem Formulation: Question Ranking
- Expected Value of Perfect Information (EVPI) inspired model
- Evaluation
- Conclusion

StackExchange Question-Answer Forum



Clarification Questions Dataset



How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.

I'm having this error message while running `./configure`

<error message>

So I have the library but the program installation isn't finding it.

Any help? Thanks in advance!



Clarification Questions Dataset



How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

Finding: Questions go unanswered for a long time if they are not clear enough

Asaduzzaman, Muhammad, et al. "Answering questions about unanswered questions of stack overflow." Working Conference on Mining Software Repositories. IEEE Press, 2013.



Clarification Questions Dataset



How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

What version of ubuntu do you have?

**Question
comment**



Clarification Questions Dataset



How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

← **Initial Post**

**Question
comment**

What version of ubuntu do you have?

I'm aiming to install ape in Ubuntu 14.04 LTS, a simple code for
pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

← **Updated Post**

Clarification Questions Dataset



How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

← **Initial Post**

**Question
comment**

What version of ubuntu do you have?

I'm aiming to install ape **in Ubuntu 14.04 LTS**, a simple code for pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

**Edit as an answer
to the question**

← **Updated Post**

Clarification Questions Dataset



How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

Initial Post

Question comment

What version of ubuntu do you have?

I'm aiming to install ape **in Ubuntu 14.04 LTS**, a simple code for pseudopotential generation.
I'm having this error message while running ./configure
<error message>
So I have the library but the program installation isn't finding it.
Any help? Thanks in advance!

Edit as an answer to the question

Updated Post

Dataset Creation

(post , question , answer) triples

post

Original post

question

Clarification question posted in comments

answer

Edit made to the post in response to the question
OR author's reply to the question comment

Dataset Size: ~77 K triples

Domains: Askubuntu, Unix, Superuser

Note: We identify a question using the question mark (?) token. We match the edit to the answer using timestamp & word embedding similarity based heuristics.

Talk Outline

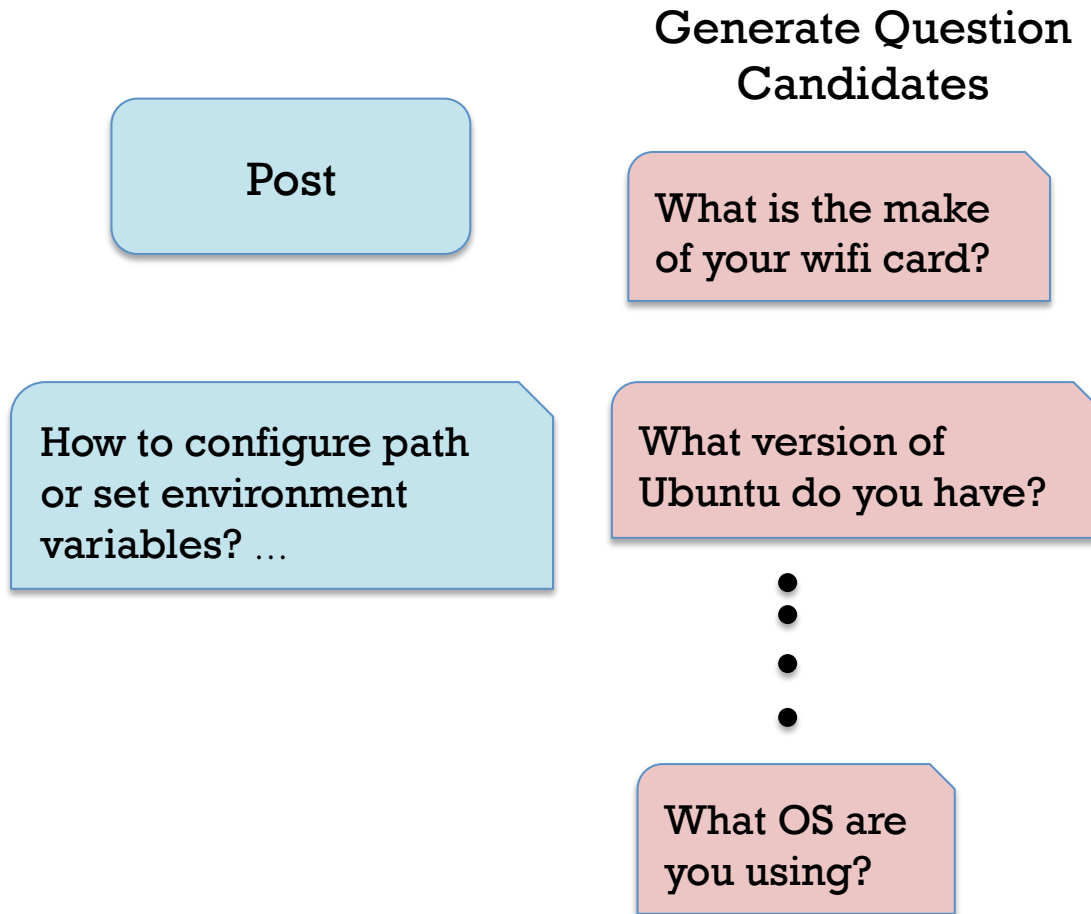
- Clarification Questions Dataset
- Problem Formulation: Question Ranking
- Expected Value of Perfect Information (EVPI) inspired model
- Evaluation
- Conclusion

Problem Formulation: Question Ranking

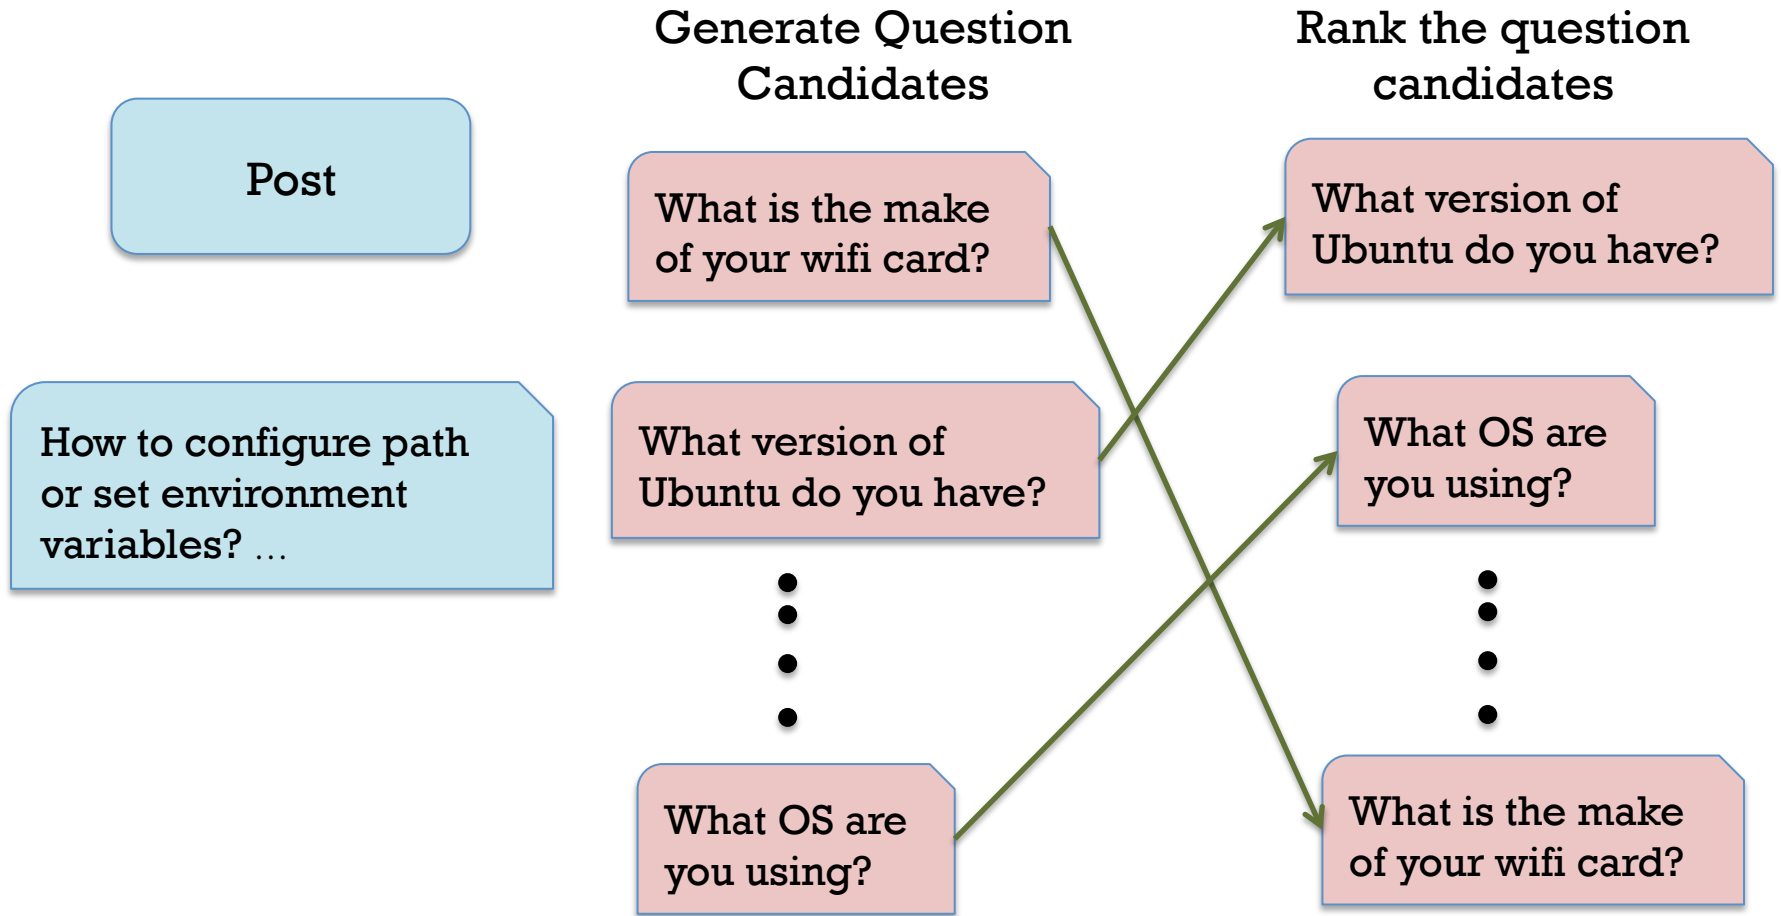
Post

How to configure path
or set environment
variables? ...

Problem Formulation: Question Ranking



Problem Formulation: Question Ranking



Problem Formulation: Question Ranking

How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.

I'm having this error message while running ./configure

<error message>

So I have the library but the program installation isn't finding it.

Any help? Thanks in advance!

What version of Ubuntu do you have?

How are you installing ape?

Do you have GSL installed?

Shortlist of
useful questions

Talk Outline

- Clarification Questions Dataset
- Problem Formulation: Question Ranking
- Expected Value of Perfect Information (EVPI) inspired model
- Evaluation
- Conclusion

Key Idea

How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.

I'm having this error message while running ./configure

<error message>

So I have the library but the program installation isn't finding it.

Any help? Thanks in advance!

Possible questions

(a) *What version of Ubuntu do you have?*

→ **Just right**

Key Idea

How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.

I'm having this error message while running ./configure

<error message>

So I have the library but the program installation isn't finding it.

Any help? Thanks in advance!

Possible questions

- (a) *What version of Ubuntu do you have?* → **Just right**
- (b) *What is the make of your wifi card?* → **Not useful**

Key Idea

How to configure path or set environment variables for installation?

I'm aiming to install ape, a simple code for pseudopotential generation.

I'm having this error message while running ./configure

<error message>

So I have the library but the program installation isn't finding it.

Any help? Thanks in advance!

Possible questions

- (a) *What version of Ubuntu do you have?* → **Just right**
- (b) *What is the make of your wifi card?* → **Not useful**
- (c) *Are you running Ubuntu 14.10 kernel 4.4.0-59-generic on an x86 64 architecture?* → **Unlikely to add value**

- Use EVPI to identify questions that add the most value to the given post

Avriel, Mordecai, and A. C. Williams. "The value of information and stochastic programming." *Operations Research* 18.5 (1970)



- Use EVPI to identify questions that add the most value to the given post
- Definition: Value of Perfect Information VPI (x)

How much value does x add to a given information content c ?

- Use EVPI to identify questions that add the most value to the given post
- Definition: Value of Perfect Information VPI (x)
How much value does x add to a given information content c ?
- Since we have not acquired x , we define its value in expectation

- Use EVPI to identify questions that add the most value to the given post
- Definition: Value of Perfect Information VPI (x)
How much value does x add to a given information content c?
- Since we have not acquired x, we define its value in expectation

$$\text{EVPI}(x|c) = \sum_{x \in X} P(x|c) \text{Utility}(x, c)$$

Likelihood of x given c
↓
↑
Value of updating c with x

Avriel, Mordecai, and A. C. Williams. "The value of information and stochastic programming." *Operations Research* 18.5 (1970)

EVPI formulation for our problem

EVPI formulation for our problem

$$\text{EVPI} \left(\boxed{q_i} \mid \boxed{p} \right)$$

\boxed{p} : given post

$\boxed{q_i}$: question from set of question candidates Q

EVPI formulation for our problem

Likelihood of a_j being the answer to q_i on post p



$$\mathbf{EVPI} (\boxed{q_i} \mid \boxed{p}) = \mathbf{P} (\boxed{a_j} \mid \boxed{p}, \boxed{q_i})$$

\boxed{p} : given post

$\boxed{q_i}$: question from set of question candidates Q

EVPI formulation for our problem

Likelihood of a_j being the answer to q_i on post p

$$\mathbf{EVPI} (\boxed{q_i} \mid \boxed{p}) = \mathbf{P} (\boxed{a_j} \mid \boxed{p}, \boxed{q_i}) \mathbf{U} (\boxed{p} + \boxed{a_j})$$

Utility of updating the post p with answer a_j

\boxed{p} : given post

$\boxed{q_i}$: question from set of question candidates Q

EVPI formulation for our problem

Likelihood of a_j being the answer to q_i on post p

$$\mathbf{EVPI} (\boxed{q_i} \mid \boxed{p}) = \sum_{\boxed{a_j} \in \mathbf{A}} \mathbf{P} (\boxed{a_j} \mid \boxed{p}, \boxed{q_i}) \mathbf{U} (\boxed{p} + \boxed{a_j})$$

Utility of updating the post p with answer a_j

- p : given post
- q_i : question from set of question candidates Q
- a_j : answer from set of answer candidates A



EVPI formulation for our problem

We rank questions based on their EVPI value

Likelihood of a_j being the answer to q_i on post p

$$\text{EVPI} (q_i \mid p) = \sum_{a_j \in A} P (a_j \mid p, q_i) U (p + a_j)$$

Utility of updating the post p with answer a_j

- p : given post
- q_i : question from set of question candidates Q
- a_j : answer from set of answer candidates A



EVPI formulation for our problem

Likelihood of a_j being the answer to q_i on post p

$$\mathbf{EVPI} (q_i \mid p) = \sum_{a_j \in \mathbf{A}} \mathbf{P} (a_j \mid p, q_i) \mathbf{U} (p + a_j)$$

Utility of updating the post p with answer a_j

p : given post

q_i : question from set of question candidates Q

a_j : answer from set of answer candidates A

1

EVPI formulation for our problem

② Likelihood of a_j being the answer to q_i on post p

$$\text{EVPI} (q_i \mid p) = \sum_{a_j \in A} P (a_j \mid p, q_i) U (p + a_j)$$

Utility of updating the post p with answer a_j

p : given post

q_i : question from set of question candidates Q

a_j : answer from set of answer candidates A

①

EVPI formulation for our problem

2 Answer Modeling
 ↓

$$\text{EVPI} (q_i \mid p) = \sum_{a_j \in \mathbf{A}} \mathbf{P} (a_j \mid p, q_i) \mathbf{U} (p + a_j)$$

↑
 Utility of updating the post p with answer a_j

p : given post

q_i : question from set of question candidates Q

a_j : answer from set of answer candidates A

1



EVPI formulation for our problem

2 Answer Modeling

↓

$$\text{EVPI} (q_i \mid p) = \sum_{a_j \in \mathbf{A}} \mathbf{P} (a_j \mid p, q_i) \mathbf{U} (p + a_j)$$

3 Utility of updating the post p with answer a_j

↑

p : given post

q_i : question from set of question candidates Q

a_j : answer from set of answer candidates A

1

EVPI formulation for our problem

$$\text{EVPI} (q_i \mid p) = \sum_{a_j \in \mathbf{A}} \mathbf{P} (a_j \mid p, q_i) \mathbf{U} (p + a_j)$$

2 Answer Modeling

3 Utility Calculator

p : given post

q_i : question from set of question candidates Q

a_j : answer from set of answer candidates A

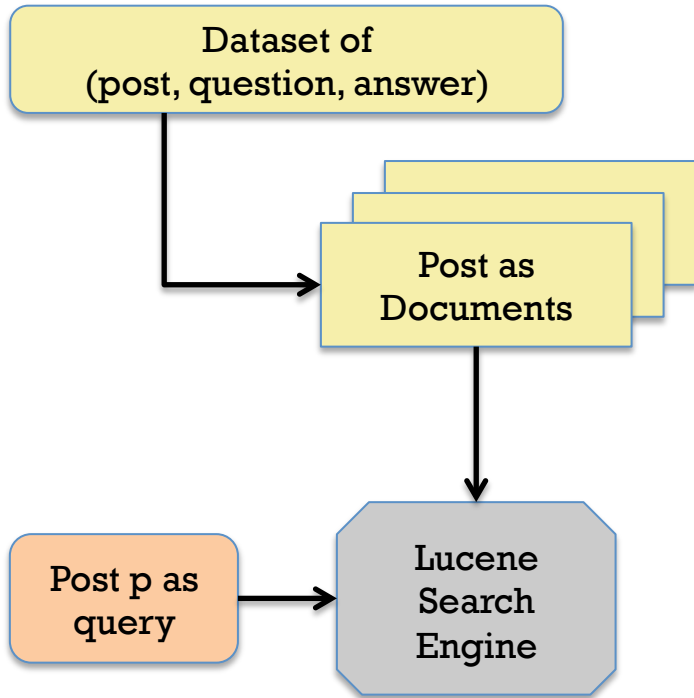
1

1

Question & Answer Candidate Generator

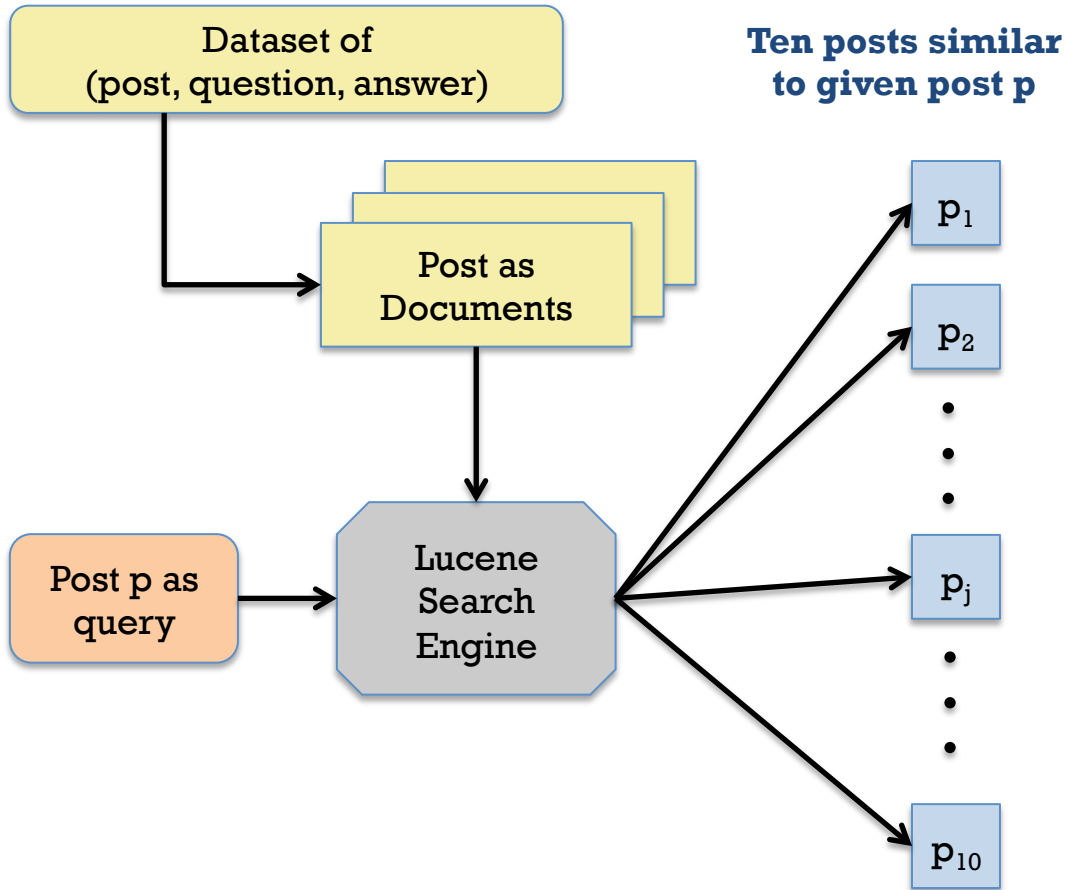
1

Question & Answer Candidate Generator



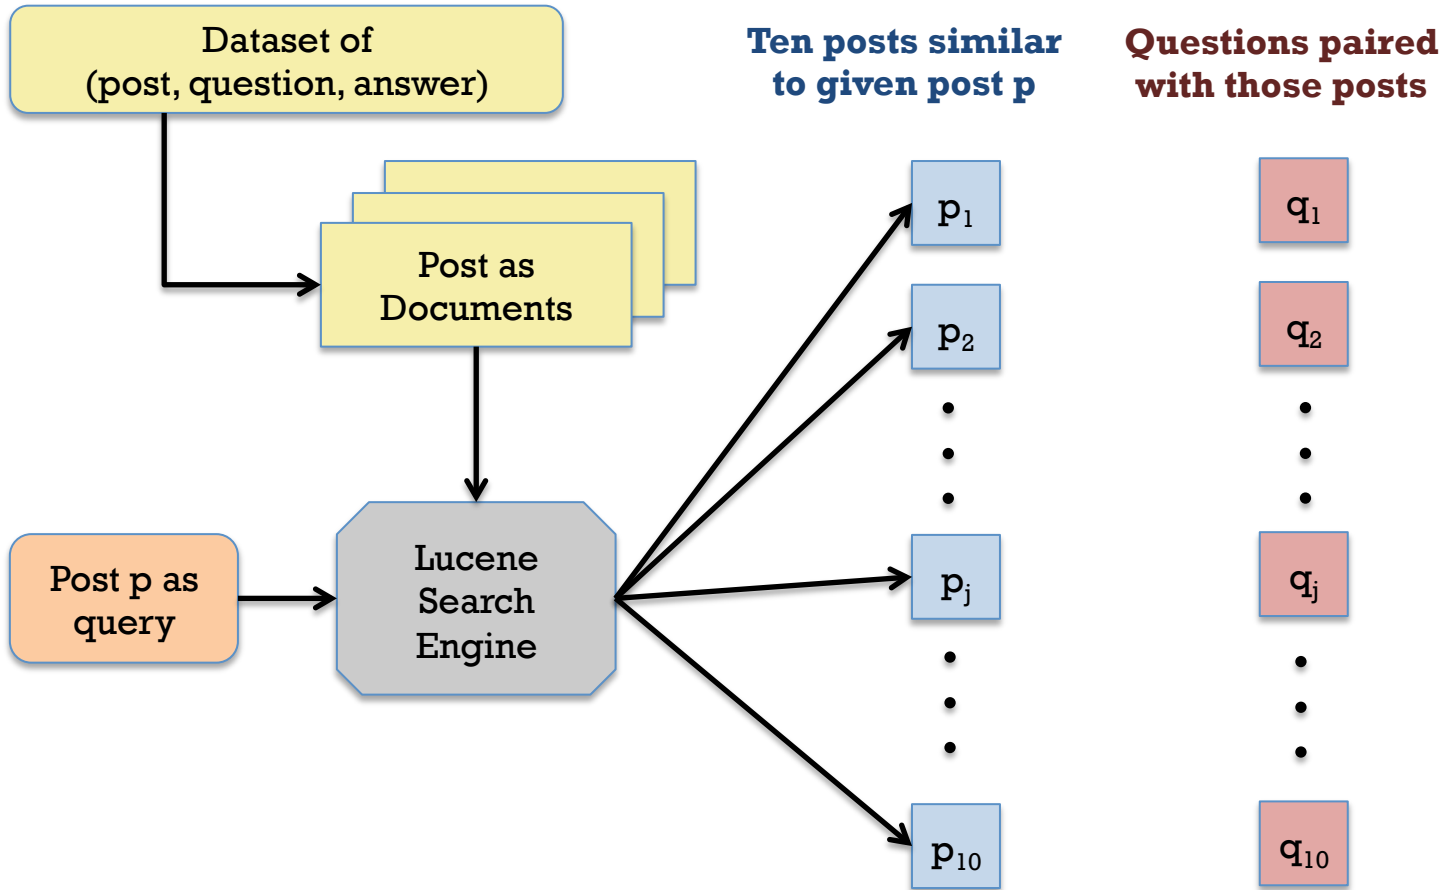
1

Question & Answer Candidate Generator



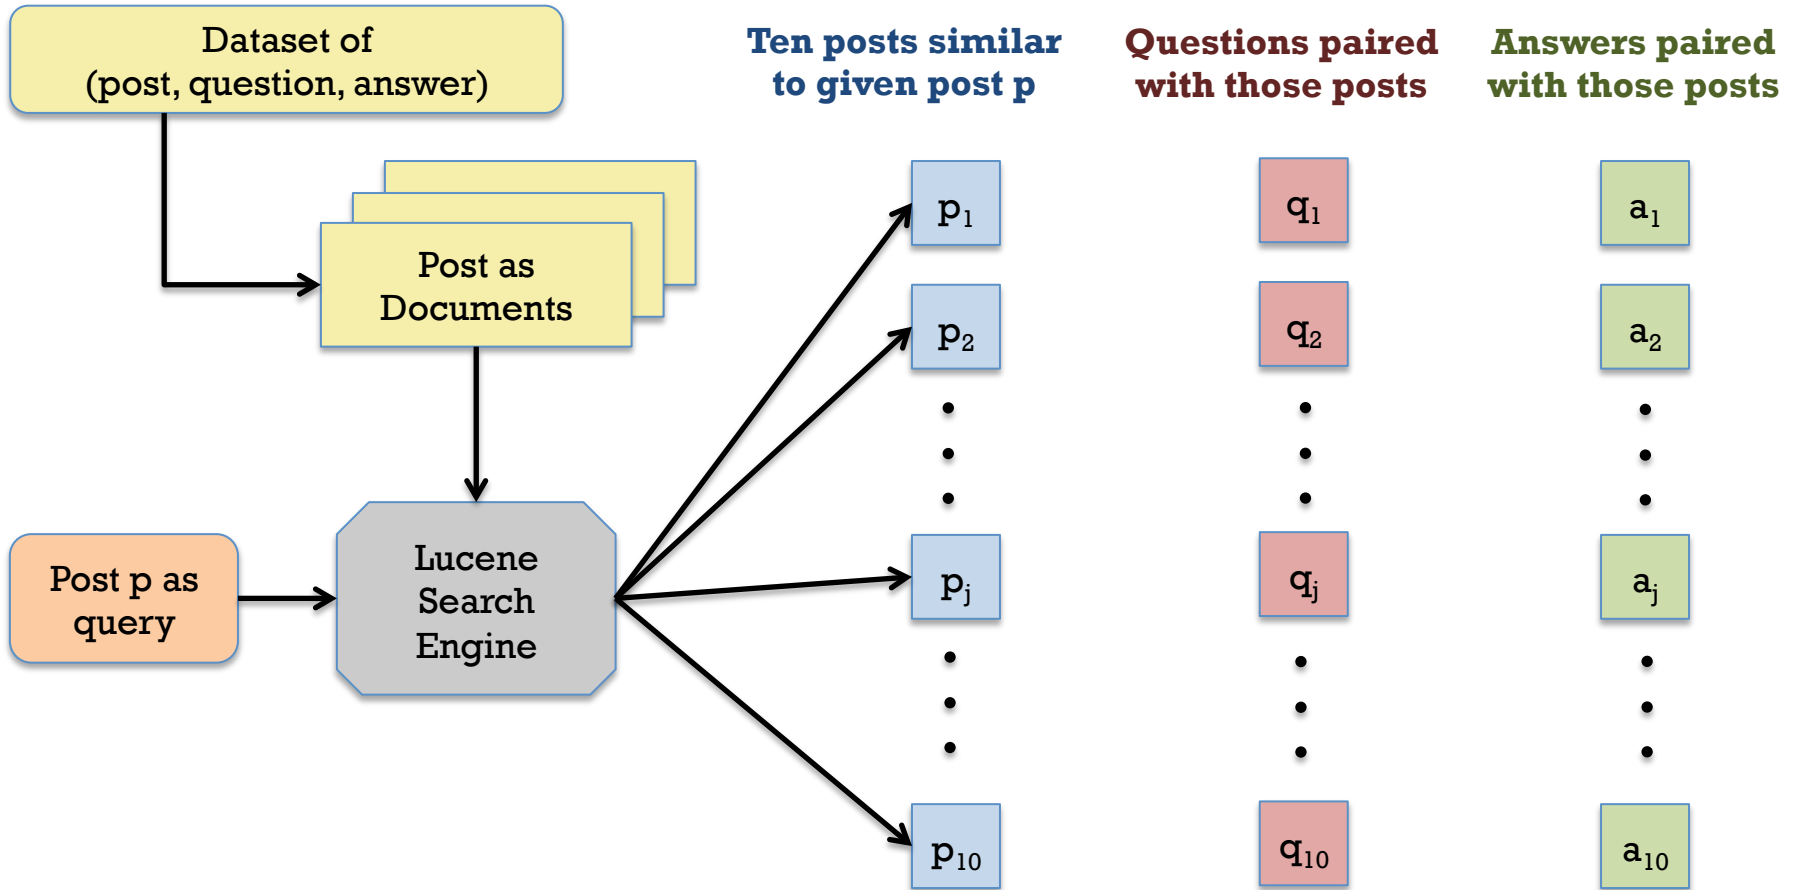
1

Question & Answer Candidate Generator



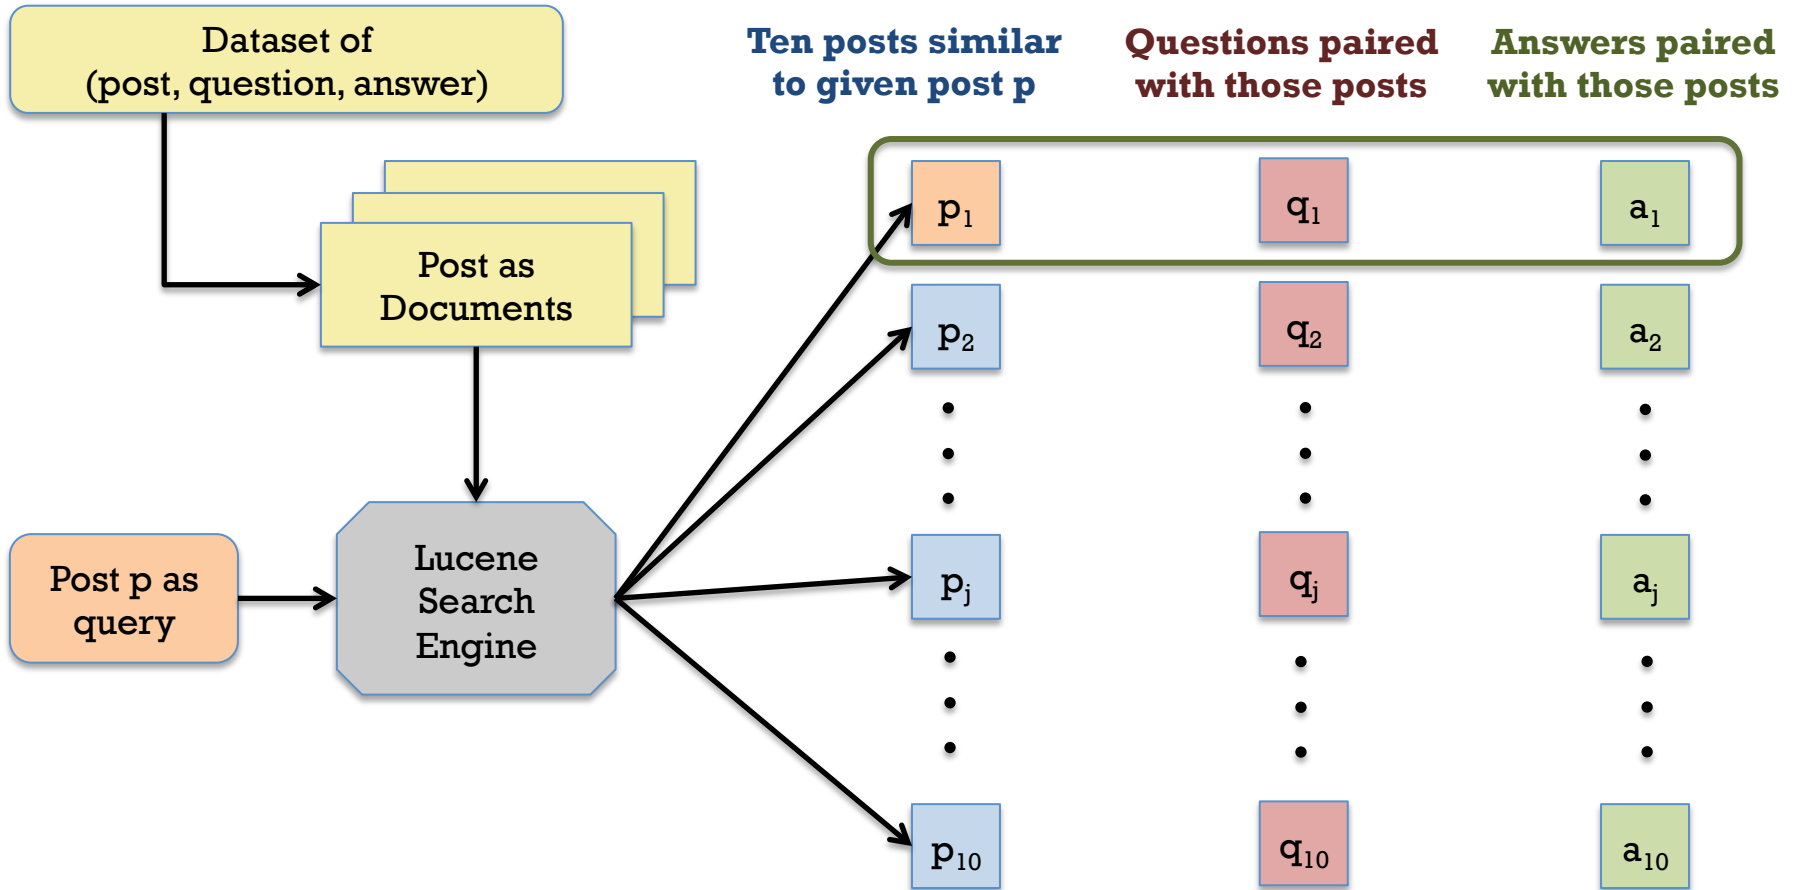
1

Question & Answer Candidate Generator



1

Question & Answer Candidate Generator



2

Answer Modeling



$$\mathbf{EVPI} (q_i | p) = \sum_{a_j \in \mathbf{A}} \mathbf{P} (a_j | p, q_i) \mathbf{U} (p + a_j)$$



2

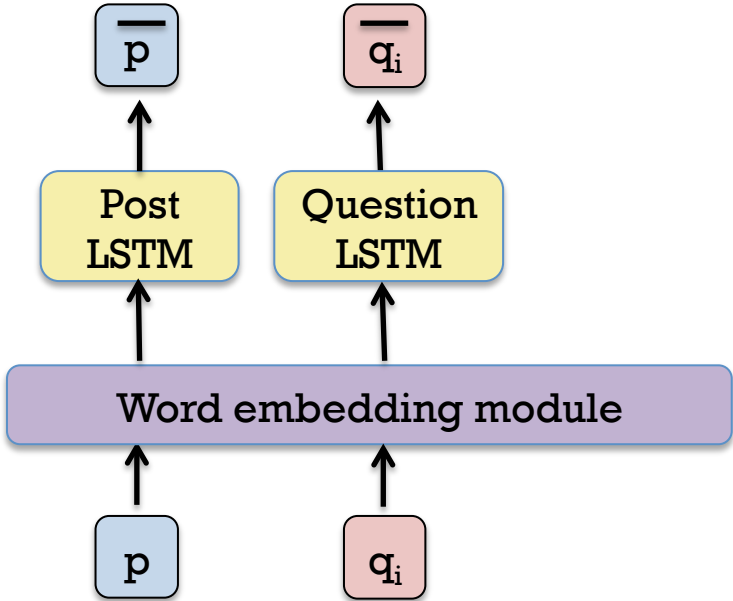
Answer Modeling

$$\mathbf{P} (a_j \mid p, q_i)$$

2

Answer Modeling

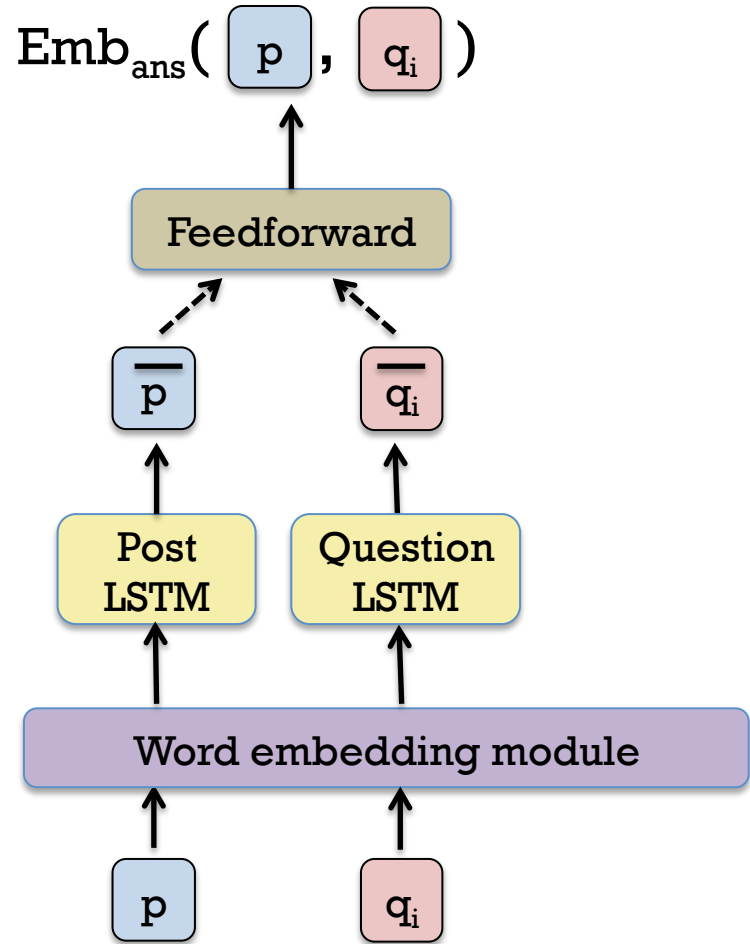
$$\mathbf{P} (a_j \mid p, q_i)$$



2

Answer Modeling

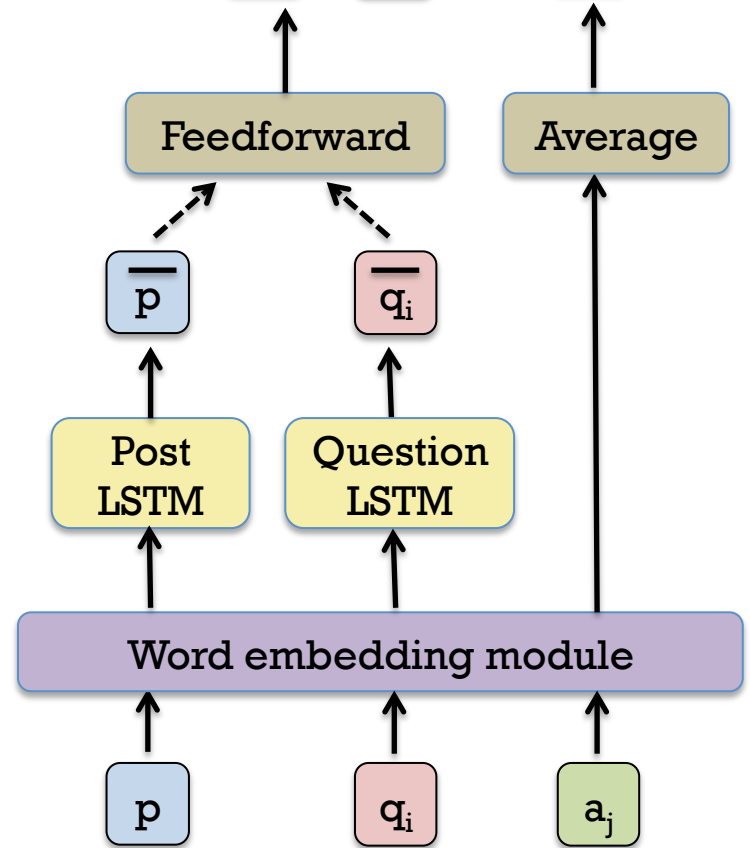
$$\mathbf{P} (a_j \mid p, q_i)$$



2

Answer Modeling

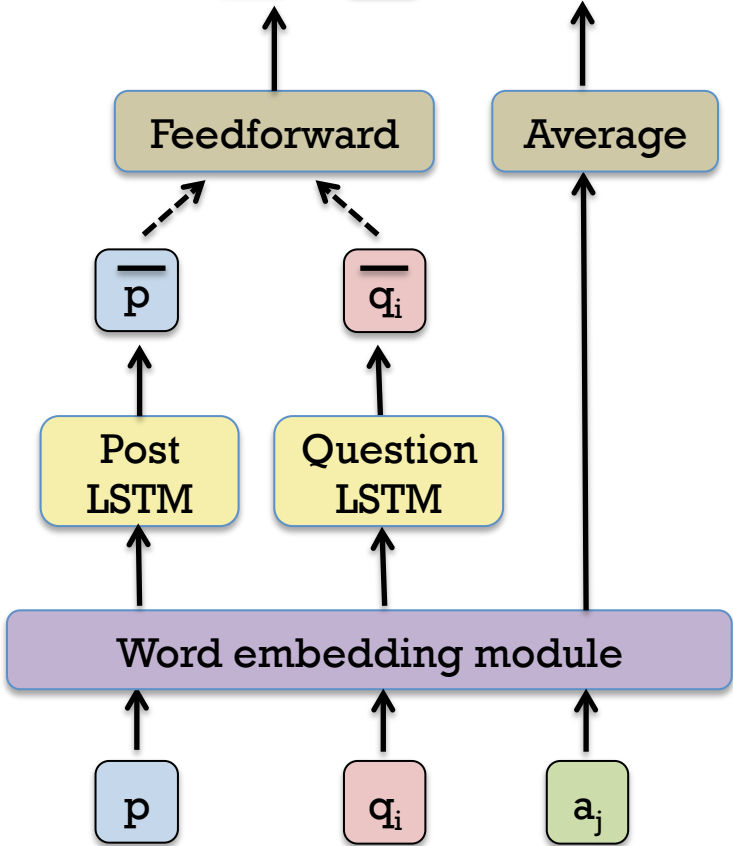

$$\mathbf{P} (a_j | p, q_i) \approx \text{cosine_sim} (\text{Emb}_{\text{ans}} (p, q_i), a_j)$$



2

Answer Modeling

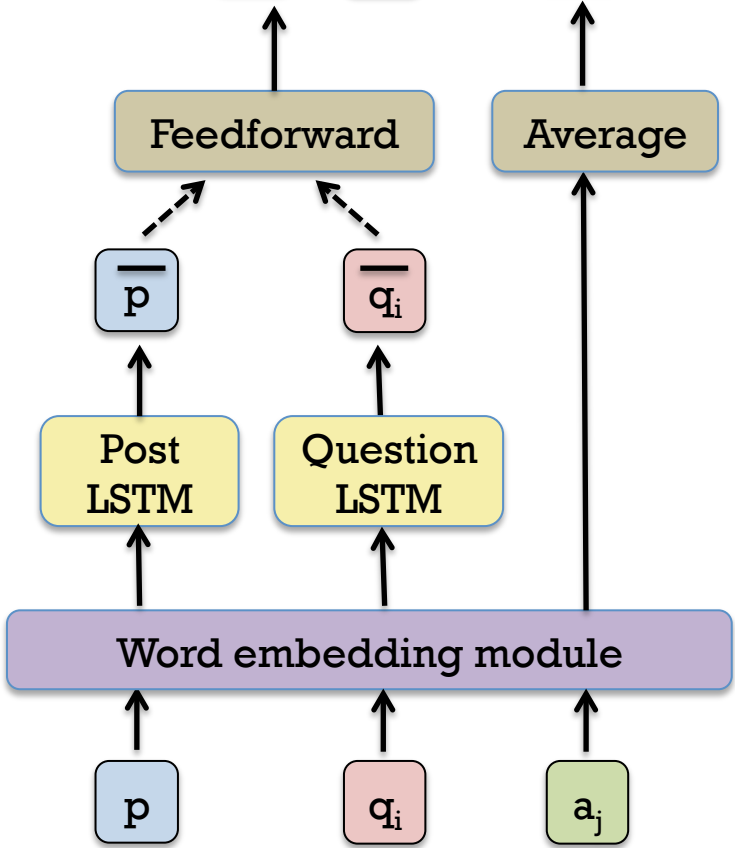
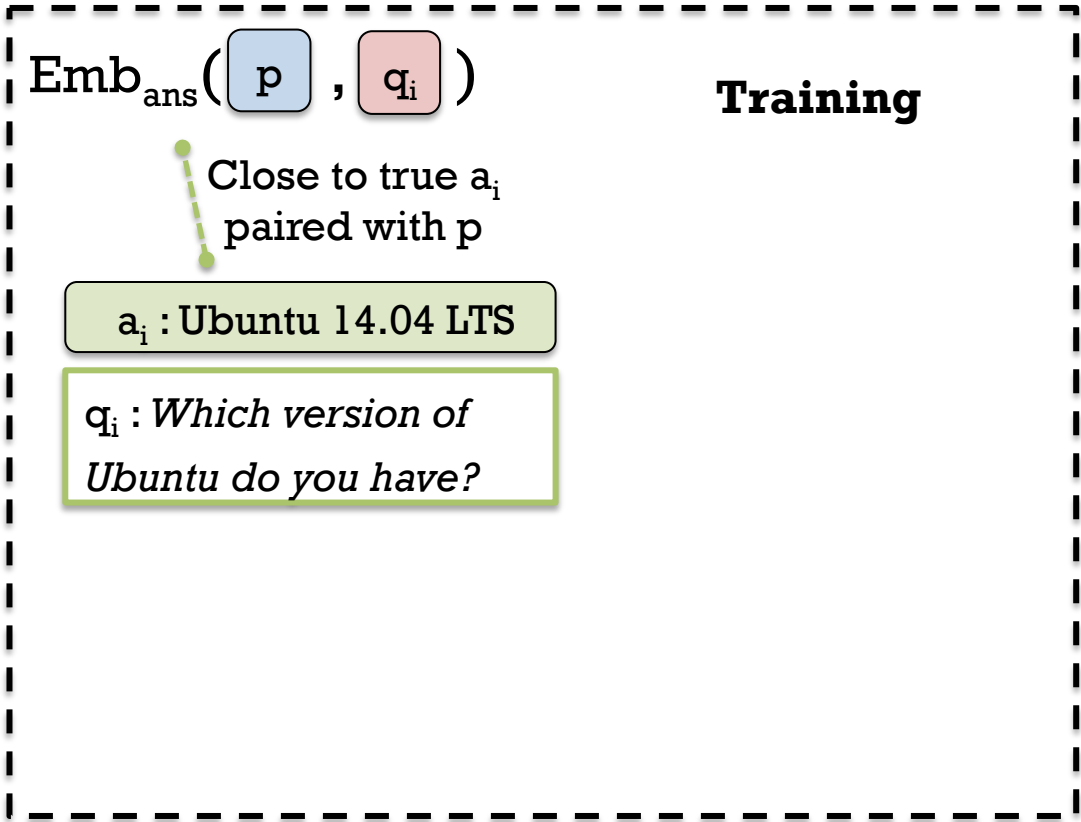
$$\mathbf{P} (a_j | p, q_i) \approx \text{cosine_sim} (\text{Emb}_{\text{ans}} (p, q_i), a_j)$$



2

Answer Modeling

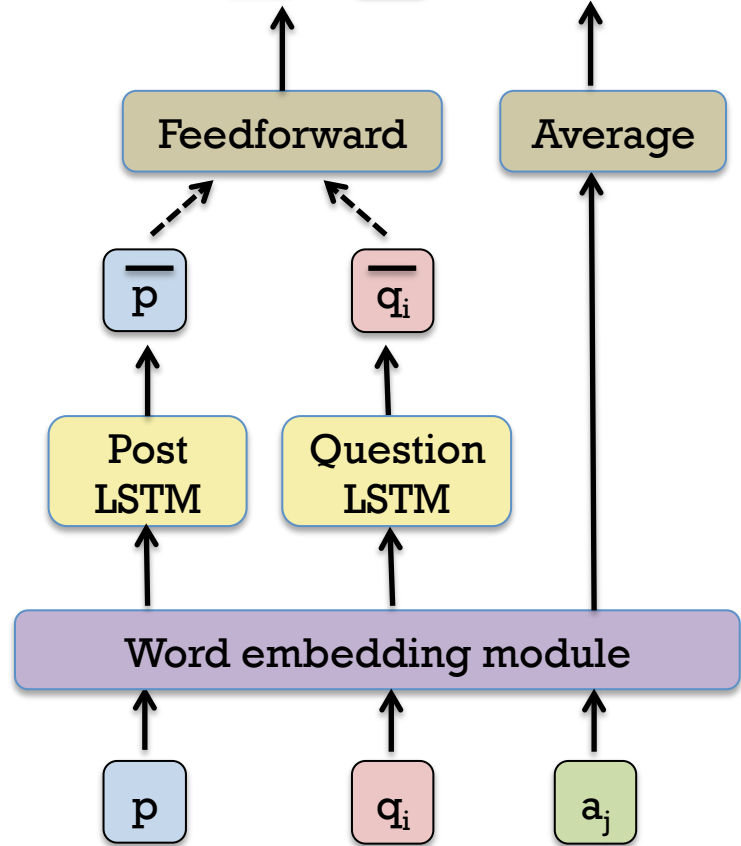
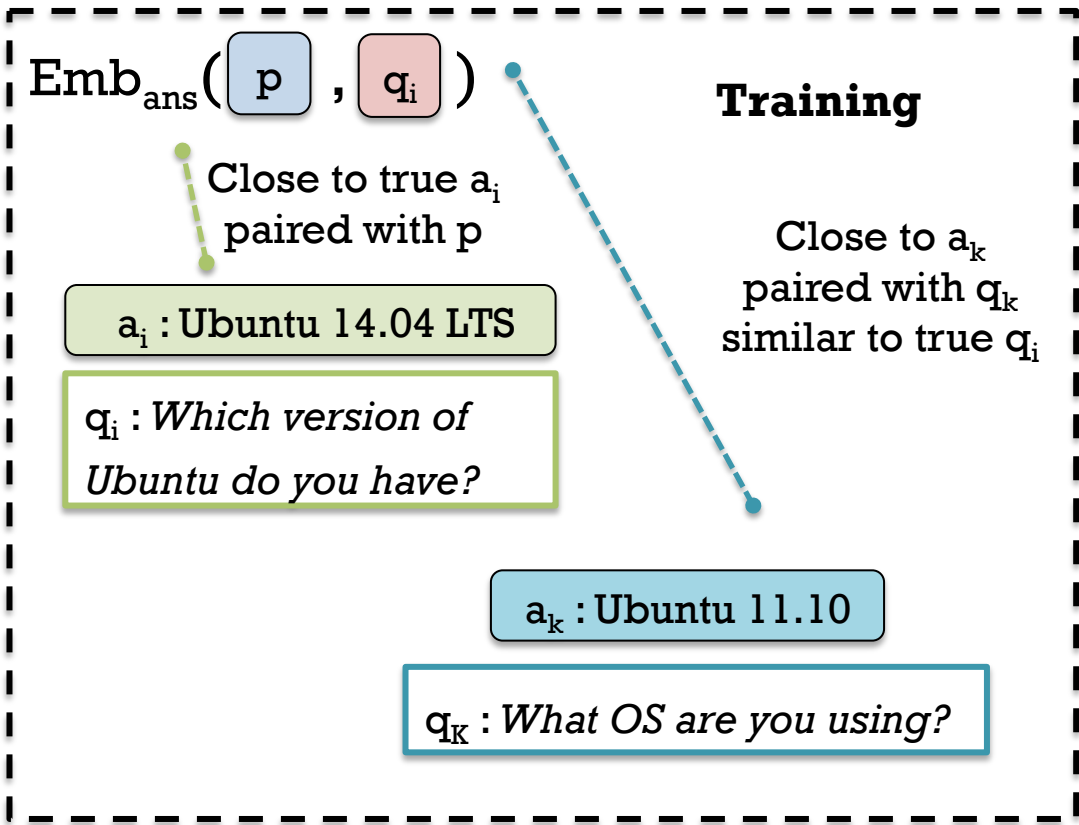
$$P(a_j | p, q_i) \approx \text{cosine_sim}(\text{Emb}_{\text{ans}}(p, q_i), a_j)$$



2


Answer Modeling

$$P(a_j | p, q_i) \approx \text{cosine_sim}(\text{Emb}_{\text{ans}}(p, q_i), a_j)$$



3

Utility Calculator



$$\mathbf{EVPI} (q_i | p) = \sum_{a_j \in \mathbf{A}} \mathbf{P} (a_j | p, q_i) \mathbf{U} (p + a_j)$$

3

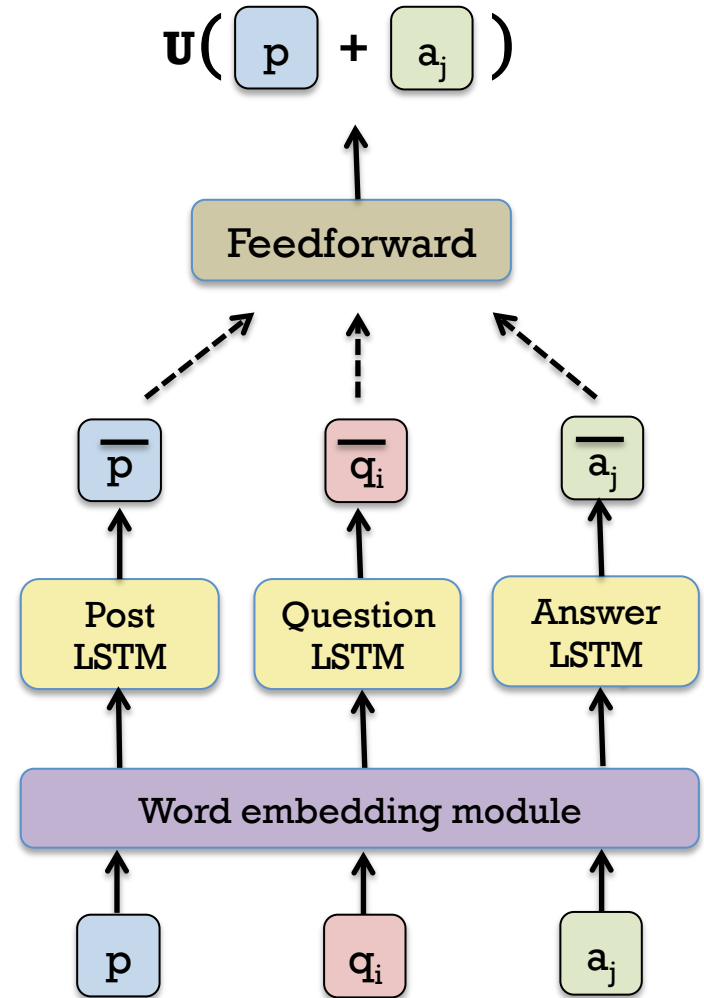
Utility Calculator

$$u(p + a_j)$$

3

Utility Calculator

Value between 0 and 1

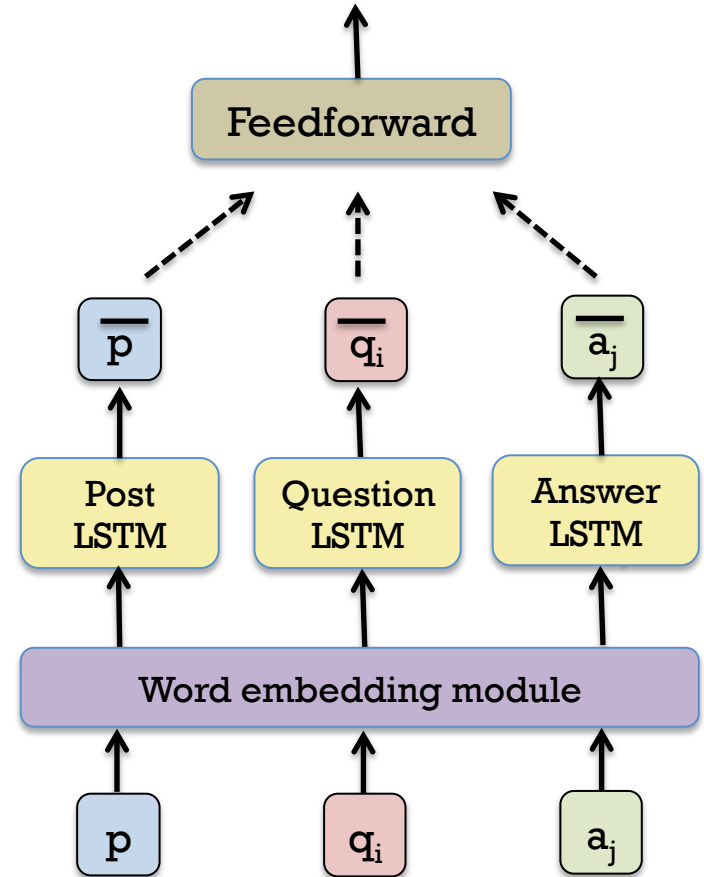



3

Utility Calculator

Value between 0 and 1

$$U(p + a_j)$$



3

Utility Calculator

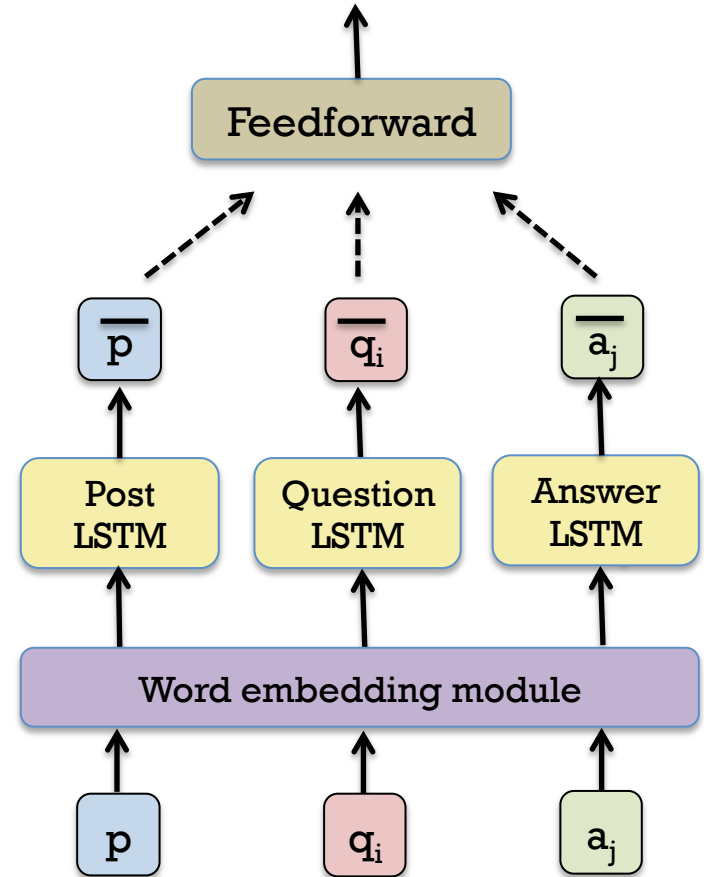
Value between 0 and 1

$$u(p + a_j)$$

Training

q_i : Which version of Ubuntu do you have?
 a_i : Ubuntu 14.04 LTS

Label
 $y = 1$

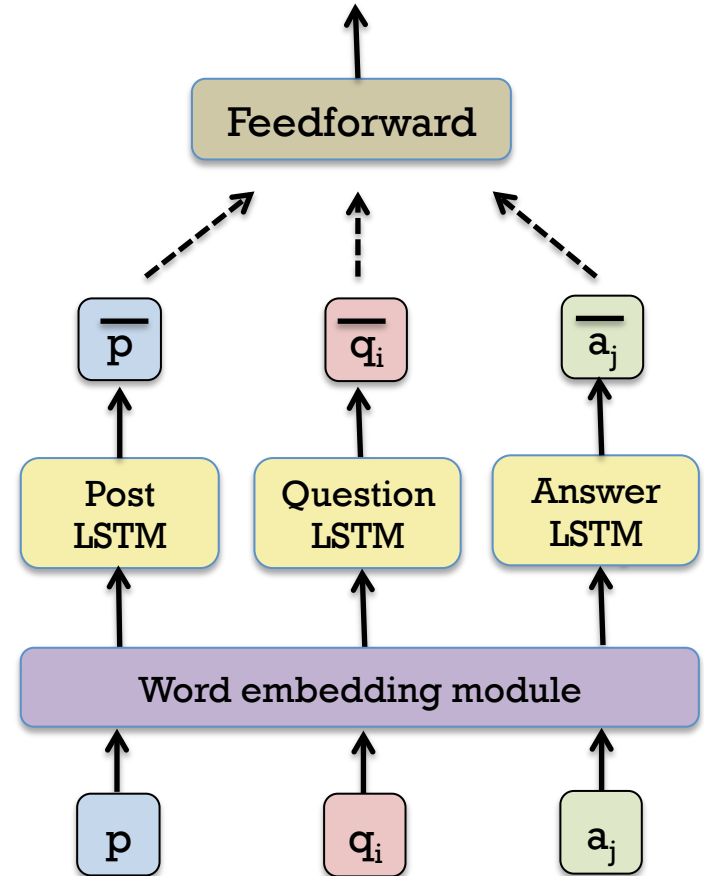
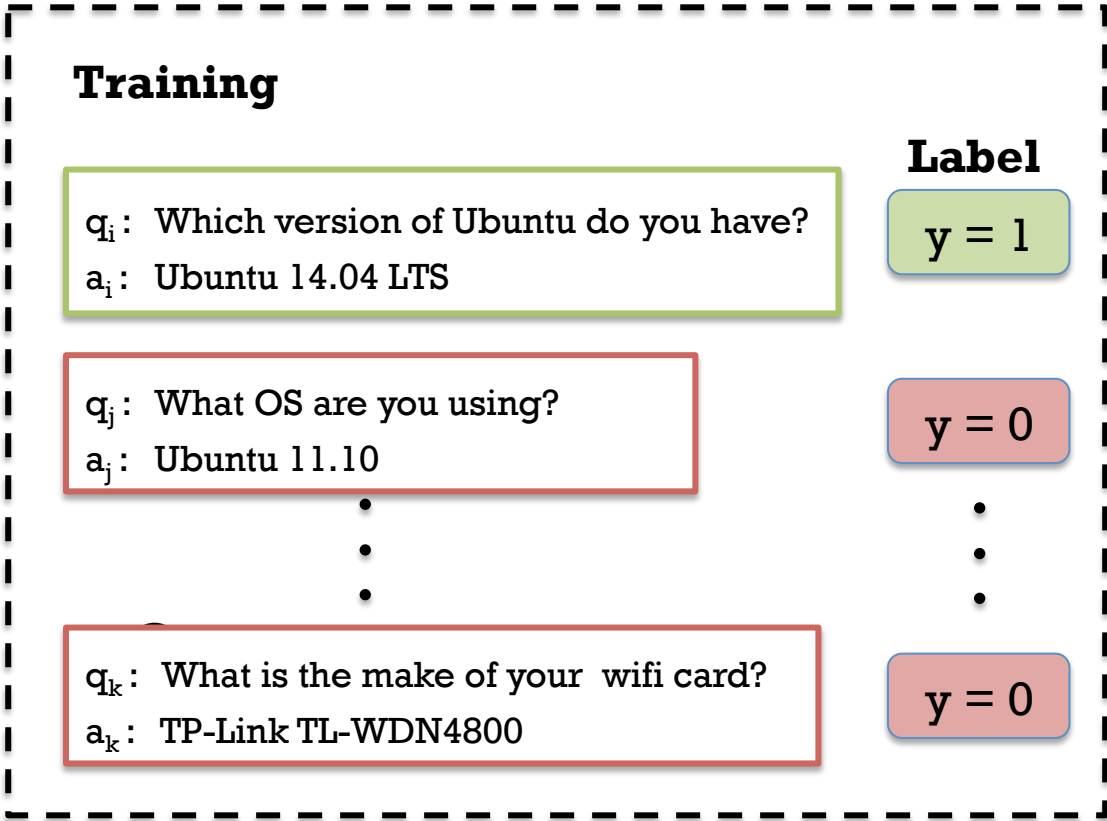


3

Utility Calculator

Value between 0 and 1

$$u(p + a_j)$$



3

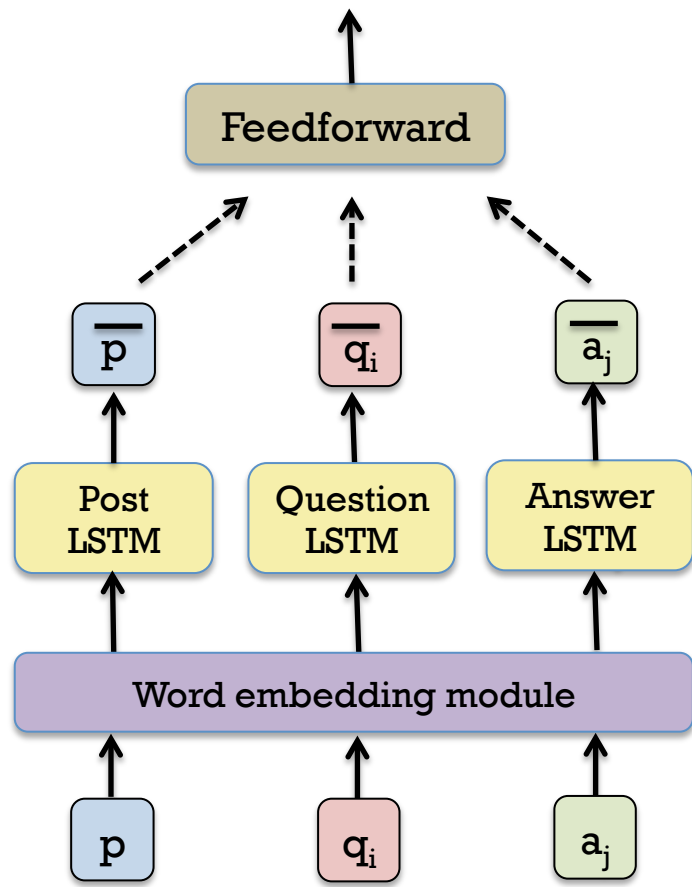
Utility Calculator

Value between 0 and 1

$$u(p + a_j)$$

Training (Minimize binary cross-entropy)

q_i : Which version of Ubuntu do you have? a_i : Ubuntu 14.04 LTS	Label $y = 1$
q_j : What OS are you using? a_j : Ubuntu 11.10 ⋮	$y = 0$ ⋮
q_k : What is the make of your wifi card? a_k : TP-Link TL-WDN4800	$y = 0$



Expected Value of Perfect Information (EVPI) inspired model

1

Question & Answer
Candidate Generator


2

Answer
Modeling

3

Utility
Calculator

Expected Value of Perfect Information (EVPI) inspired model

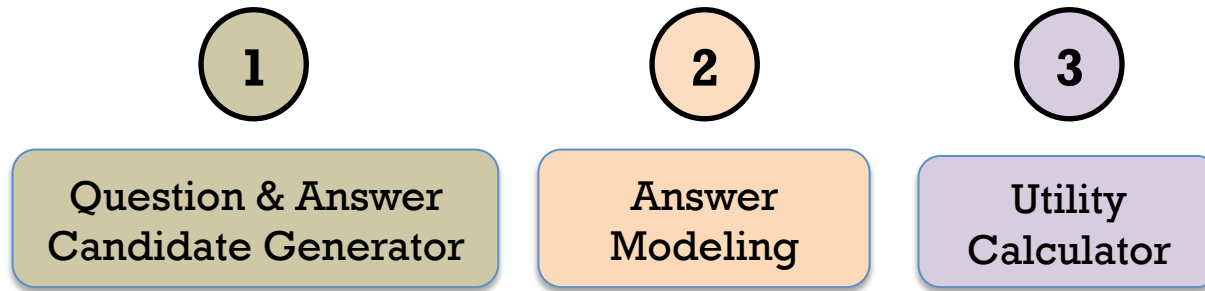


Train time behavior: For each (p, q, a) in our train set

1. Generate question candidates (Q) and answer candidates (A)
2. Train Answer Model and Utility Calculator

using joint loss function : $\text{loss}_{\text{ans}}(p, q, a, Q) + \text{loss}_{\text{util}}(y, p, q, a)$

Expected Value of Perfect Information (EVPI) inspired model



Train time behavior: For each (p, q, a) in our train set

1. Generate question candidates (Q) and answer candidates (A)
2. Train Answer Model and Utility Calculator
using joint loss function : $\text{loss}_{\text{ans}}(p, q, a, Q) + \text{loss}_{\text{util}}(y, p, q, a)$

Test time behavior: Given a post from our test set

1. Generate question candidates (Q) and answer candidates (A)
2. Calculate $P(a_j | p, q_i)$ for each $q \in Q$ using Answer Model
3. Calculate $U(p + a_j)$ for each $a \in A$ using Utility Calculator
4. Rank questions by EVPI $(q_i | p) = \sum_{a_j \in A} P(a_j | p, q_i) U(p + a_j)$

Talk Outline

- Clarification Questions Dataset
- Problem Formulation: Question Ranking
- Expected Value of Perfect Information (EVPI) inspired model
- Evaluation
- Conclusion

Too much disk read/write when launching an application

I have Xubuntu 13.04 on an old Dell Inspiron.

When I launch an application it takes a pretty long time to be launched and I see a lot of disk read/write.

If the system was short on memory, this would be understandable as the system would use swap. But that's not the case in my situation (i.e. I have this problem even when the RAM is almost empty).

Post

Question Candidates

1. How much ram do you have installed ? and what size is the swap disk partition ?
2. If you do not have any problem with getting a little techy then may i suggest a method ?
3. How is it slow exactly ? boot time ? hdd read/write ? cpu time ? graphics rendering ?
4. What is the longest time you have let it run ?
5. This may be a silly question but ... did you make your usb stick bootable ?
6. Do your system were recently updated ?
7. Why not have two ssds in raid 1 for redundancy ?
8. Is that a `parted -- list` on the synology device ?
9. Can you tell us a little about your configuration ?
10. Did you turn hardware virtualization on in bios/efi ?

Too much disk read/write when launching an application

I have Xubuntu 13.04 on an old Dell Inspiron.

When I launch an application it takes a pretty long time to be launched and I see a lot of disk read/write.

If the system was short on memory, this would be understandable as the system would use swap. But that's not the case in my situation (i.e. I have this problem even when the RAM is almost empty).

Question Candidates

Contains more than one good question!

1. **How much ram do you have installed ? and what size is the swap disk partition ?**
2. If you do not have any problem with getting a little techy then may i suggest a method ?
3. **How is it slow exactly ? boot time ? hdd read/write ? cpu time ? graphics rendering ?**
4. What is the longest time you have let it run ?
5. This may be a silly question but ... did you make your usb stick bootable ?
6. Do your system were recently updated ?
7. Why not have two ssds in raid 1 for redundancy ?
8. Is that a `parted -- list` on the synology device ?
9. **Can you tell us a little about your configuration ?**
10. **Did you turn hardware virtualization on in bios/efi ?**

Evaluation set creation process

- We recruit 10 Unix admin experts using UpWork
- Given a post and the set of ten question candidates
 - Mark the one best question
 - Mark any other valid questions

Evaluation set creation process

- We recruit 10 Unix admin experts using UpWork
- Given a post and the set of ten question candidates
 - Mark the one best question
 - Mark any other valid questions
- We annotate a total of 500 posts from our test set
- Each post is annotated by two experts

Evaluation set creation process

- We recruit 10 Unix admin experts using UpWork
- Given a post and the set of ten question candidates
 - Mark the one best question
 - Mark any other valid questions
- We annotate a total of 500 posts from our test set
- Each post is annotated by two experts
- **Union of Bests:** Questions marked as best by either of the annotators
- **Intersection of Valids:** Questions marked as valid by both annotators

	Best		Valid	
	A1	A2	A1	A2
Q1			✓	✓
Q2		✓	✓	
Q3	✓		✓	✓
Q4				✓
Q5			✓	✓

Union of Bests: { Q2, Q3 }

Intersection of Valids: { Q1, Q3, Q5 }

Baseline Models

1. **Random:** Randomly permute the 10 candidate questions

Baseline Models

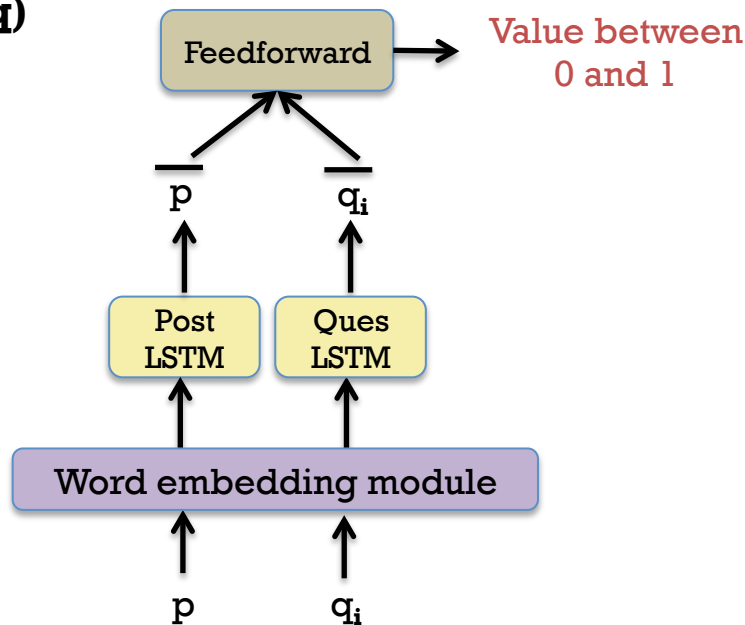
1. **Random:** Randomly permute the 10 candidate questions
2. **Bag-of-ngrams:** Train linear classifier using bag-of-ngrams of p, q and a

Baseline Models

1. **Random:** Randomly permute the 10 candidate questions
2. **Bag-of-ngrams:** Train linear classifier using bag-of-ngrams of p, q and a
3. **Community QA** (Nakov et al., 2017) :
 - SemEval Task: Rank comments by relevance to post on Qatar Living
 - Winning model: Logistic regression trained with string similarity & word embedding based features (Nandi et al., 2017)
 - Our baseline: We retrain this model on our dataset

Baseline Models

1. **Random:** Randomly permute the 10 candidate questions
2. **Bag-of-ngrams:** Train linear classifier using bag-of-ngrams of p , q and a
3. **Community QA** (Nakov et al., 2017) :
 - SemEval Task: Rank comments by relevance to post on Qatar Living
 - Winning model: Logistic regression trained with string similarity & word embedding based features (Nandi et al., 2017)
 - Our baseline: We retrain this model on our dataset
4. **Neural (p , q)**

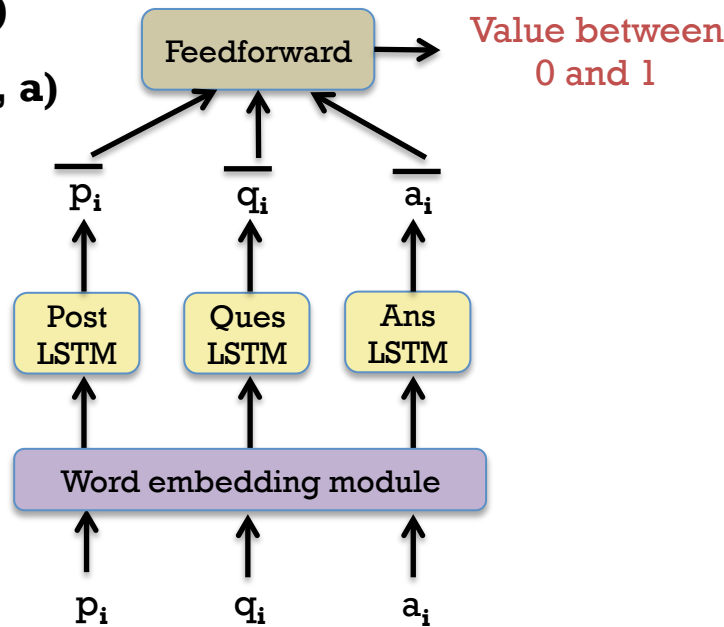


Baseline Models

1. **Random:** Randomly permute the 10 candidate questions
2. **Bag-of-ngrams:** Train linear classifier using bag-of-ngrams of p , q and a
3. **Community QA** (Nakov et al., 2017) :
 - SemEval Task: Rank comments by relevance to post on Qatar Living
 - Winning model: Logistic regression trained with string similarity & word embedding based features (Nandi et al., 2017)
 - Our baseline: We retrain this model on our dataset

4. Neural (p , q)

5. Neural (p , q , a)

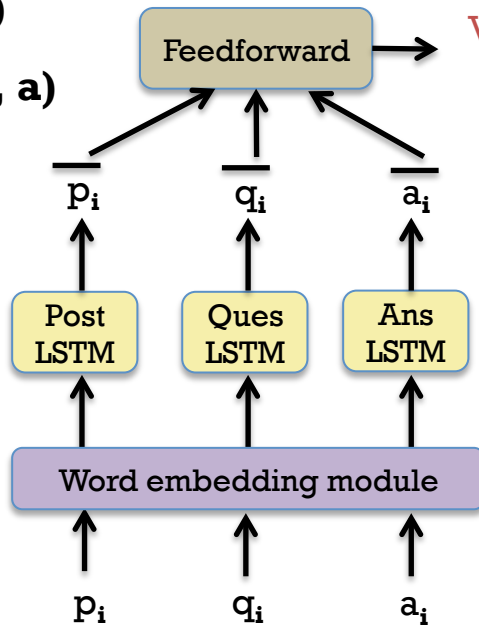


Baseline Models

1. **Random:** Randomly permute the 10 candidate questions
2. **Bag-of-ngrams:** Train linear classifier using bag-of-ngrams of p , q and a
3. **Community QA** (Nakov et al., 2017) :
 - SemEval Task: Rank comments by relevance to post on Qatar Living
 - Winning model: Logistic regression trained with string similarity & word embedding based features (Nandi et al., 2017)
 - Our baseline: We retrain this model on our dataset

4. Neural (p , q)


5. Neural (p , q , a)



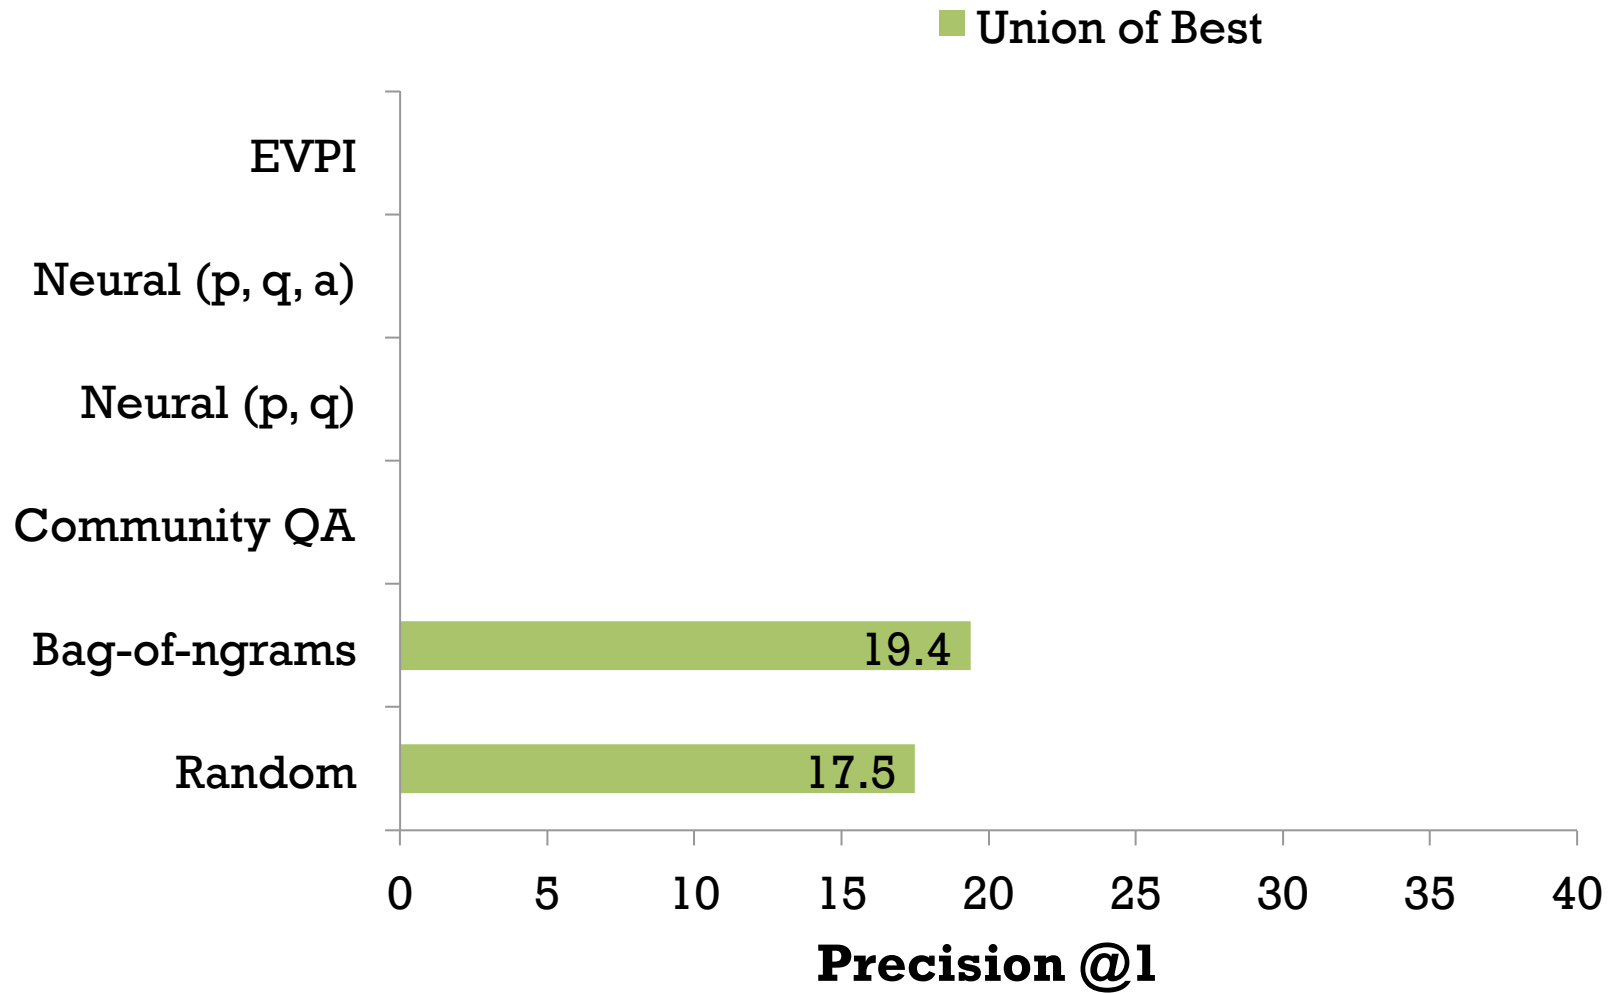
Value between
0 and 1

Both Neural (p , q , a) and EVPI (q | p , a)
have similar no. of parameters


RESULTS




RESULTS




RESULTS




RESULTS



RESULTS




RESULTS




Note: Difference between EVPI and all baselines is statistically significant with $p < 0.05$

RESULTS




Note: Difference between EVPI and all baselines is statistically significant with $p < 0.05$

RESULTS



Note: Difference between EVPI and all baselines is statistically significant with $p < 0.05$

RESULTS



Talk Outline

- Clarification Questions Dataset
- Problem Formulation: Question Ranking
- Expected Value of Perfect Information (EVPI) inspired model
- Evaluation
- Conclusion

✓ **Key Contributions**

- Create a **dataset** of ~77K clarification questions (and answers) with context
- Introduce novel **model** that integrates deep learning with classic notion of expected value of perfect information
- Create an **evaluation set** of size 500 with expert human annotations

✓ **Key Contributions**

- Create a **dataset** of ~77K clarification questions (and answers) with context
- Introduce novel **model** that integrates deep learning with classic notion of expected value of perfect information
- Create an **evaluation set** of size 500 with expert human annotations

✓ **Key findings**

- A context can have multiple good clarification question
- Explicitly modeling the answer helps in identifying good questions
- EVPI formalism provides leverage over similarly expressive feedforward network

✓ **Key Contributions**

- Create a **dataset** of ~77K clarification questions (and answers) with context
- Introduce novel **model** that integrates deep learning with classic notion of expected value of perfect information
- Create an **evaluation set** of size 500 with expert human annotations

✓ **Key findings**

- A context can have multiple good clarification question
- Explicitly modeling the answer helps in identifying good questions
- EVPI formalism provides leverage over similarly expressive feedforward network

✓ **Future work**

- Sequence-to-sequence based question generation model
- Multi-turn question generation
- How to automatically evaluate performance?

✓ **Key Contributions**

- Create a **dataset** of ~77K clarification questions (and answers) with context
- Introduce novel **model** that integrates deep learning with classic notion of expected value of perfect information
- Create an **evaluation set** of size 500 with expert human annotations

✓ **Key findings**

- A context can have multiple good clarification question
- Explicitly modeling the answer helps in identifying good questions
- EVPI formalism provides leverage over similarly expressive feedforward network

✓ **Future work**


- Sequence-to-sequence based question generation model
- Multi-turn question generation
- How to automatically evaluate performance?


CODE + DATA: https://github.com/raosudha89/ranking_clarification_questions

Backup Slides

RESULTS

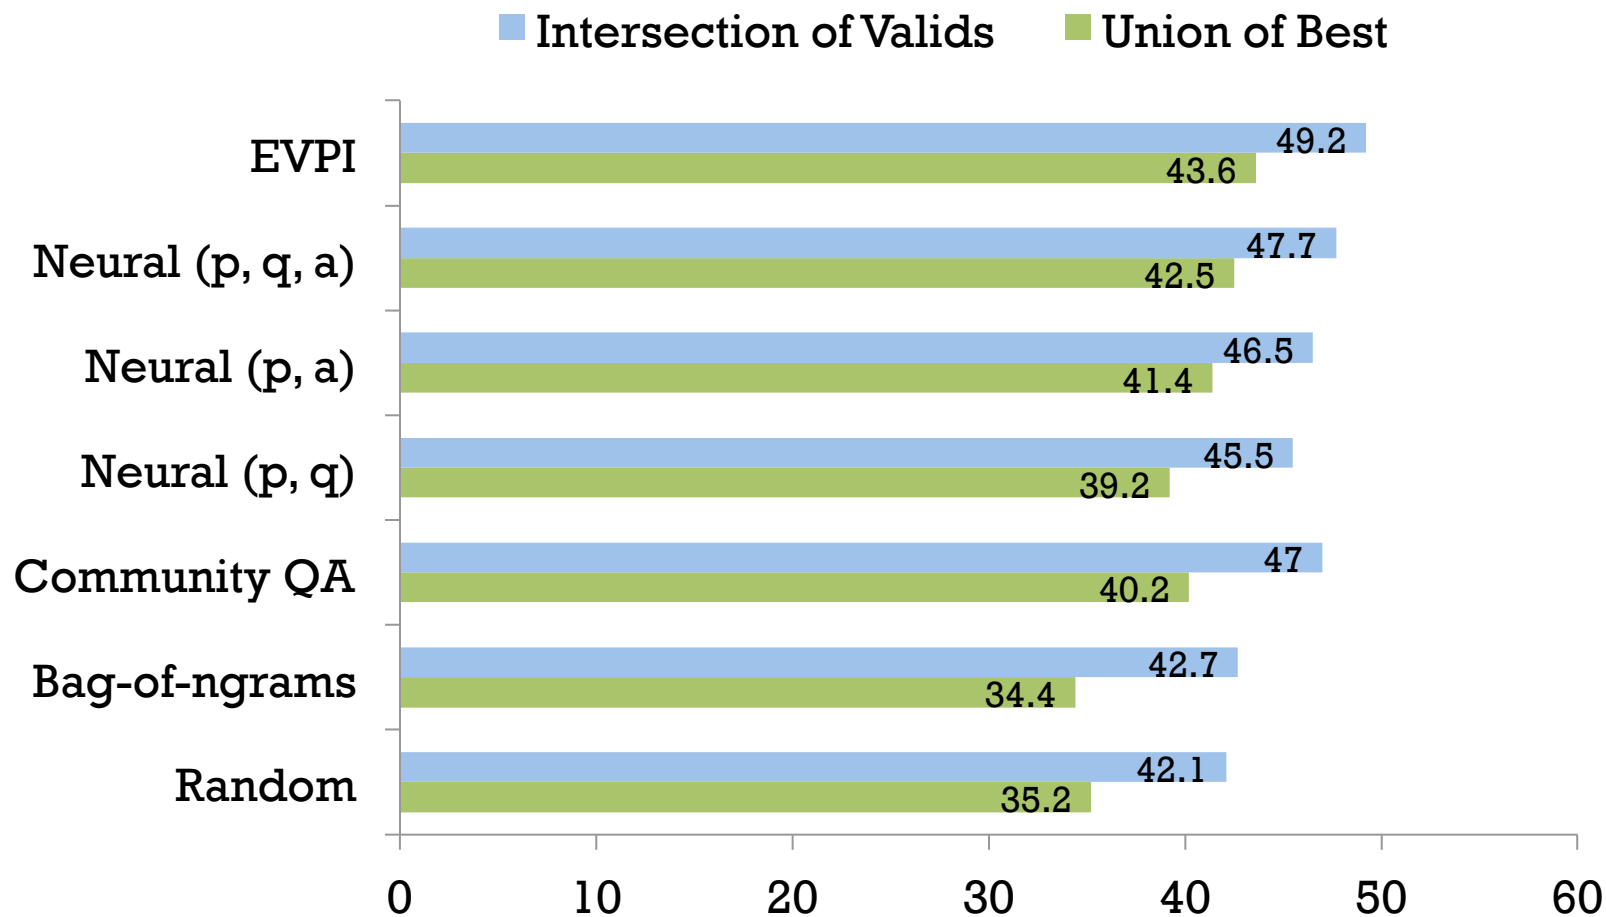
Precision @ 1



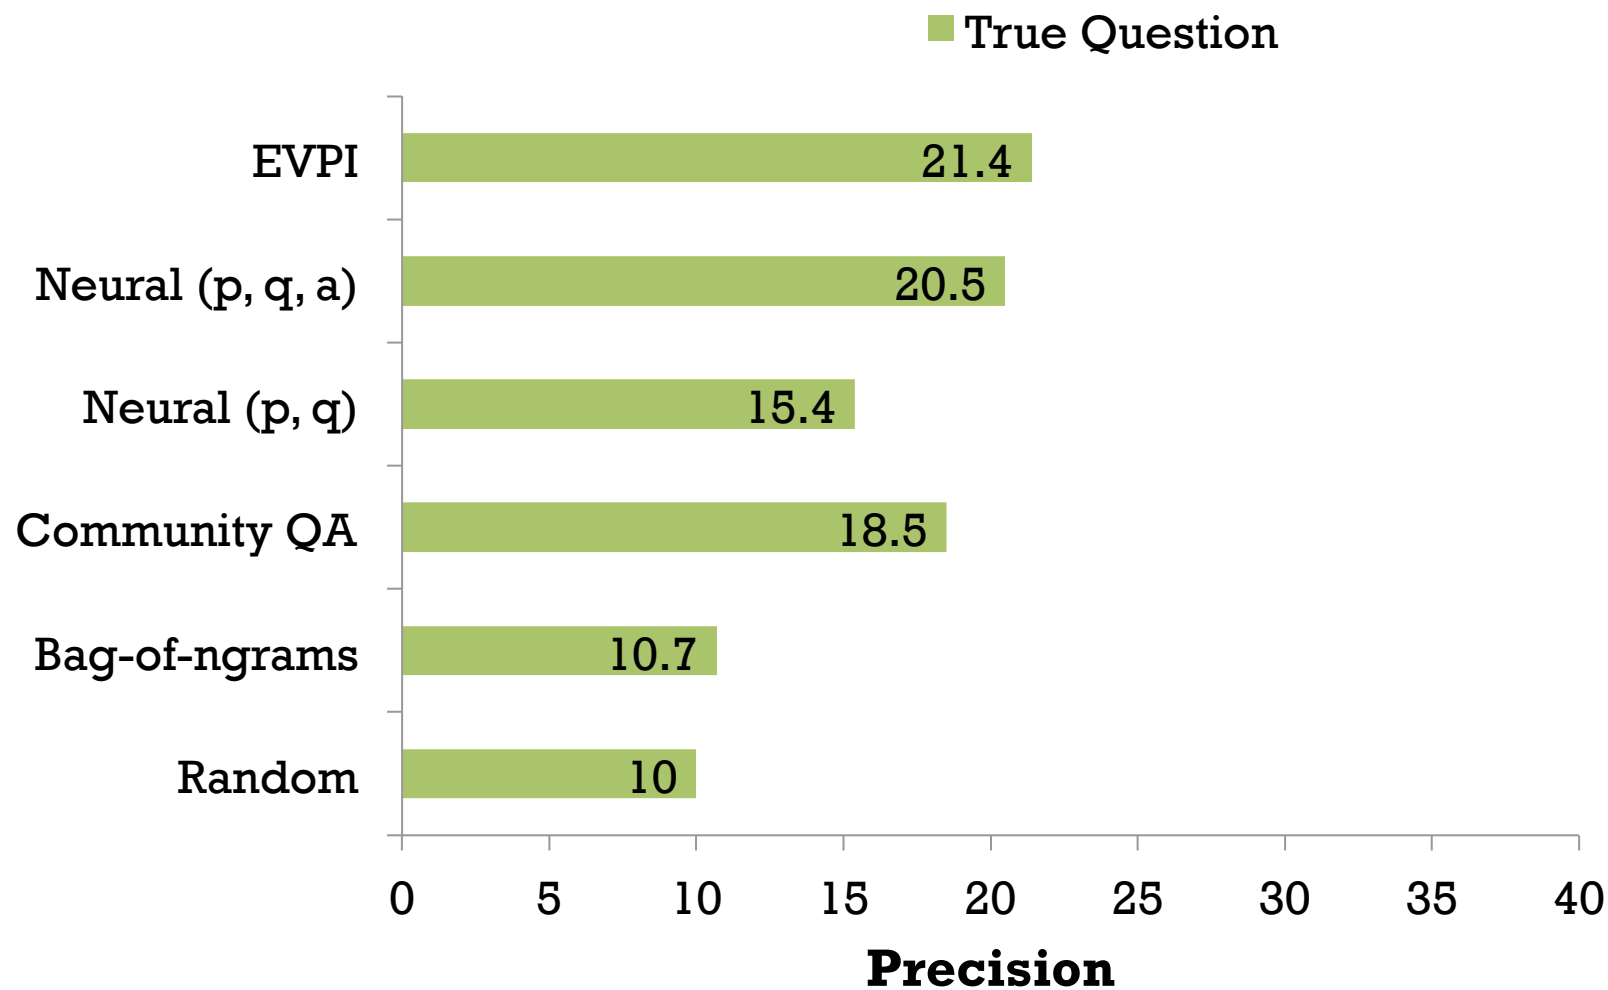
RESULTS**Precision @3**

RESULTS

Mean Average Precision




RESULTS



Note: Difference between EVPI and all baselines is statistically significant with $p < 0.05$

RESULTS



Too much disk read/write when launching an application

I have Xubuntu 13.04 on an old Dell Inspiron.

When I launch an application it takes a pretty long time to be launched and I see a lot of disk read/write.

If the system was short on memory, this would be understandable as the system would use swap. But that's not the case in my situation (i.e. I have this problem even when the RAM is almost empty).

Post

EVPI value

Ranking of Question Candidates

Best Valid

- | | | | |
|------|----------------------------------|-----------------------|---|
| 0.21 | <input checked="" type="radio"/> | <input type="radio"/> | How much ram do you have installed? and what size is the swap disk partition |
| 0.18 | <input type="radio"/> | <input type="radio"/> | Can you tell us a little about your configuration ? |
| 0.17 | | | What is the longest time you have let it run ? |
| 0.11 | <input checked="" type="radio"/> | <input type="radio"/> | How is it slow exactly ? boot time ? hdd read/write ? cpu time ? |
| 0.00 | | | If you do not have any problem with getting a little techy may i suggest a method ? |
| 0.00 | | | This may be a silly question but ... did you make your usb stick bootable ? |
| 0.00 | | | Do your system were recently updated ? |
| 0.00 | | | Why not have two ssds in raid 1 for redundancy ? |
| 0.00 | | | Is that a `parted -- list` on the synology device ? |
| 0.00 | <input type="radio"/> | <input type="radio"/> | Did you turn hardware virtualization on in bios/efi ? |

SAMPLE OUTPUT

No wifi after restart in Ubuntu 16.04

After upgrading to 16.04, there is no wifi whenever I restart the system. My wireless interface of Ubuntu is RT3290 Wireless 802.11n 1T/1R PCIe

On iwconfig I got the following
eth0 no wireless extensions....

Currently to start wifi again I have to shutdown, then boot the system again. How to fix the problem?

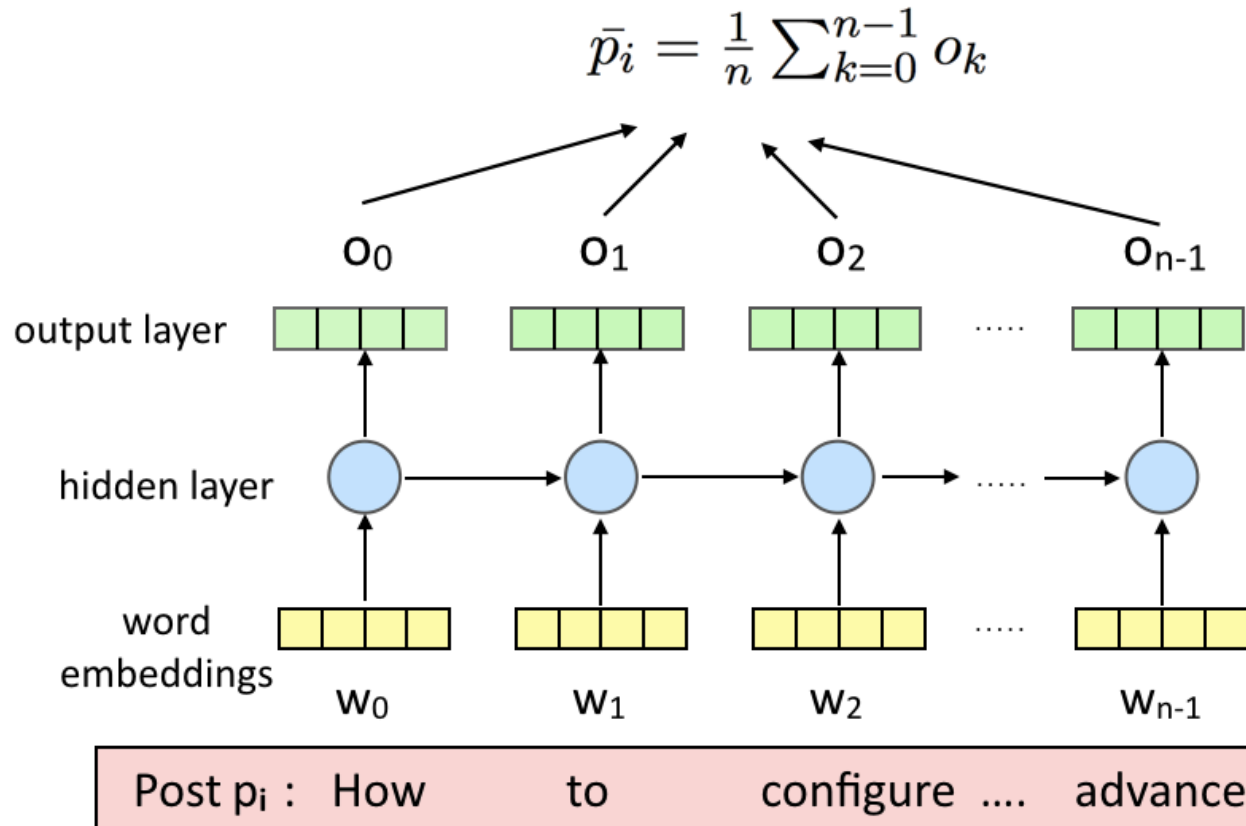
EVPI value

Ranking of Question Candidates

Best Valid

- 0.24 **I doubt it, shutdown and reboot are exactly identical! are you really rebooting?**
- 0.13 **Be clear about the problem. Is Ubuntu not showing them even though they are present?**
- 0.11 What is 4g wifi connection?
- 0.09 **Can you type `iwconfig` in terminal and paste what it returns here?**
- 0.09 What does this tell us?
- 0.08 If I post it as an answer, would you kindly mark as such?
- 0.06 Which Ubuntu 15?
- 0.06 What exactly do you mean by make fails?
- 0.05 **Welcome to ask Ubuntu! ; -) Is the wireless lan disabled in the bios?**
- 0.00 **Is Ubuntu detecting your wireless card ? `**iwconfig**` does list your card?**

LONG SHORT TERM MEMORY (LSTM)



Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. "Glove: Global vectors for word representation" In *Empirical Methods on Natural Language Processing*.