



Open Questions for Next Generation Chatbots

By Winson Ye and Qun Li
College of William and Mary



Introduction

1. Roadmap
 - a. Knowledge Grounding
 - i. Knowledge Domain
 - ii. Fact Collection
 - iii. Knowledge Grounding Architecture
 - b. Persona Consistency
 - i. Persona Collection
 - ii. Persona Definition
 - iii. Persona Training
 - c. Emotional Intelligence
 - i. Dimensions of Emotional Intelligence
 - ii. Emotion Databases
 - iii. Emotion Grounding Architecture
 - d. Evaluation Methods
 - i. Human Evaluation
 - ii. Automatic Evaluation
 - iii. New Metrics

Knowledge Grounding Domain

- Types of Information
 - Generic and domain specific knowledge
- Sources
 - Wikipedia
 - IMDb
 - Reddit
 - Twitter
- What are some future research challenges?
 - Can we learn from both general purpose and domain specific facts?
 - General purpose: Wikipedia
 - Domain specific: Stack Overflow
 - Can we take advantage of question and answer forums?
 - How can we measure the quality of a knowledge base?

The image shows a screenshot of a Stack Overflow question and its top answer. The question is titled "Get the current URL with JavaScript?" and has 2688 votes. The question body asks for the current website URL. The top answer, with 3327 votes, provides the code `window.location.href` and is marked as an "Accepted Answer". Annotations with lines point to the title, question body, tags, score, and accepted answer.

Title: Get the current URL with JavaScript?

Question Body: All I want is to get the website URL. Not the URL as taken from a link. On the page loading I need to be able to grab the full, current URL of the website and set it as a variable to do with as I please.

Tags: javascript url

Score: 2688

Score: 3327

Accepted Answer: Use:
`window.location.href`
As noted in the comments, the line below works, but it is bugged for Firefox.
`document.URL;`
See [URL of type DOMString, readonly.](#)

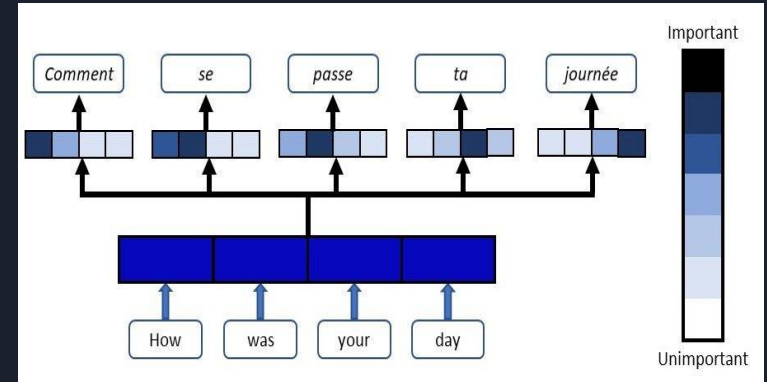


Knowledge Grounding Fact Collection

- Two datasets are used for knowledge grounding
 - Dataset of external facts
 - Dataset of conversational exchanges
- What are some future research challenges?
 - How can we filter out offensive language from conversational exchange datasets?
 - How can we make it easier to collect large amounts of facts and conversations?

Knowledge Grounding Architecture

- Two main approaches
 - Encoder based
 - Memory networks
- What are some future research challenges?
 - How can we develop attention mechanisms for facts?
 - How can we speed up the training process?
 - How can we leverage Transformers?



Persona Collection

- Common Datasets
 - Persona Chat
 - Wikipedia
 - Reddit
- Techniques
 - Scrape social media forums
 - Learn the personas
- What are some future research challenges?
 - How can we filter out toxic personas?
 - Can we extract personas from live conversations?

Persona:

- I love to drink fancy tea.
- I have a big library at home.
- I'm a museum tour guide.
- I'm partly deaf.

Persona:

- I have two dogs.
- I like to work on vintage cars.
- My favorite music is country.
- I own two vintage Mustangs.



Hello, how are you doing?



Nice! I'm not much of a music fan myself, but I do love to read.



Great thanks, just listening to my favorite Johnny Cash album!



Me too! I just read a book about the history of the auto industry.

Persona Definition

- Dimensions of Personality
 - Collection of background facts
 - Behavioral traits
 - Likes/Dislikes
- What are some future research challenges?
 - How can we capture multiple dimensions of personality?
 - How do we measure how well the chatbot expresses a personality?
 - How can we teach the chatbot to change its writing style to match its persona?





Persona Training

- Three approaches for training
 - Memory network
 - Simple MLP
 - Transformers
- What are some future research challenges?
 - What features are chatbots learning to exhibit personalities?
 - How can we use reinforcement learning?
 - How can we speed up the training process?

Dimensions of Emotional Intelligence

- Dimensions of Emotional Intelligence
 - Empathy
 - Politeness
 - Understanding Mood
- What are some future research challenges?
 - How can we incorporate multiple dimensions of emotional intelligence into a chatbot?
 - What dimensions of emotional intelligence are most important for having an interesting conversation?





Emotion Databases

- Data Collection Techniques
 - Collect conversational data from emotional situations
 - Train a classifier to detect emotion, then automatically annotate text corpus
- What are some future research challenges?
 - How can the chatbot learn emotions from live conversations?
 - How can we synthesize more complex emotions from a few primary ones?
 - How can we leverage multiple emotion datasets from different domains?

Emotion Grounding Architecture

- Techniques
 - Use a seq2seq backbone
 - Use Transformers
 - Use a seq2seq backbone along with memory networks
 - Use a seq2seq backbone with CNNs
- What are some future research challenges?
 - How can we develop an adversarial approach to training the chatbot?
 - How can we develop an attention mechanism for emotions?
 - How can we use an ensemble approach?

Case 1	Case 2
	
Caption "Hanging out on Saturday." Question "Did you have a good time?"	
Response 1 "We had a great time at the beach!"	Response 2 "She just hates going for a walk!"




Human Evaluation

- Experimental Guidelines
 - Use Cohen's kappa score
 - Set expectations clearly beforehand
 - Replicate studies with larger group of participants
- What are some future research challenges?
 - How can we develop new metrics for determining the reliability of human evaluation?
 - What are some strengths and pitfalls among all the existing human evaluation schemes?
 - How can we integrate human evaluation with automatic evaluation?

Automatic Evaluation

- Weak correlation between automatic and human evaluation
- What are some future research challenges?
 - How can we develop new metrics to evaluate automatic evaluation schemes?
 - How can we secure automatic evaluation schemes against attacks?





Designing New Metrics

- Proposals So Far
 - Combine human and automatic evaluation
 - Train an evaluator
- What are some future research challenges?
 - How can we use reinforcement learning?
 - How can we design metrics that combine multiple dimensions of response quality?



Conclusion

- Contact Information
 - we@emil.wm.edu