# Conversations Based on Search Engine Result Pages

**Maarten de Rijke**

University of Amsterdam

derijke@uva.nl

# A bit of background

- University professor of AI & IR at University of Amsterdam
  - Lead a team of around 50 researchers on technology to connect people to information, as well as the implications
- Co-founder and director of national Innovation Center for AI
  - Built around labs in which universities collaborate with companies or NGO's around a shared research agenda

Innovation Center for Artificial Intelligence

# Today’s talk

# Joint work in progress with

Maartje ter Hoeve

Pengjie Ren

Maarten de Rijke

Svitlana Vakulenko

Nikos Voskarides

Yangjun Zhang

# Information retrieval

- **Technology to connect people to information**
  - Search engines
  - Recommender systems
  - Conversational assistants

# Landscape is changing

- **More mobile queries**
  - At the start of 2019, over 60% of all queries submitted to a major web search engine were mobile
- **Spoken queries**
  - Exceeding 50% in some parts of the world
  - Spoken queries longer, sessions longer

Source: https://www.statista.com/statistics/275814/mobile-share-of-organic-search-engine-visits/

# Conversational search

- Idea of **search as conversation** has been around since early 1980s (Belkin, CJIS 1980)
- Making information retrieval interfaces feel more natural and convenient for their users (Radlinski & Craswell, CHIIR 2017)
- Ongoing research and development efforts heavily skewed towards **task-oriented dialogue systems** and **(factoid) question answering tasks**

There’s more than factoids …

# Talk with a document

- Conversational agents to increase productivity
- Document-centered assistance — e.g., to help an individual quickly review a document
  - What type of functionality?
  - What type of questions?
  - Recognizing and answering document-centered questions

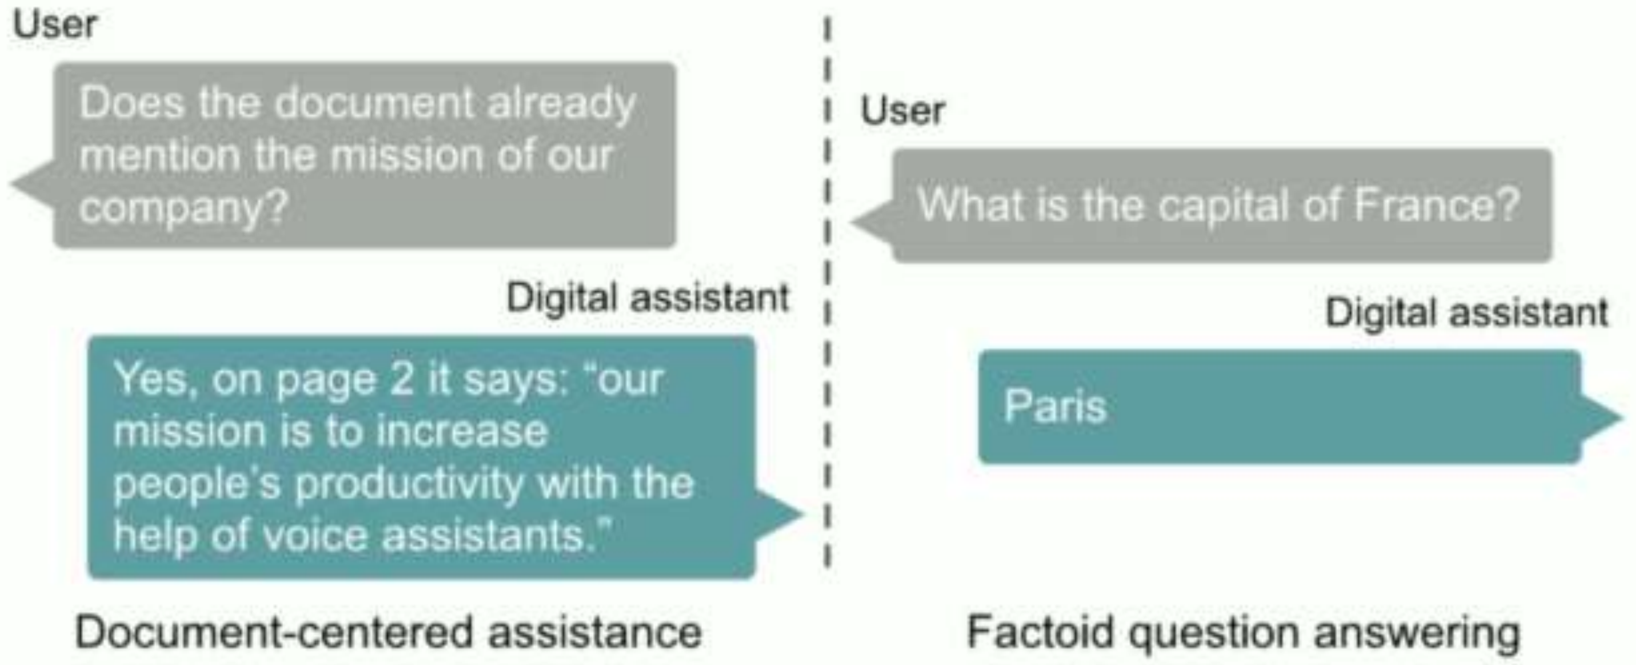


M. ter Hoeve et al., 2019. Conversations with Documents. Under review.

# Talk with structured information

- **Talk with a dataset**
- Using conversations to **browse** large collections of information objects
- User model maintains knowledge state, information goal, navigation strategy

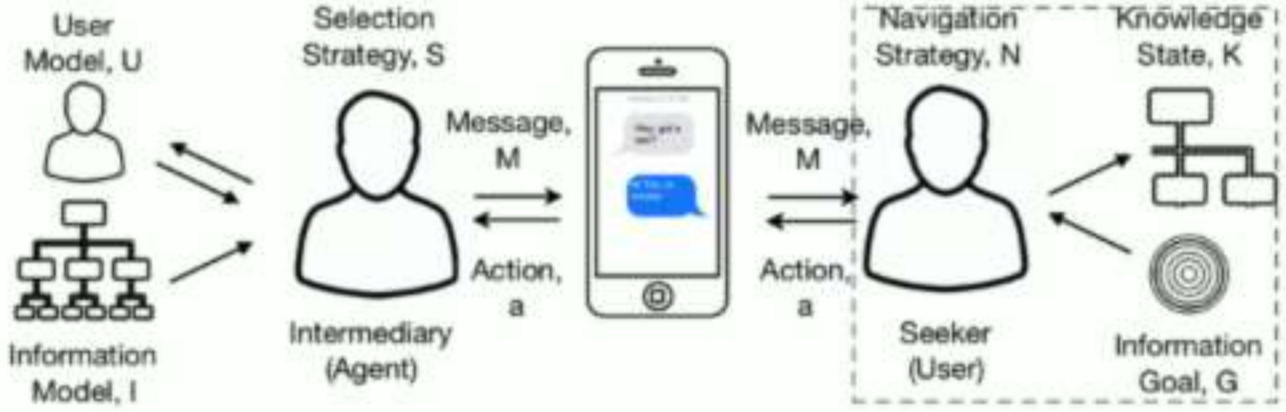




S. Vakulenko et al., 2018. Measuring semantic coherence of a conversation. *ISWC*

# Talk with a collection of items

- **Talk to support exploration**
  - Educational, entertainment purposes
  - Serendipitous discoveries of cultural artifacts – users often look for inspiration, surprises, novel ideas
  - E-commerce

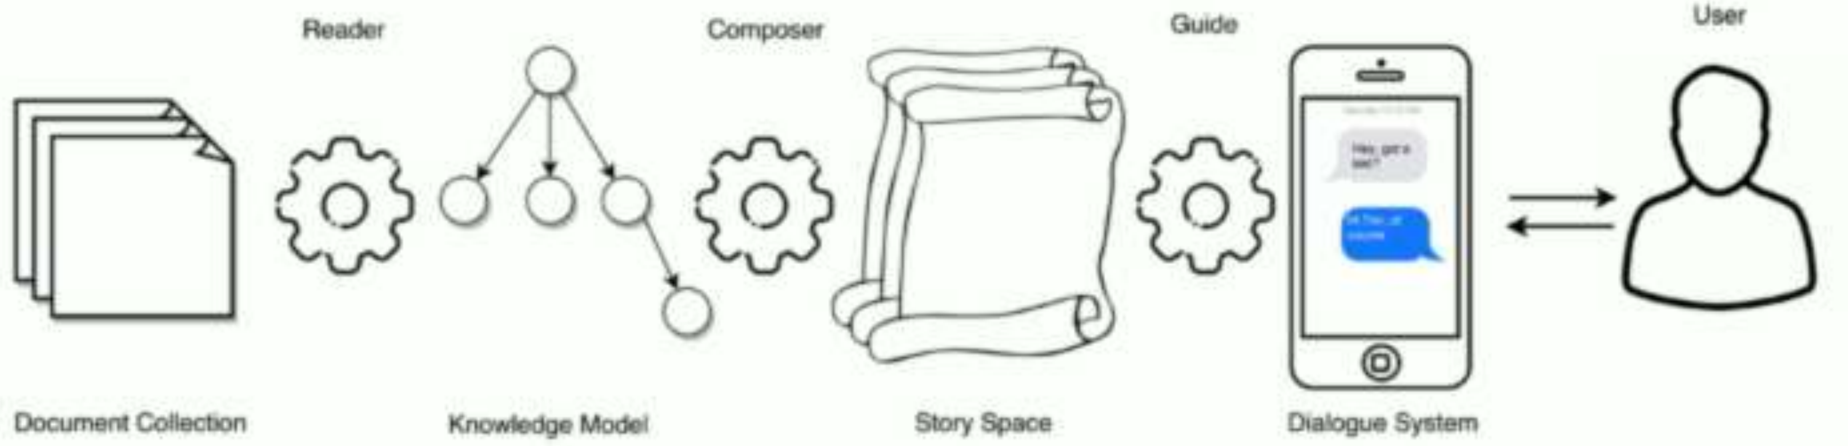


S. Vakulenko et al., 2017. Conversational Exploratory Search via Interactive Storytelling. SCAI

# Informational goals

- There is more than task oriented dialogue systems and question answering

- **Navigational**, **informational**, and **resource** goals

    - Informational goals consistently 40–60% of all goals

D.E. Rose and D. Levinson. 2004. Understanding User Goals in Web Search. In *WWW*

# Addressing informational goals

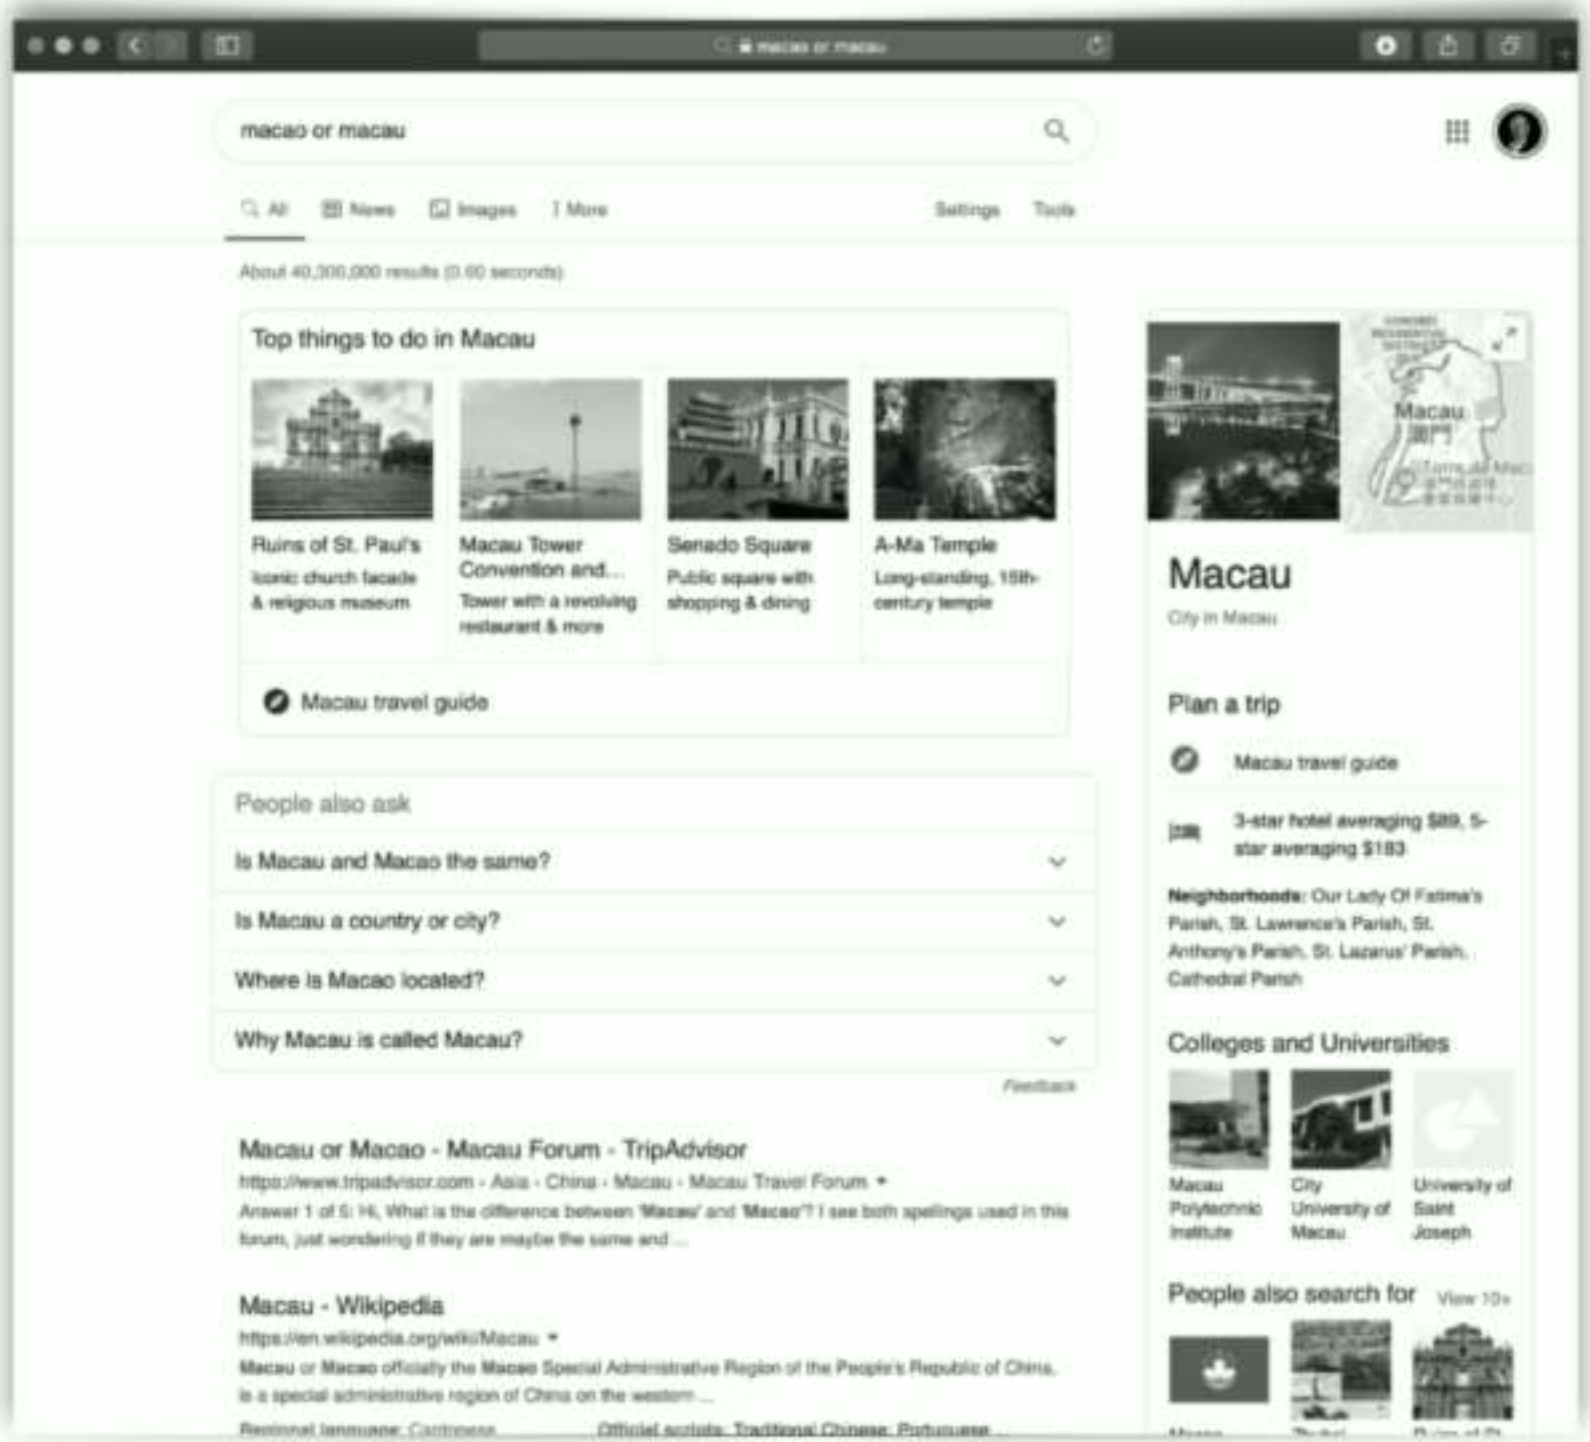

# SERP-based conversations

- As our mode of interaction changes, how can we support conversations based on **search engine result pages** with all the diversity and uncertainty there is in SERPs?

## Where are we ...?

**Mostly work in progress**

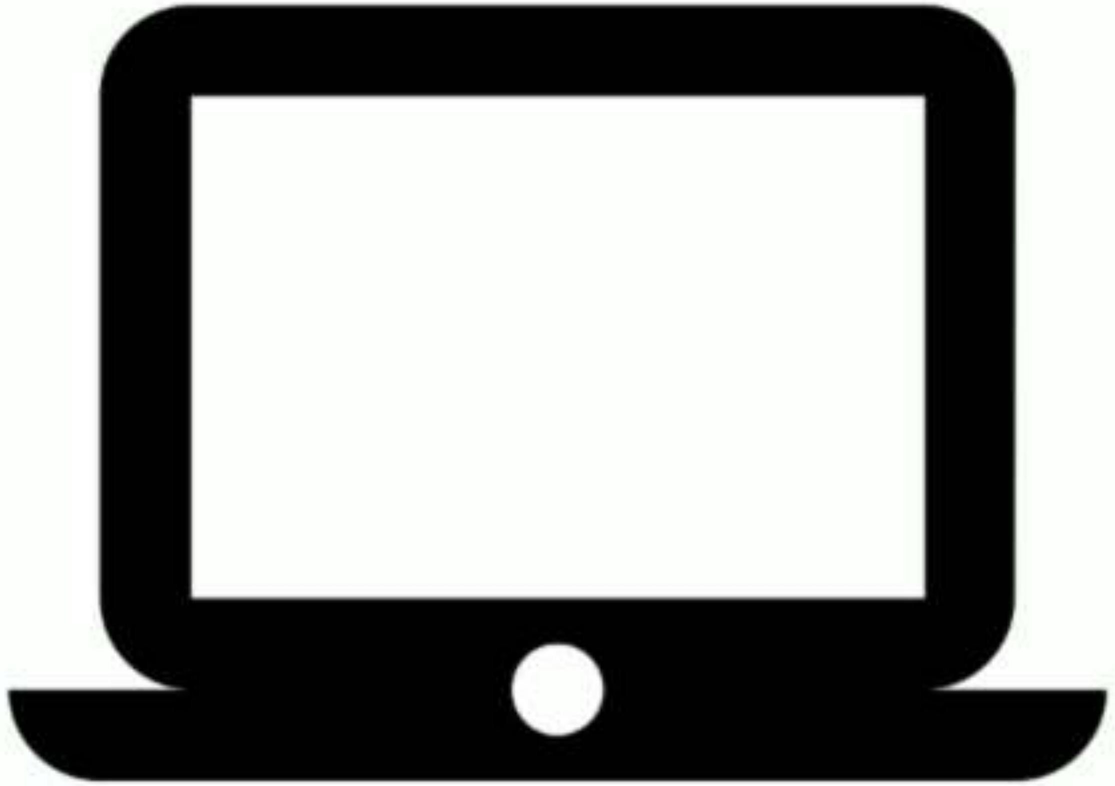


money made by <movie>

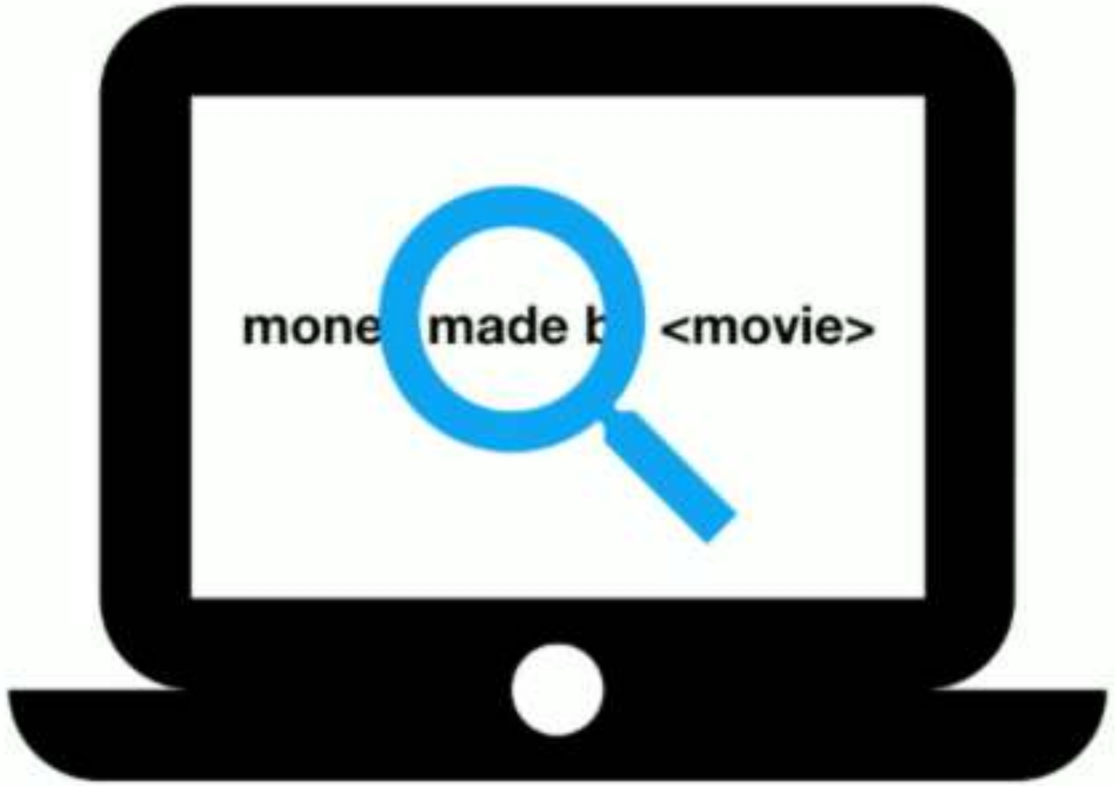
mone made b <movie>

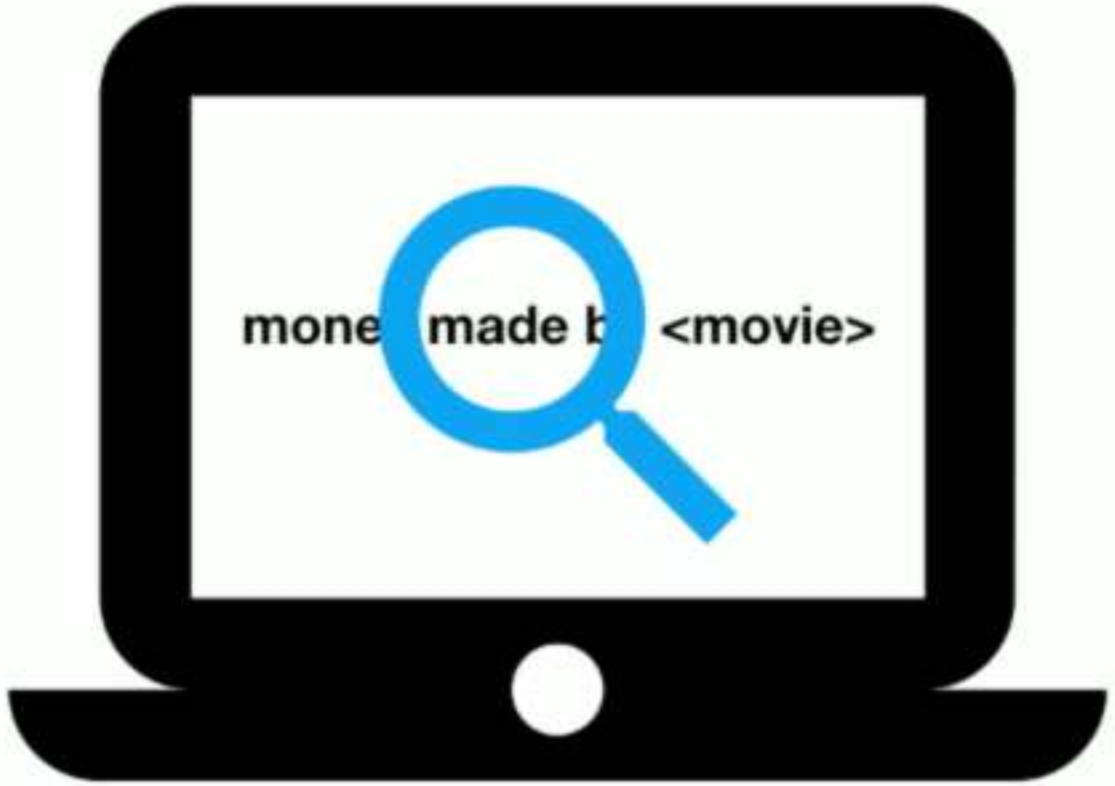


... that 's redundant because it is a ben stiller movie . clearly my expectations for this movie were not high and , maybe because of that , i found " meet the fockers " quite funny . do you remember when big box office $ 279,167,575 awards ascap film and television music awards 2005 top box office ...

what did you think of the title ?

what did you think of the title ?

the title pretty much describes the level of the humor in this ben stiller movie .

what did you think of the title ?

the title pretty much describes the level of the humor in this ben stiller movie .

haha , i agree ! do you know if it made any money ?

what did you think of the title ?

the title pretty much describes the level of the humor in this ben stiller movie .

haha , i agree ! do you know if it made any money ?

yeah , it made $ 279,167,575 . pretty good .

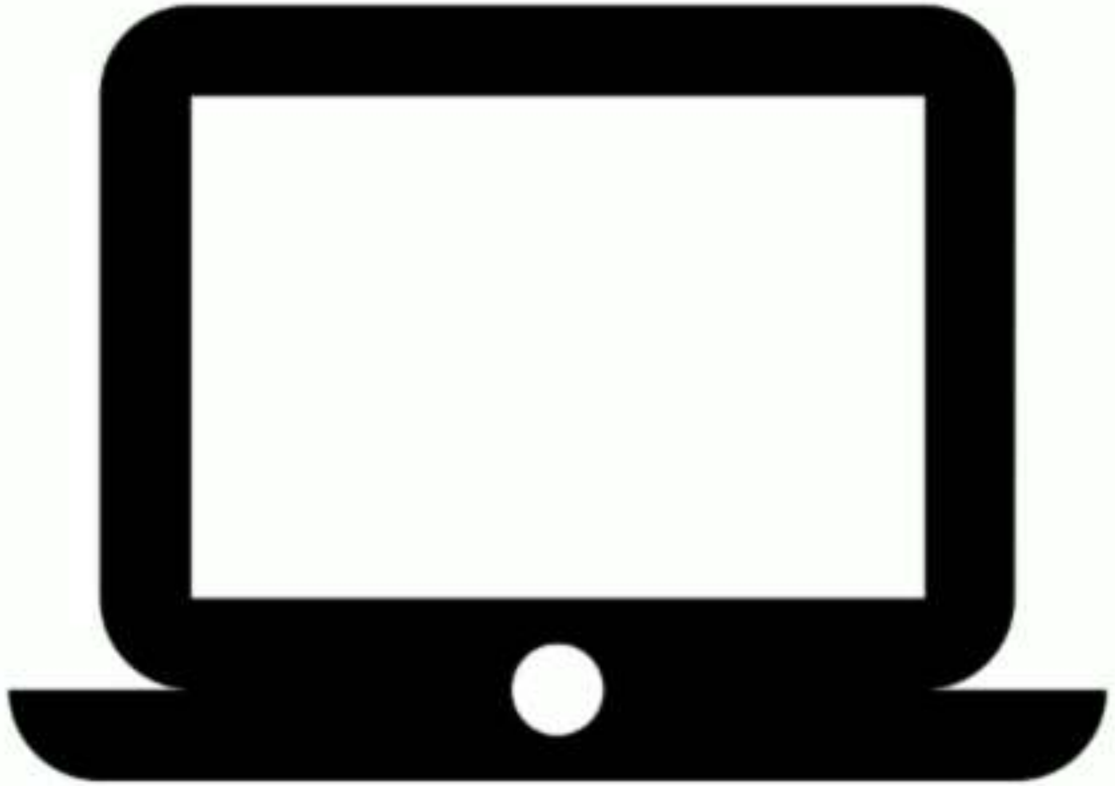

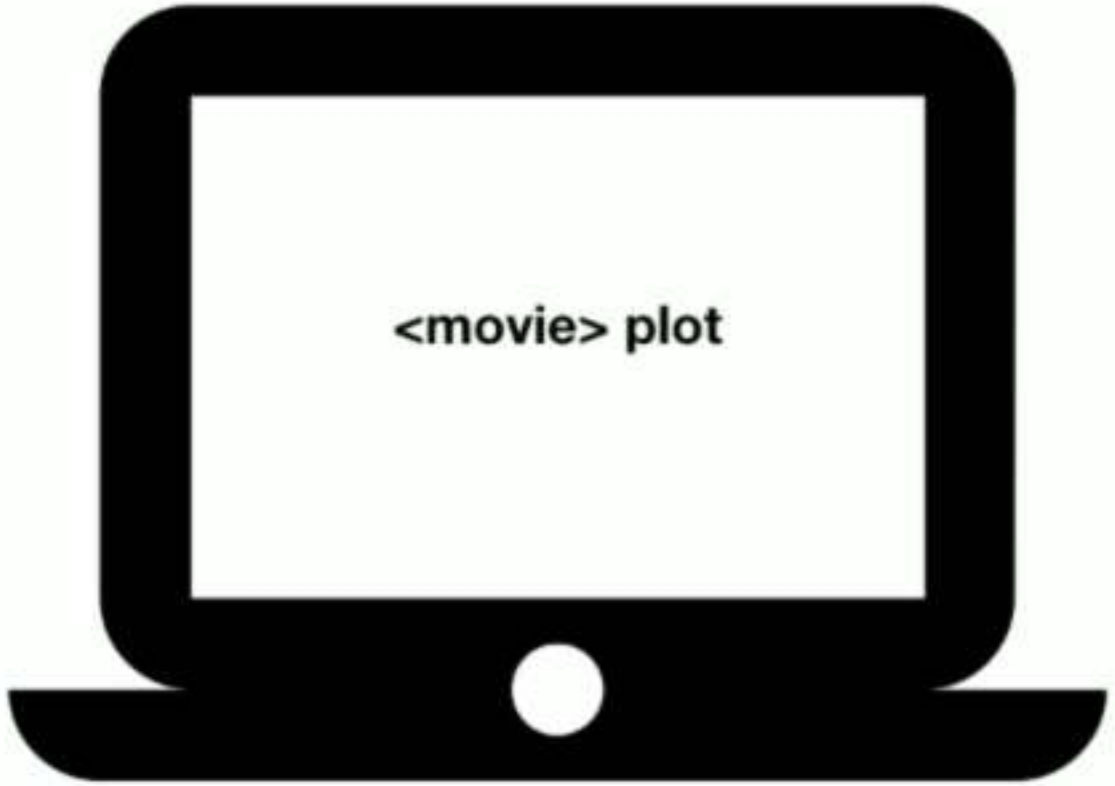
<movie> plot

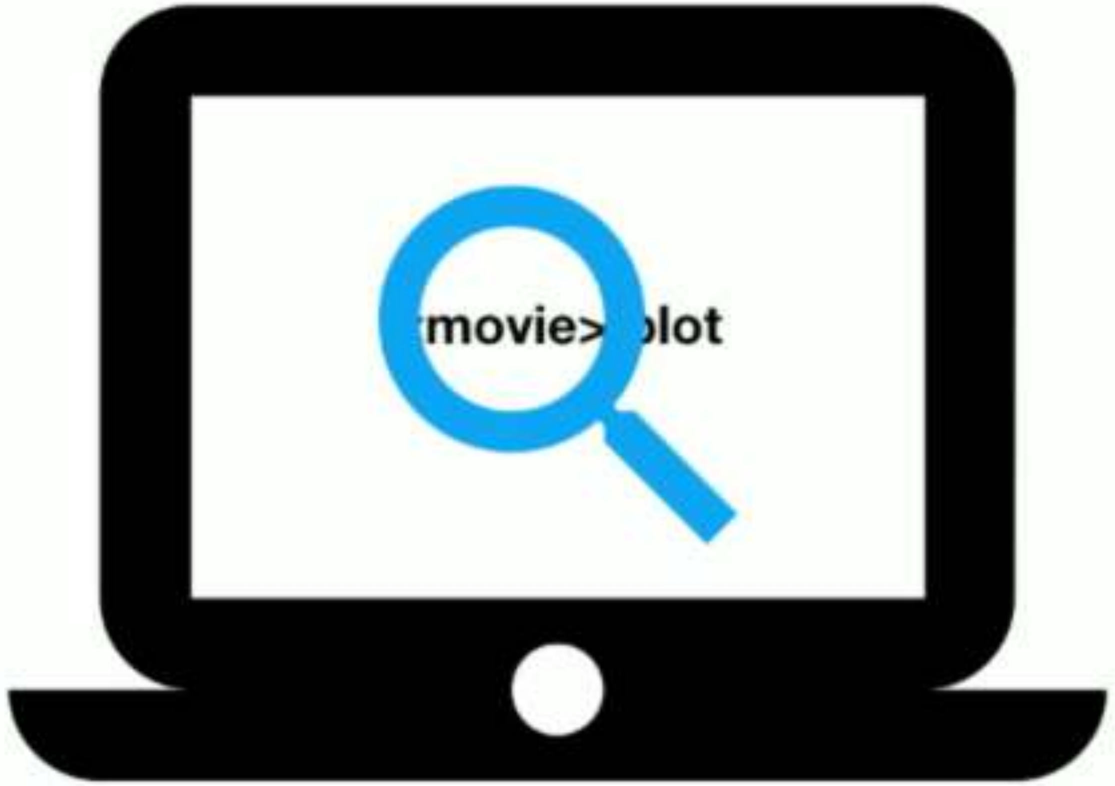
<movie> plot

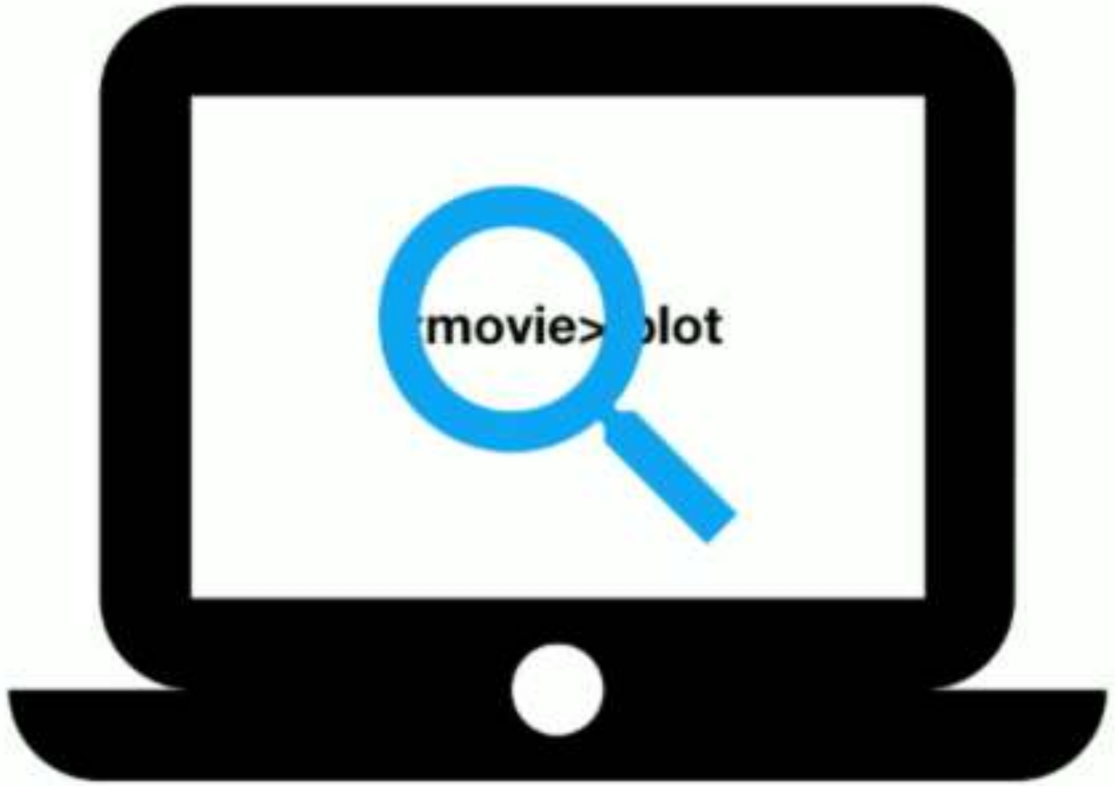


...being captured by boris and onatopp . bond arrives in st . petersburg and meets his cia contact , jack wade ( joe don baker ) . wade agrees to take bond to the hideout of a russian gangster , valentin zukovsky ( robbie coltrane ) , whom bond had shot in the leg and given a permanent limp years before ...

that was a good scene

that was a good scene

what did you like about the movie ?

that was a good scene

what did you like about the movie ?

i liked his friend , jack wade .

that was a good scene

what did you like about the movie ?

i liked his friend , jack wade .

i loved the part where bond arrives in st . petersburg and meets his cia contact , jack wade ( joe don baker ) .

# SERP-based conversation

# SERP-based conversation

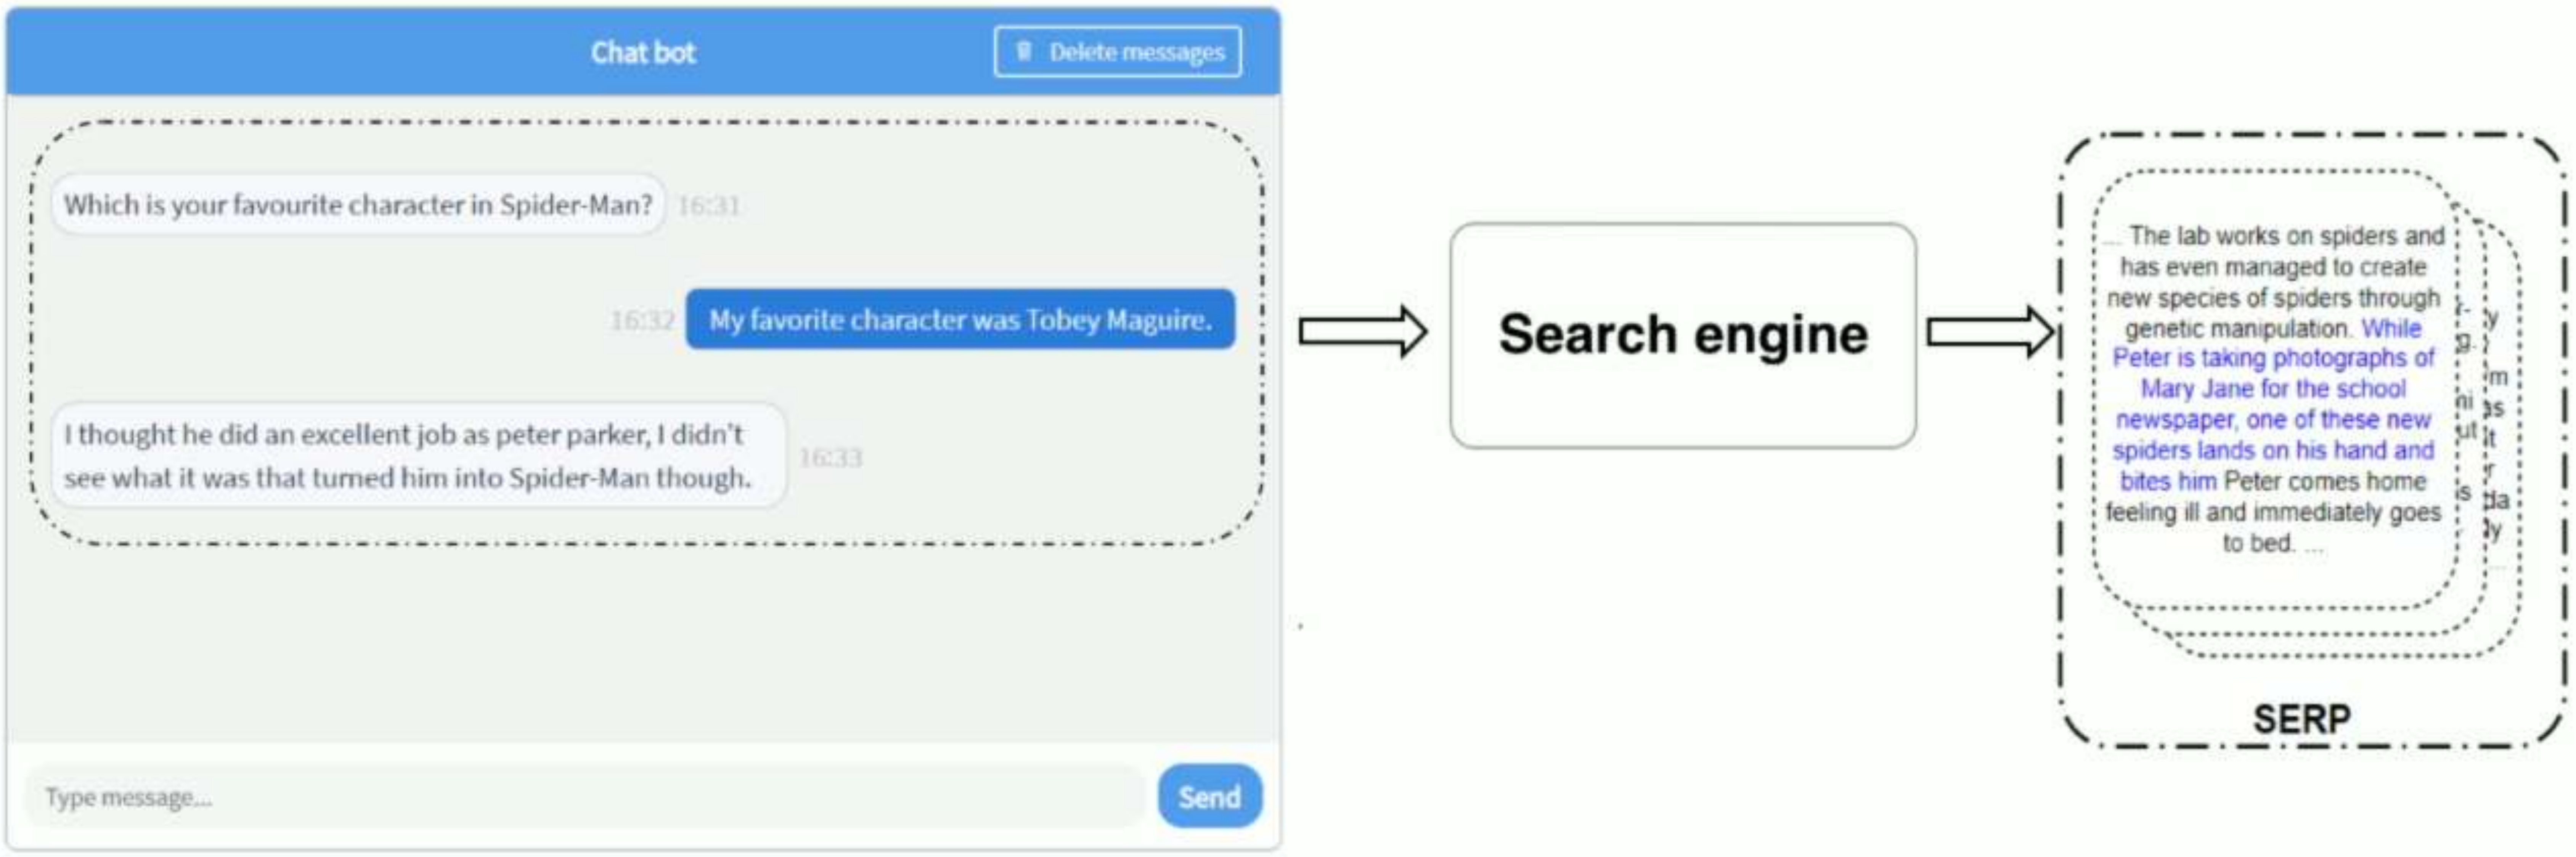

# SERP-based conversation

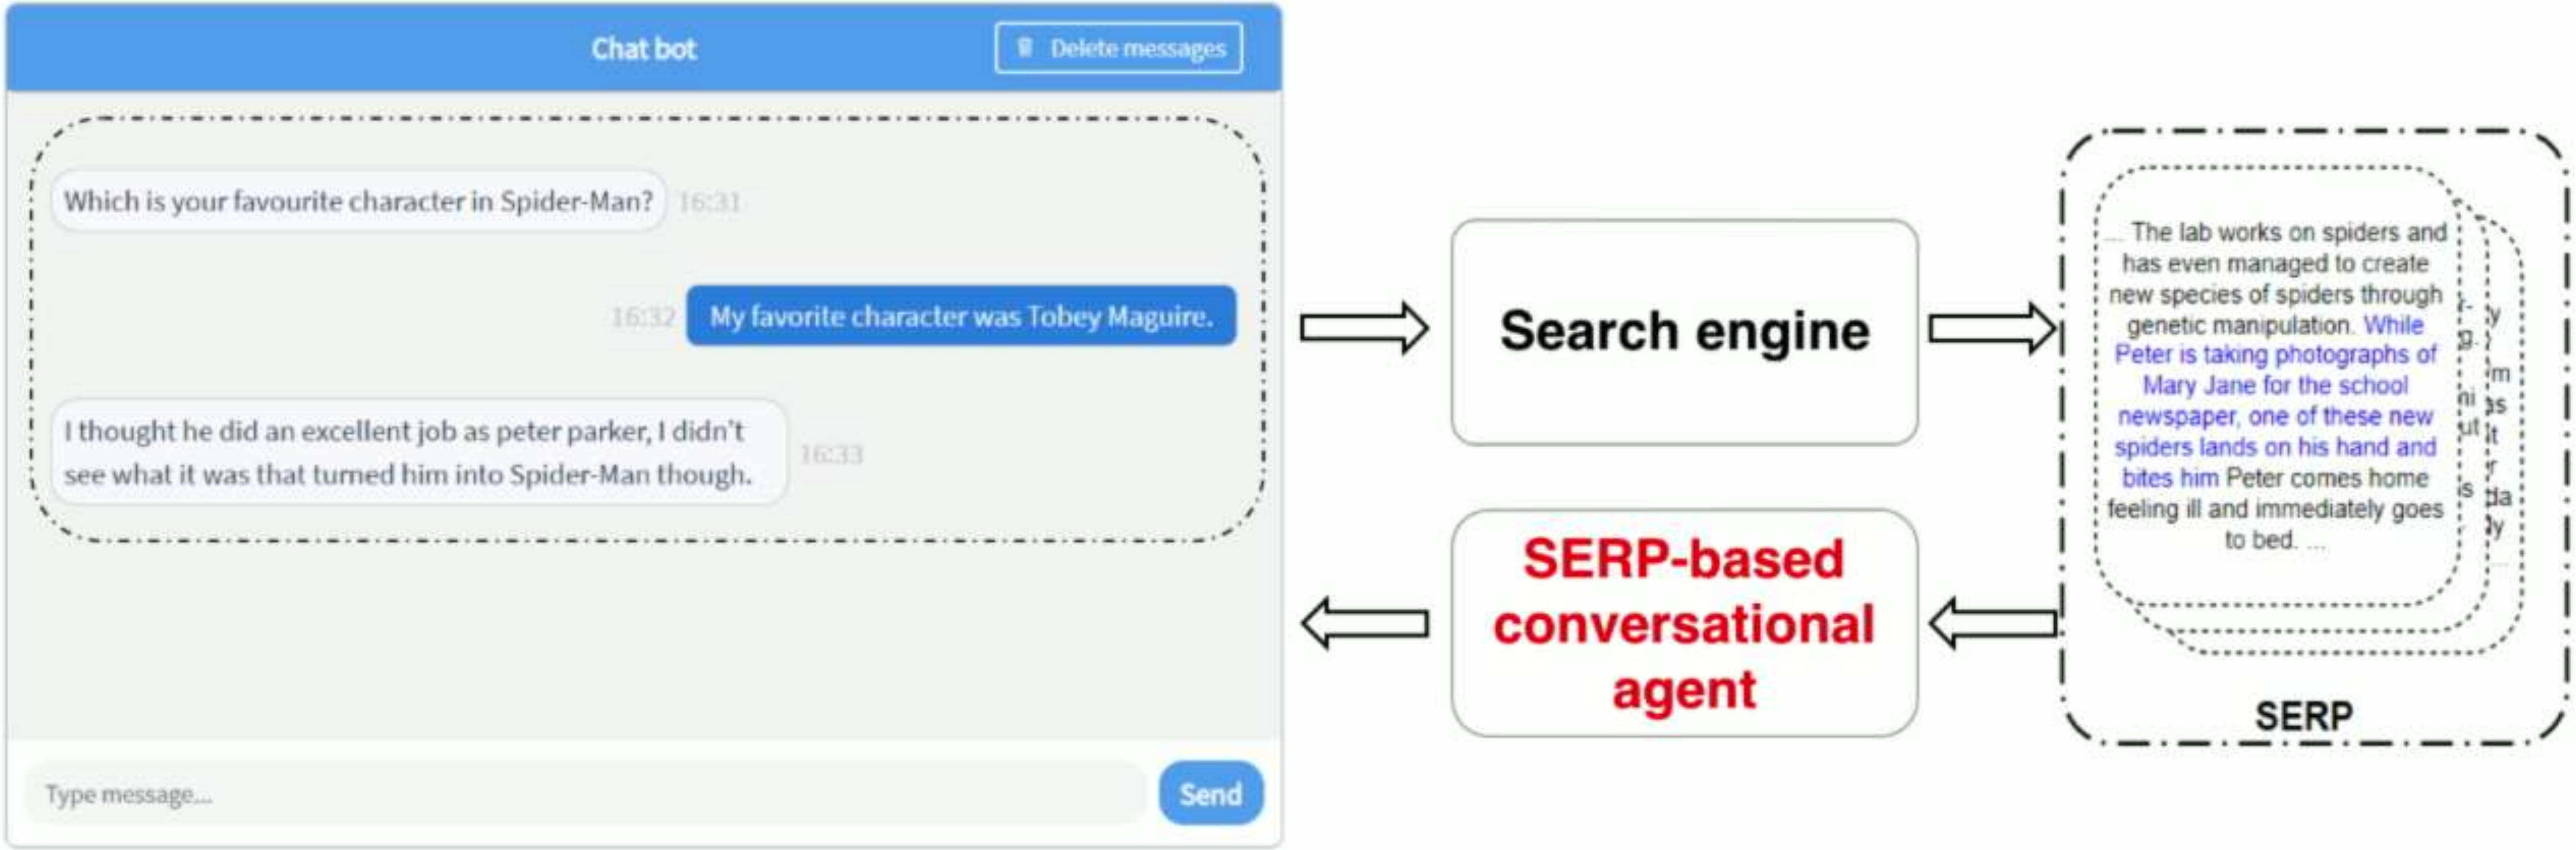

# SERP-based conversation

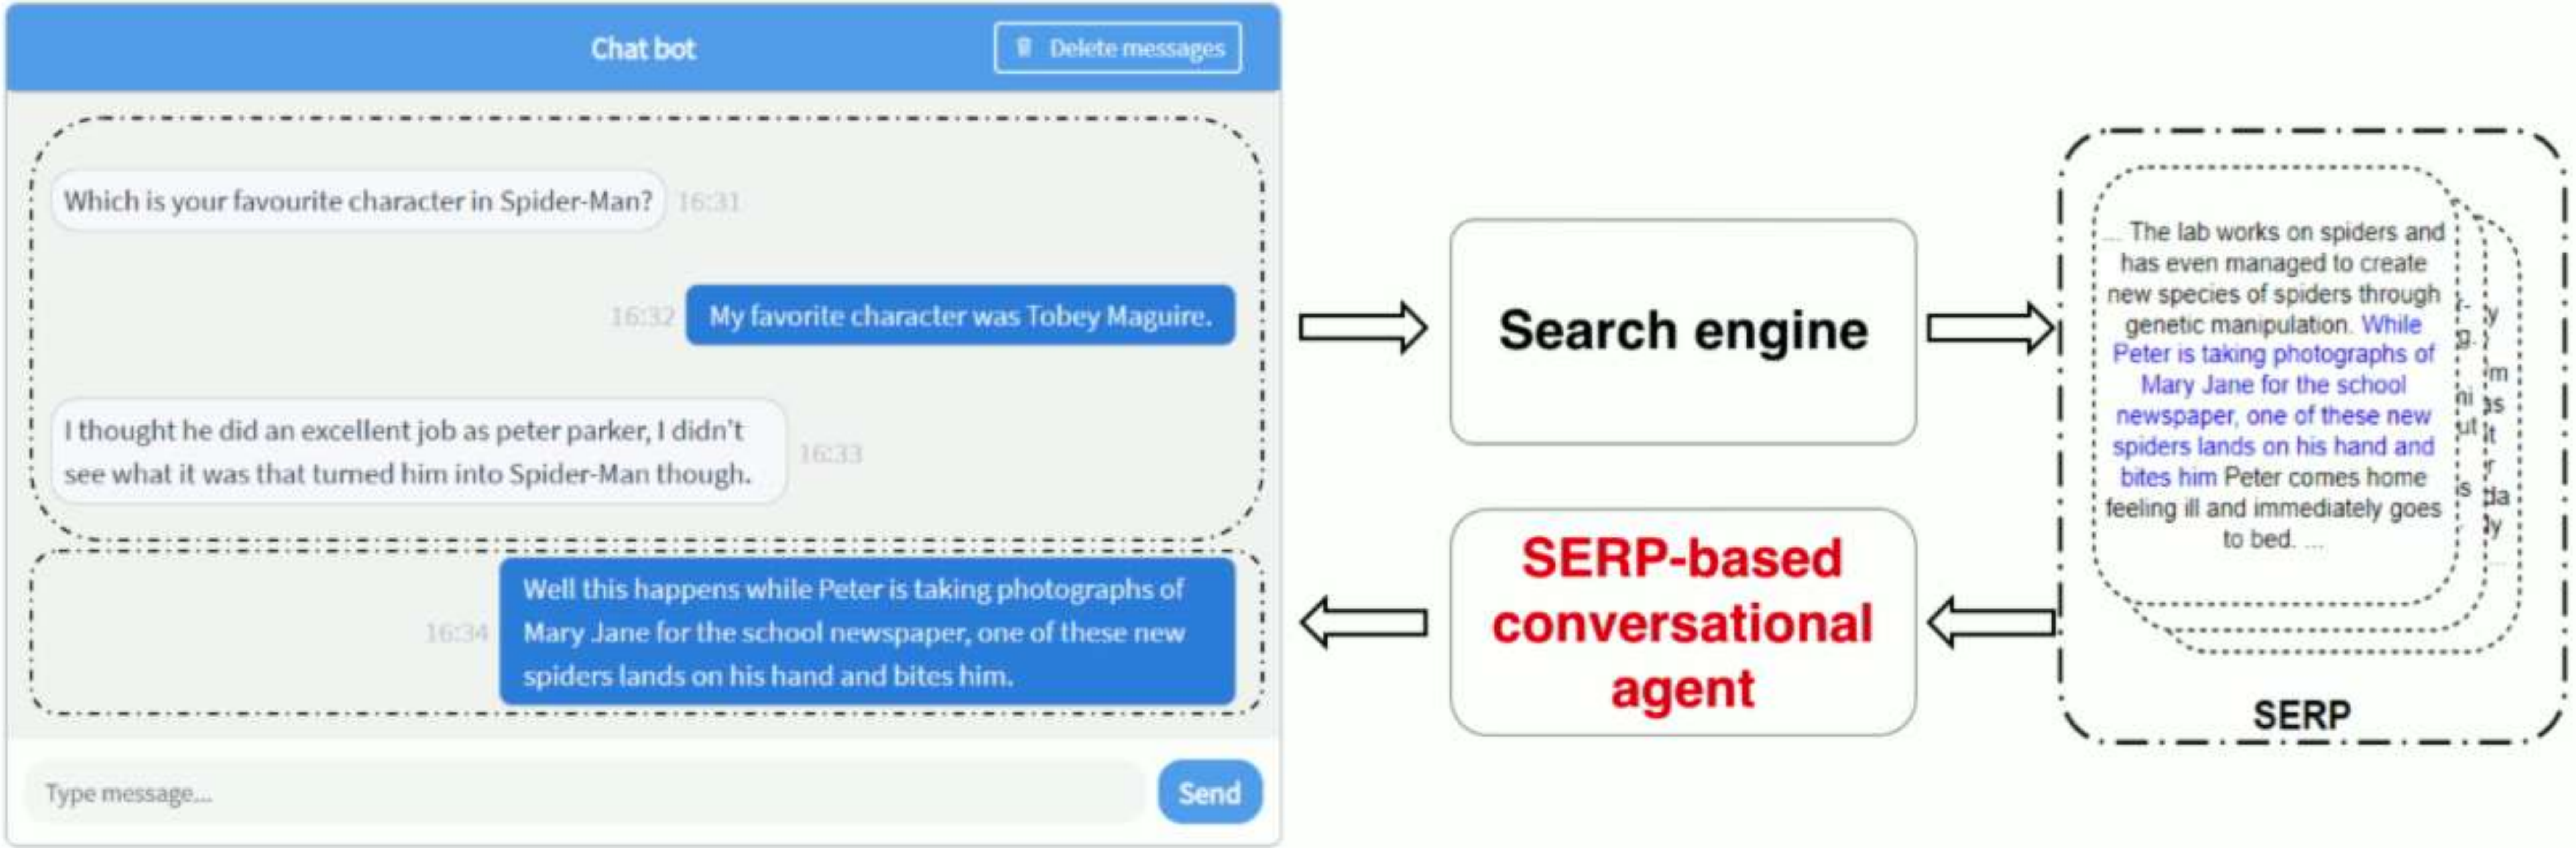

# SERP-based conversation

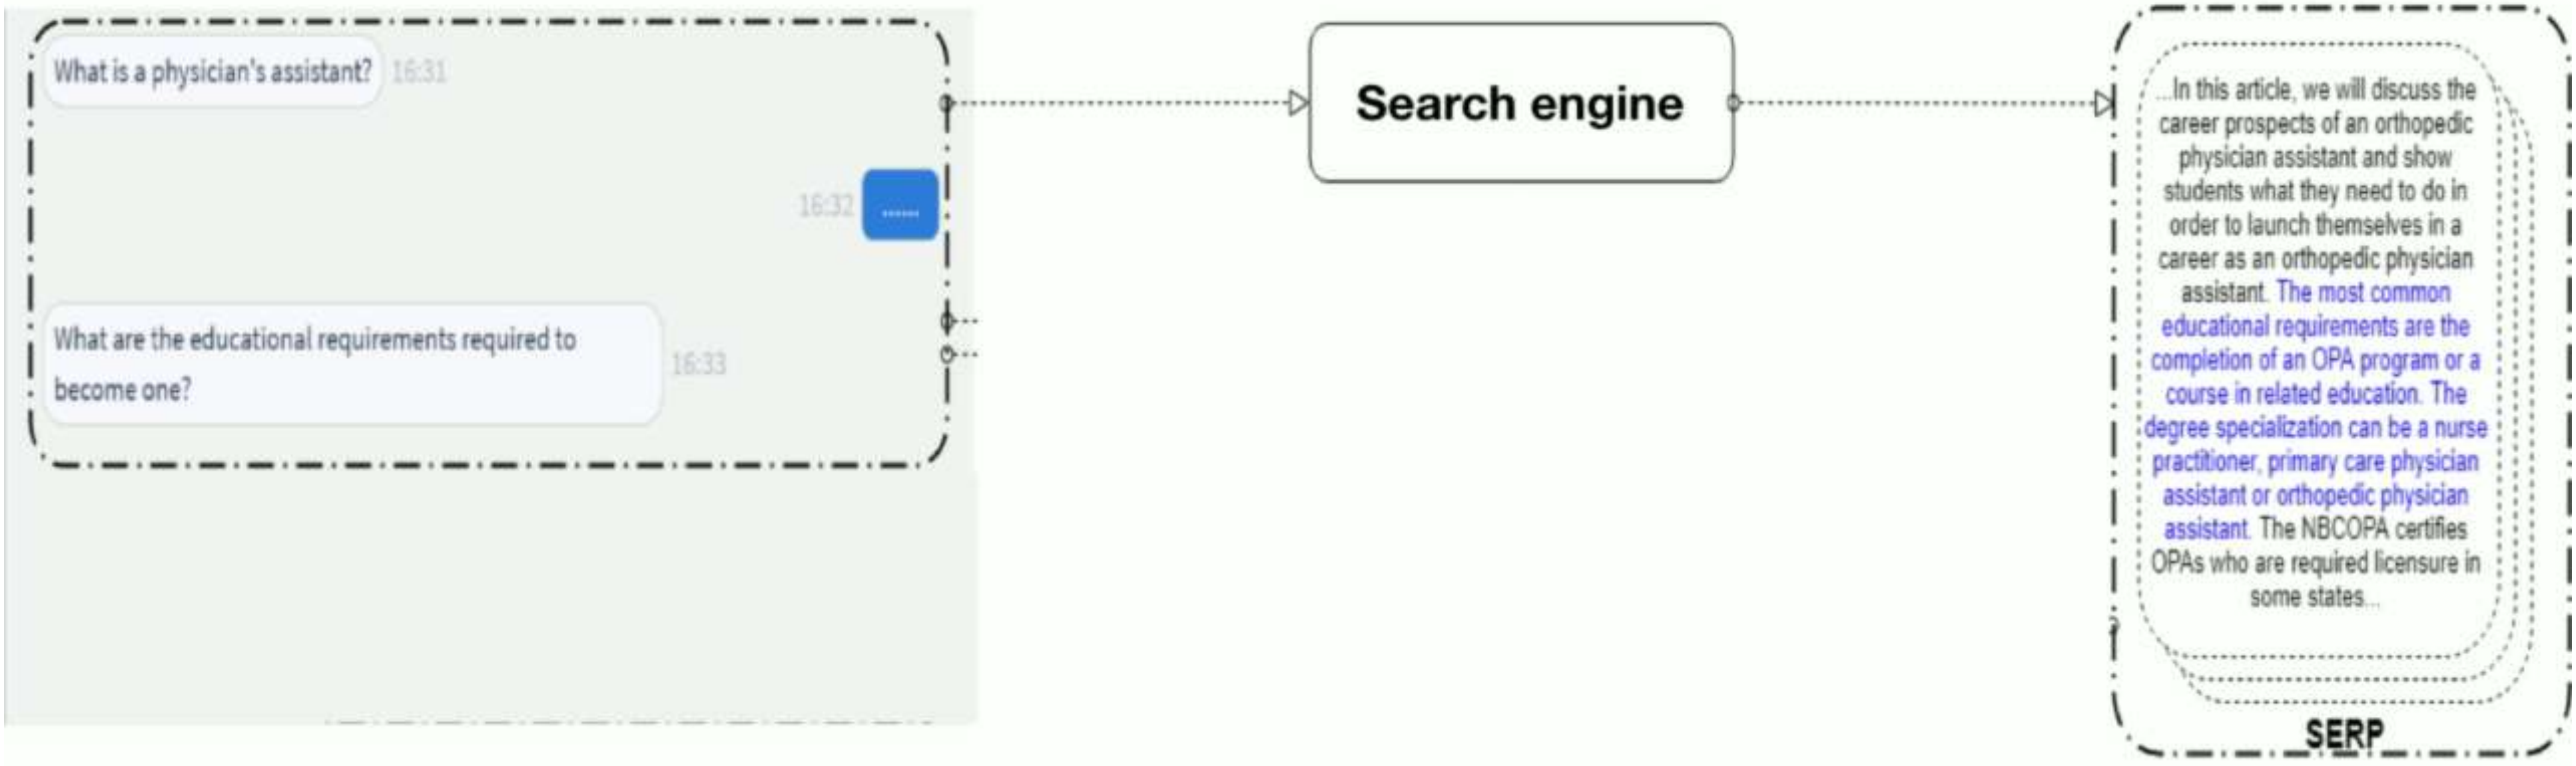

# SERP-based conversation

# SERP-based conversation

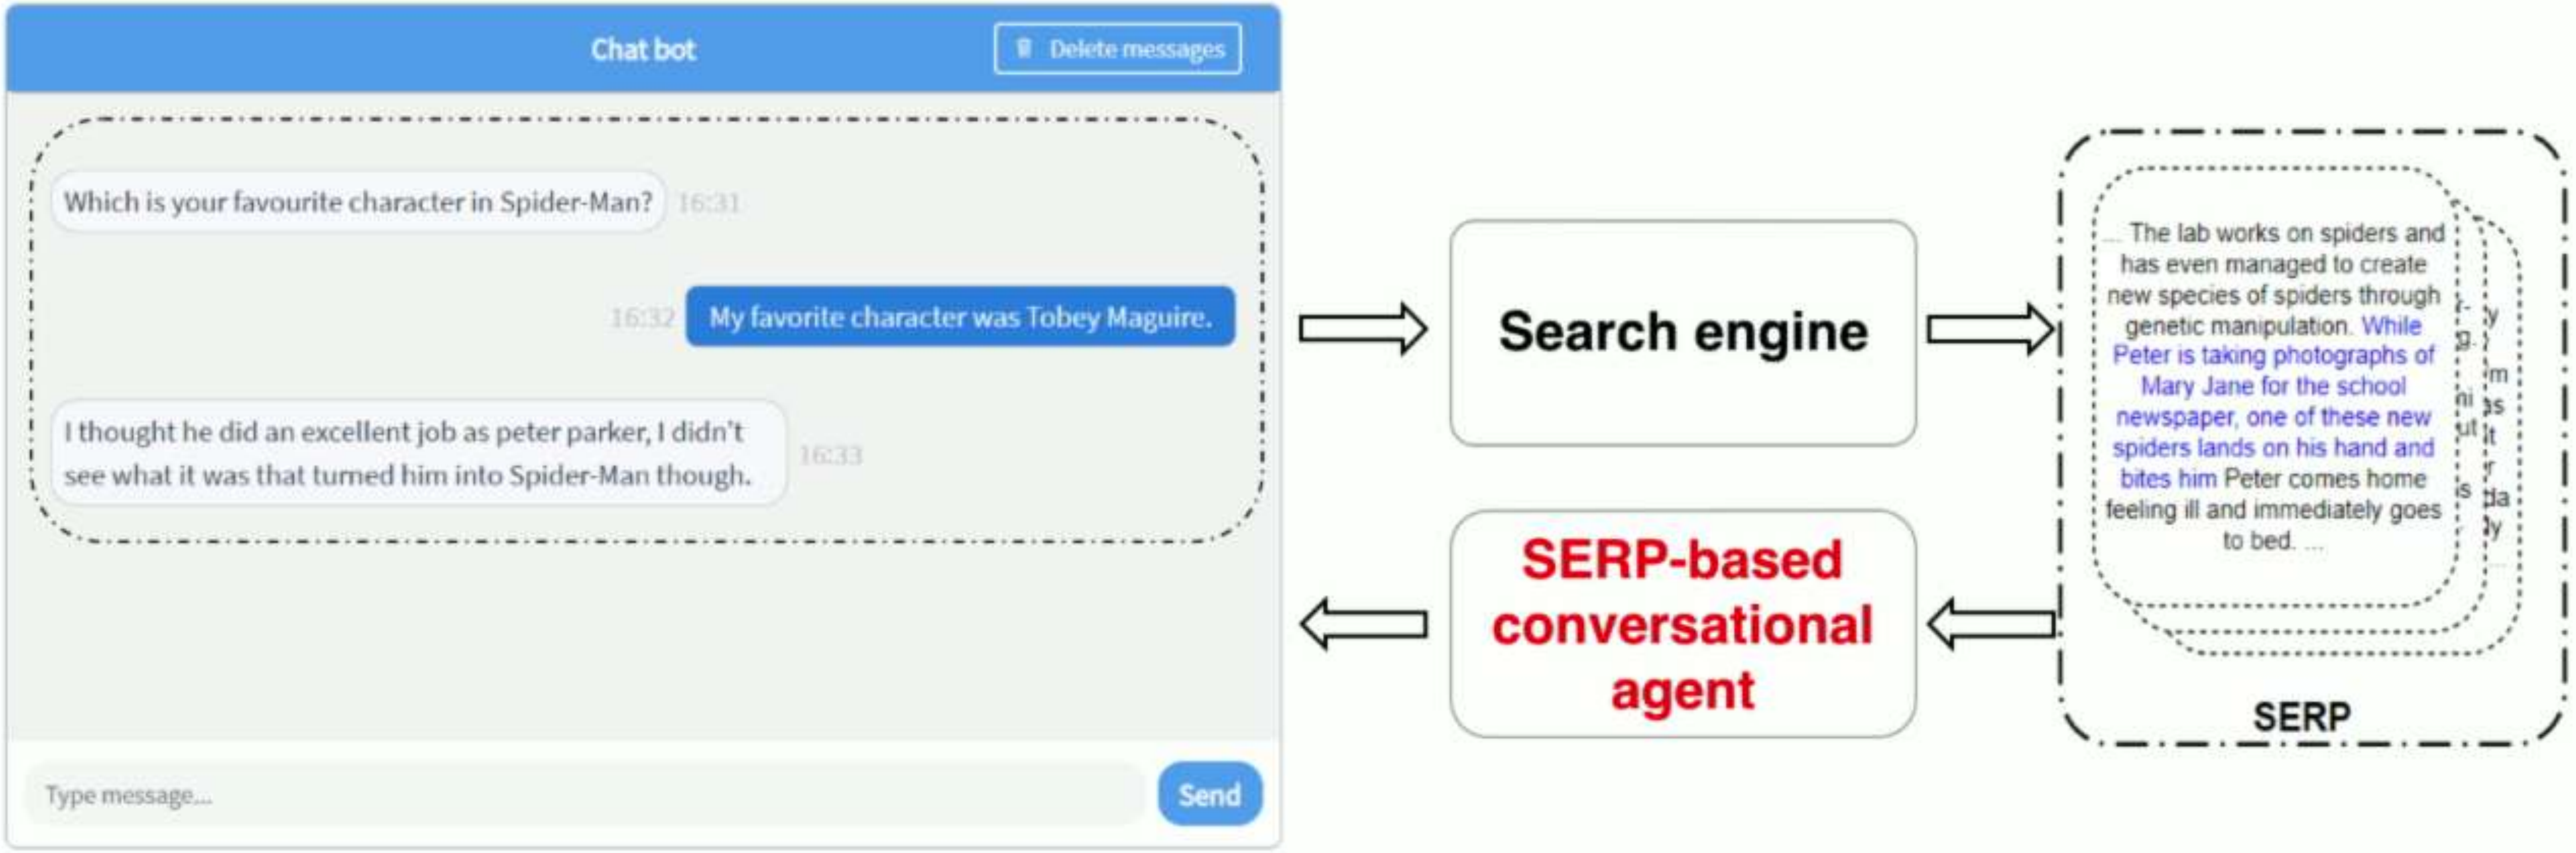

# SERP-based conversation

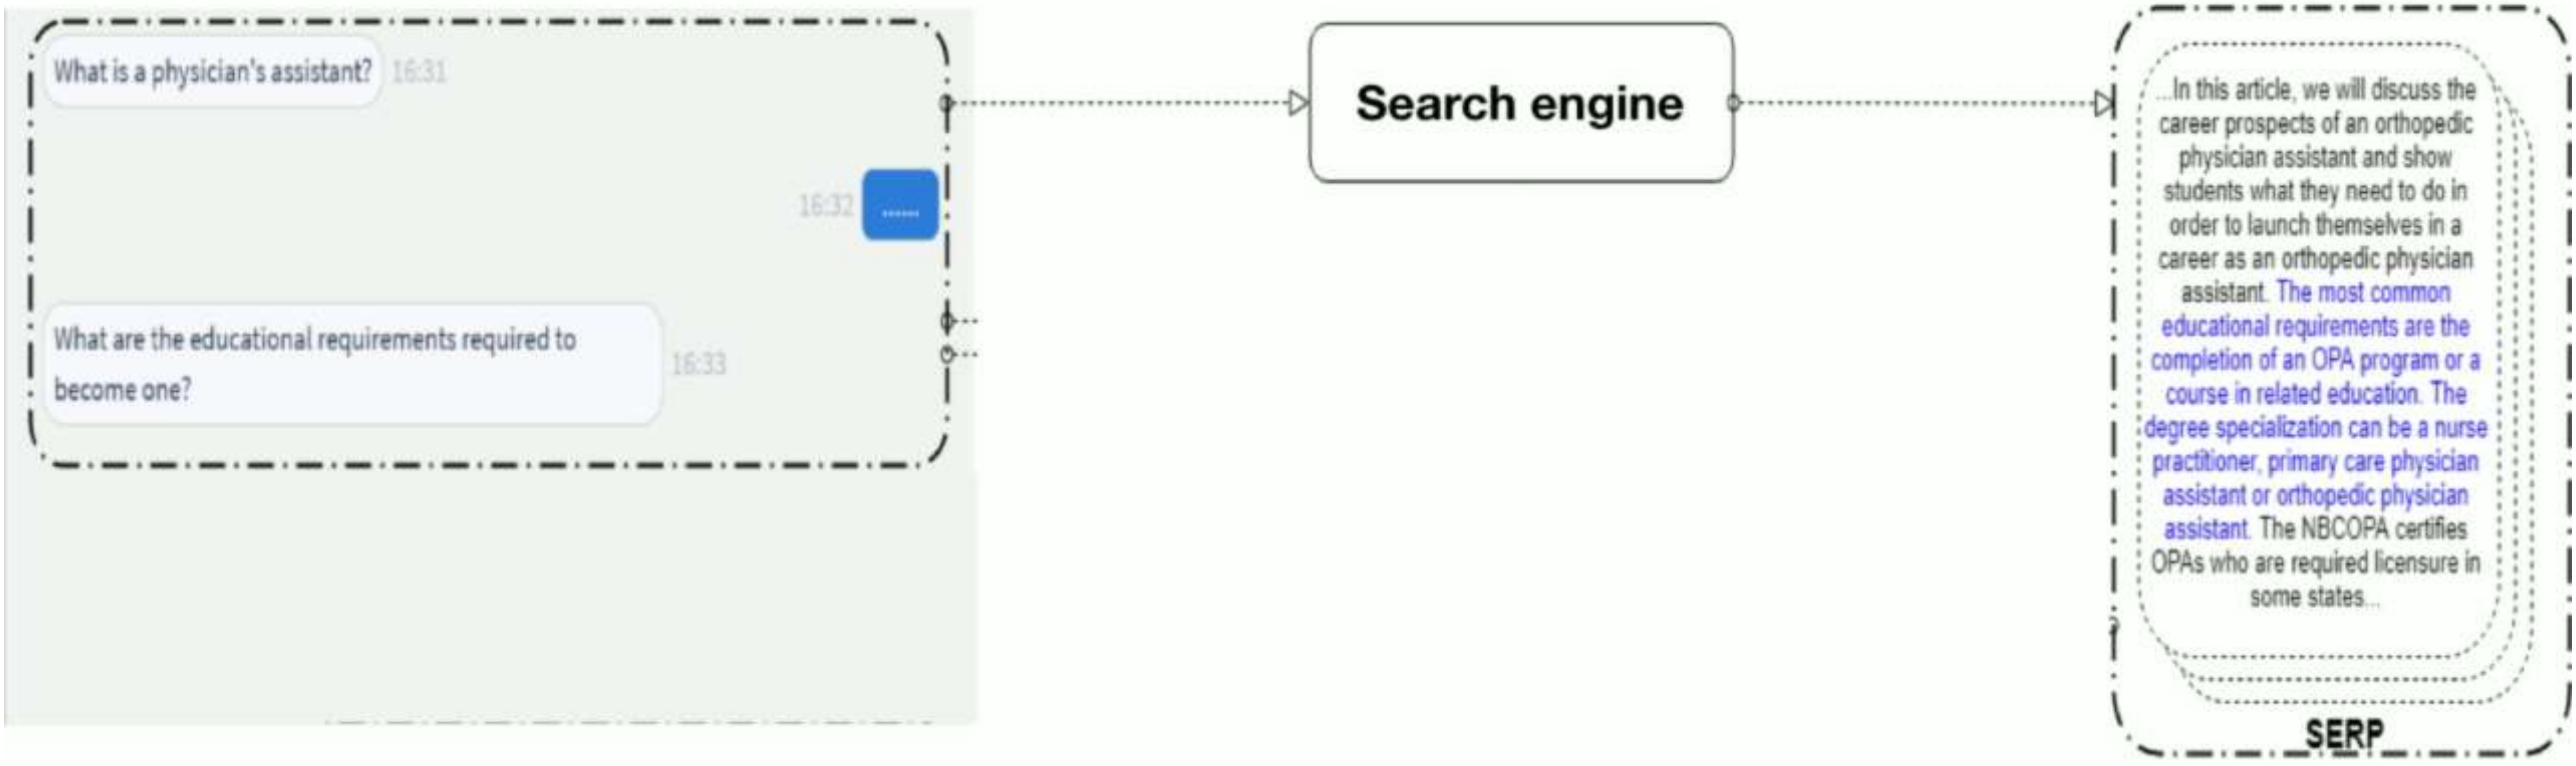

# SERP-based conversation

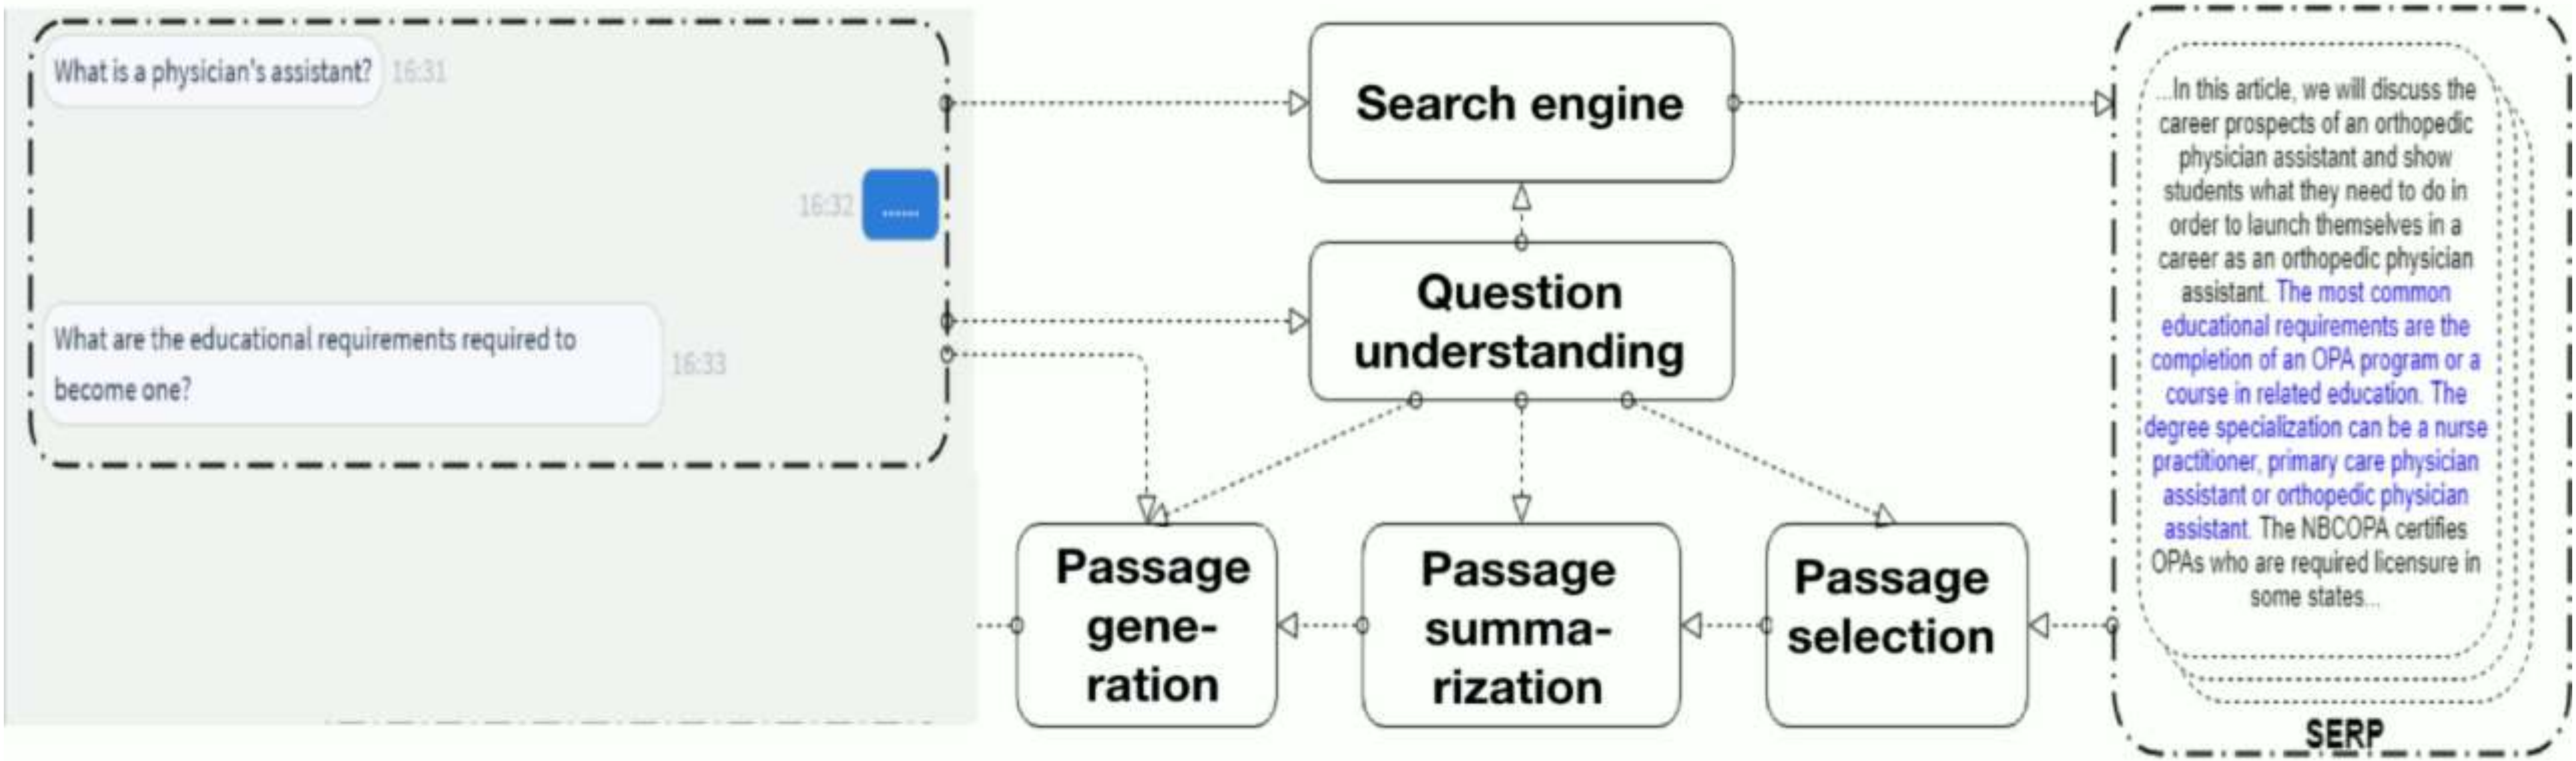

# SERP-based conversation

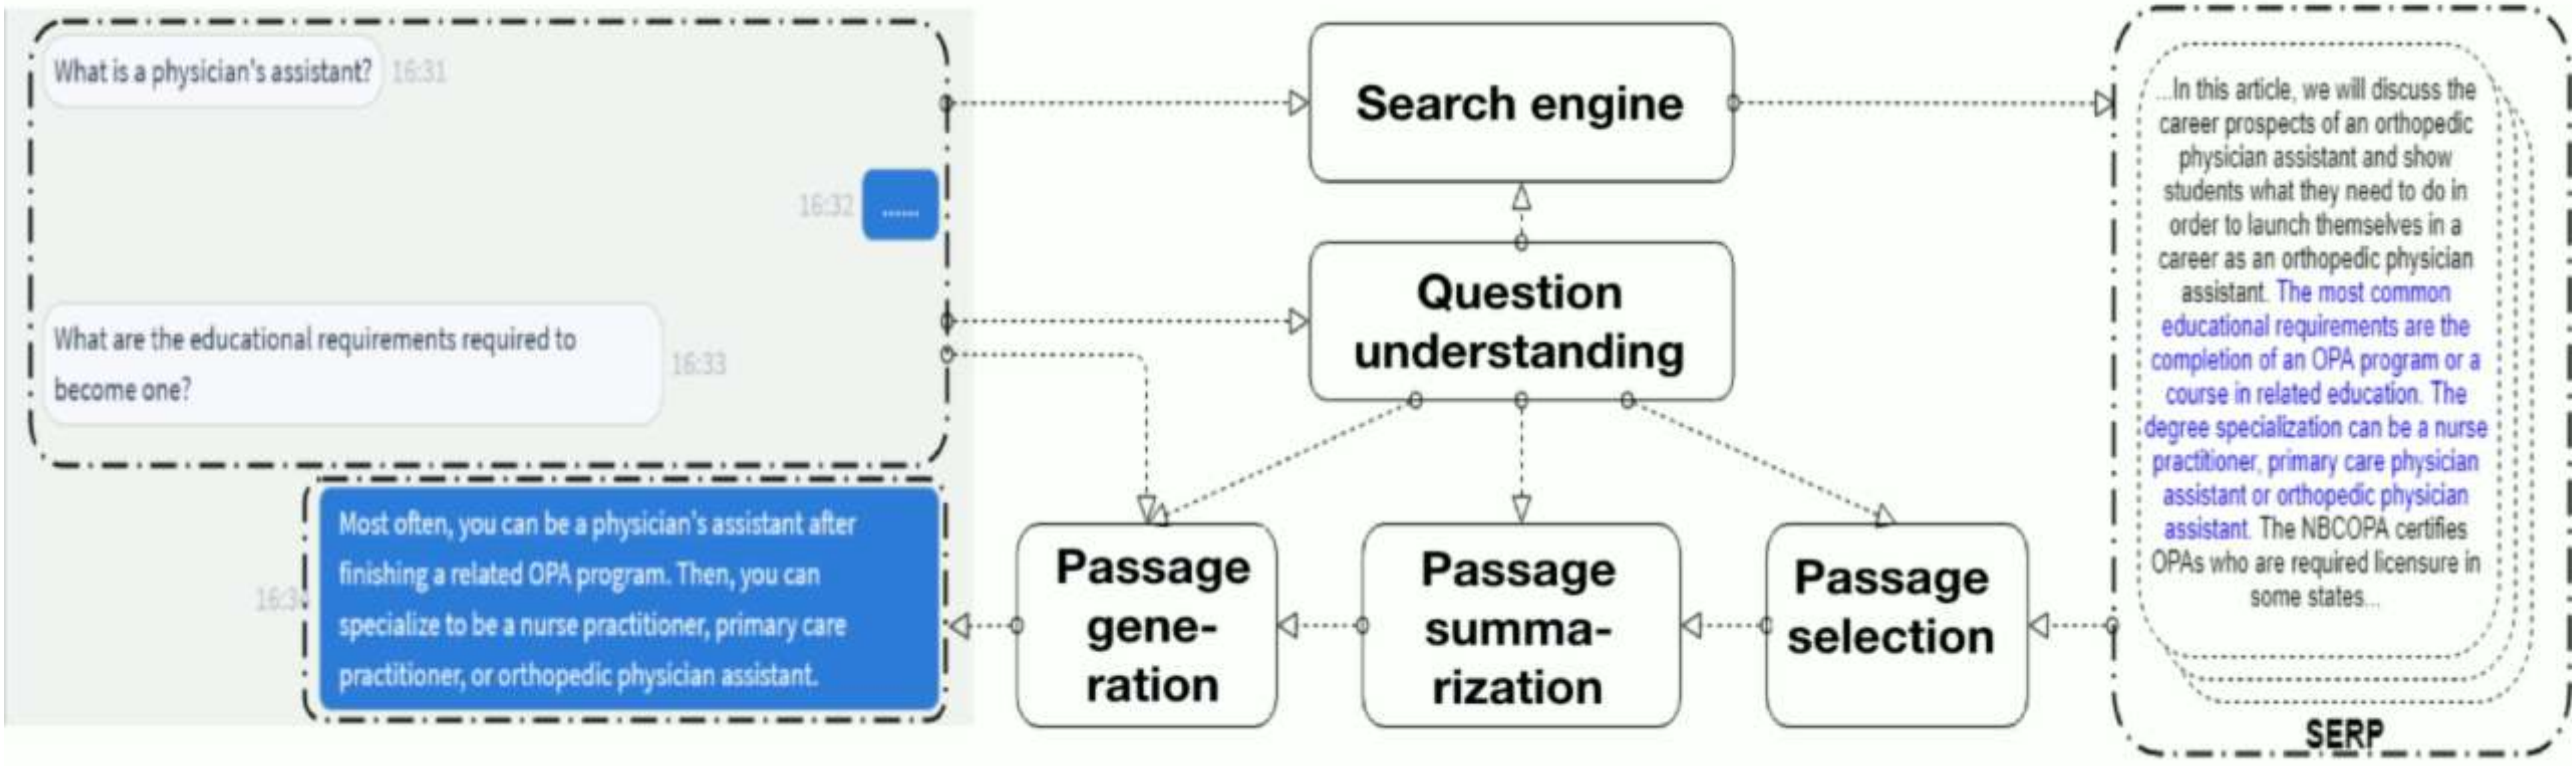

# Let's simplify things …

- **Cut out the retrieval step** and assume that an oracle has given a single document that may be used to inform the response

- **Background based conversation (BBC):** Given some background knowledge (e.g., an article in the form of free text) and a conversation, the BBC task is to generate responses by referring to the background knowledge and considering the dialogue history context at the same time

Zhou et al., 2018. A dataset for document grounded conversations. In *EMNLP 2018*.

# Background-based conversation

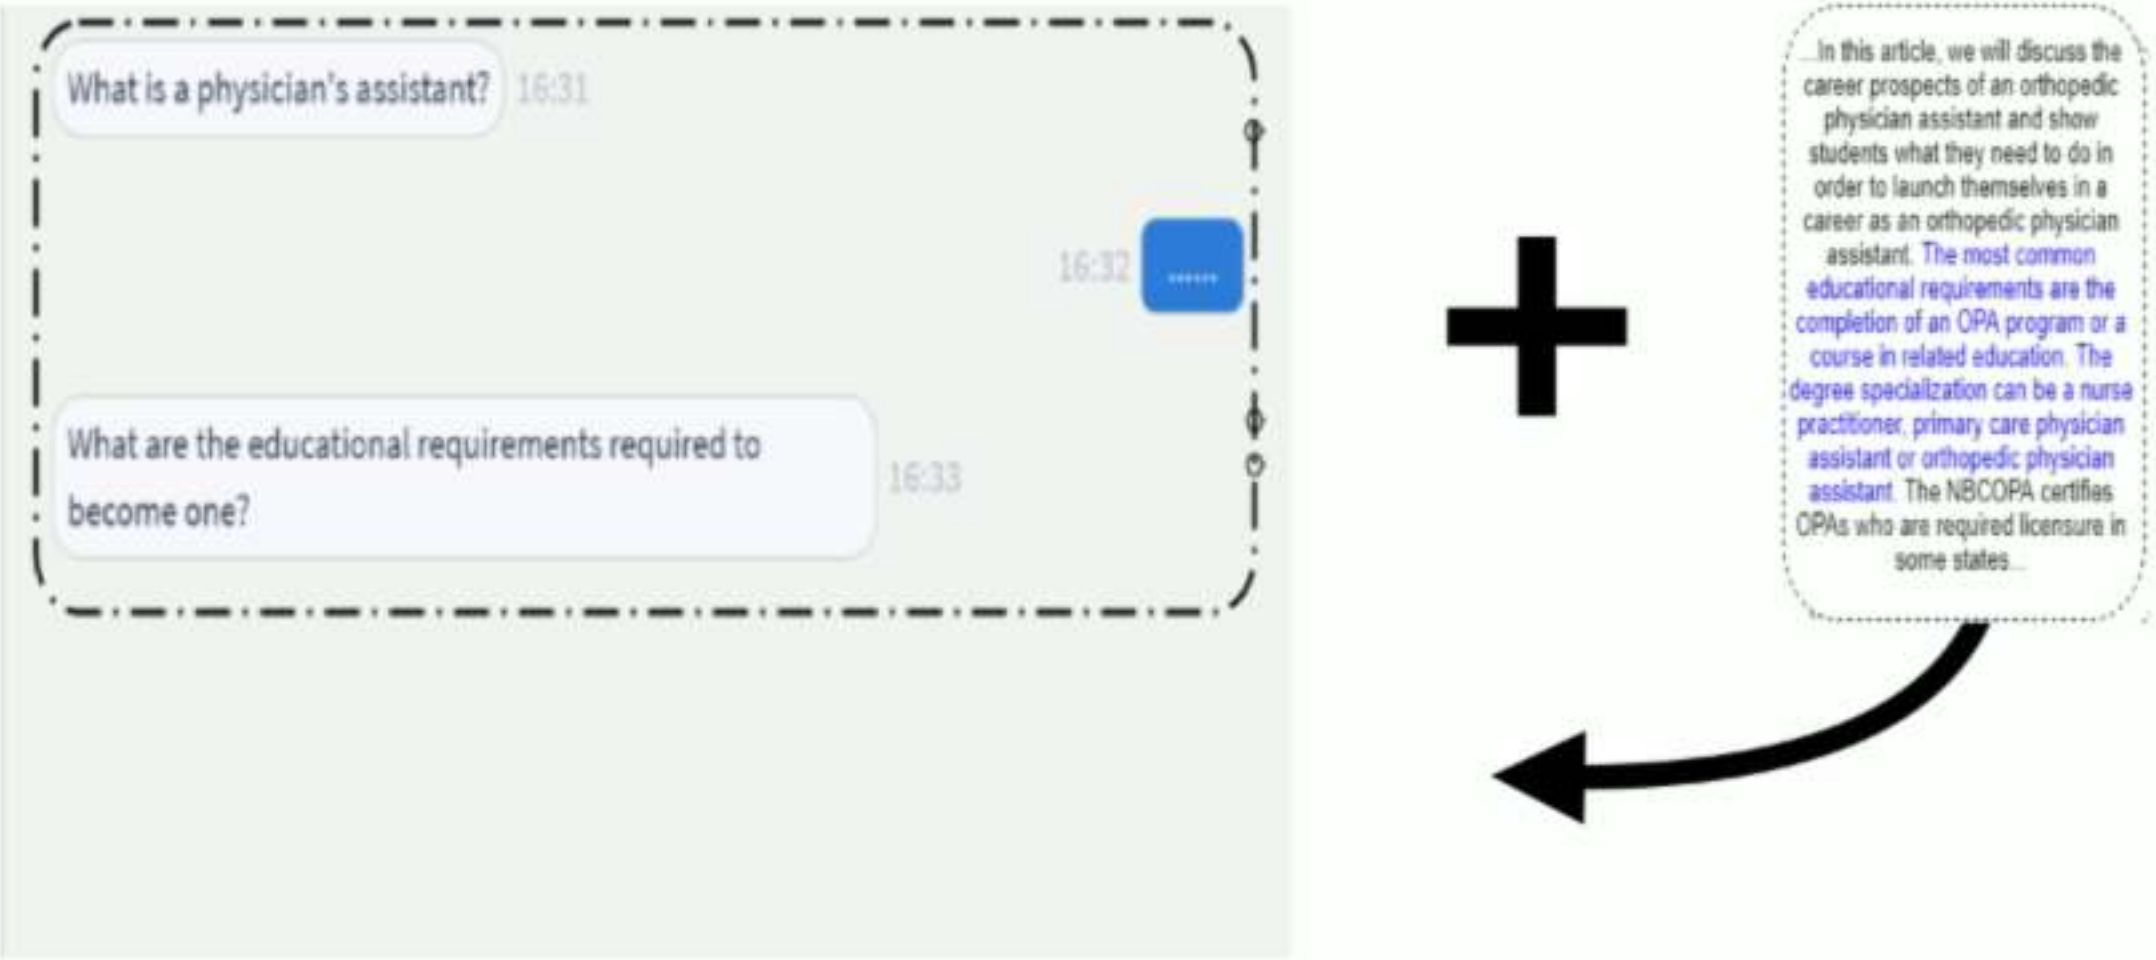

# Background-based conversation

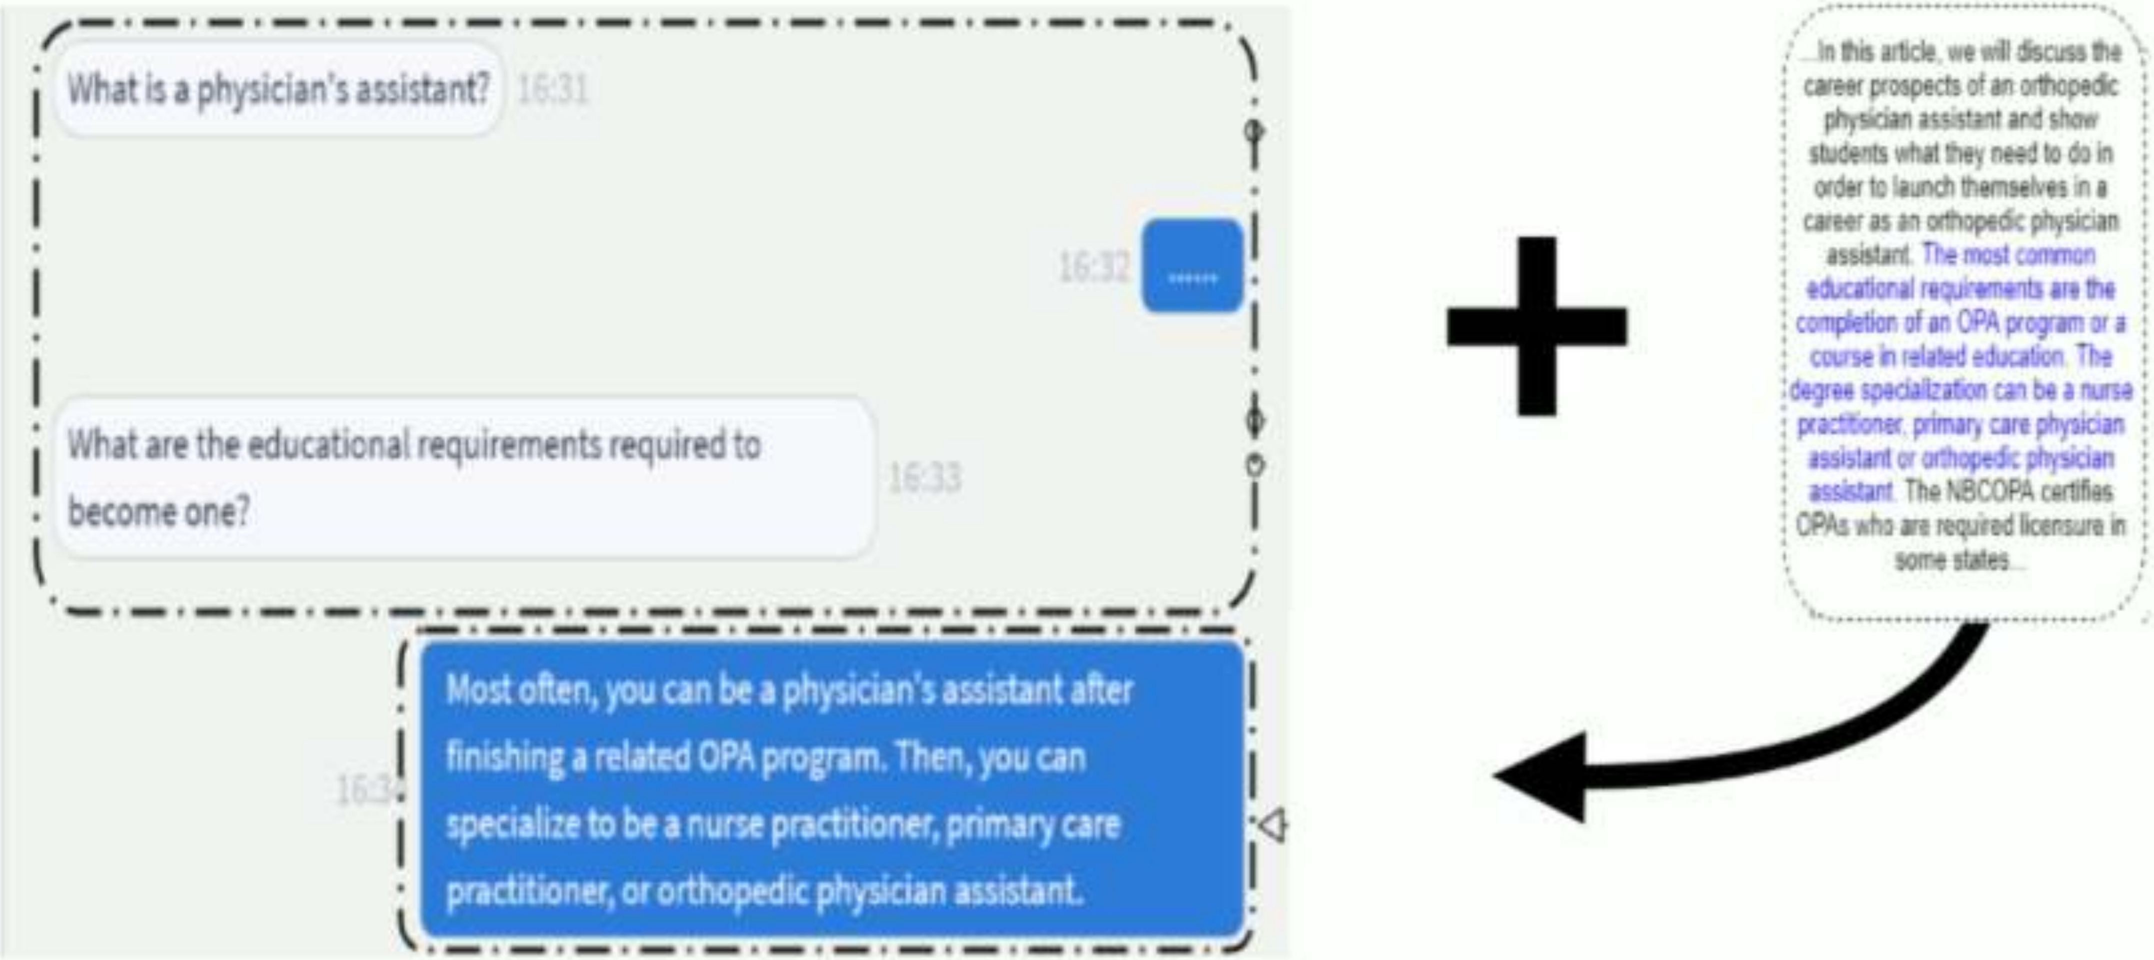

# Example BBC


*Scary Movie 3*

From Wikipedia, the free encyclopedia

This article **needs additional citations for verification**. Please help improve this article by adding citations to reliable sources. Unsourced material may be challenged and removed.
*Find sources:* "Scary Movie 3" – news · newspapers · books · scholar · JSTOR *(March 2016) (Learn how and when to remove this template message)*

***Scary Movie 3*** is a 2003 American horror sci-fi comedy film, which parodies the horror, sci-fi, and mystery genres. It is the third film in the *Scary Movie* franchise, and the first to be directed by David Zucker.

The film stars Anna Faris and Regina Hall reprising their roles as Cindy Campbell and Brenda Meeks, respectively. New cast members include Charlie Sheen, Simon Rex, Anthony Anderson, Kevin Hart, and Leslie Nielsen. It is the first film in the series to feature no involvement from the Wayans family. The characters of Shorty Meeks and Ray Wilkins, previously played by Shawn and Marlon Wayans, do not appear, nor are they referenced.

The film's plot significantly parodies the films *The Ring*, *Signs*, *The Matrix Reloaded* and *8 Mile*. The film grossed $220.7 million worldwide. It is the last film in the series to be released by The Walt Disney Company's subsidiary Miramax Films, under the brand Dimension Films. It was named the 2004 Teen Choice Awards in the category of *Choice Movie: Your Parents Didn't Want You to See*.[4]

**Contents** [show]

## Plot [ edit source ]

Katie (Jenny McCarthy) and Becca (Pamela Anderson) talk about what Katie believes is a sex tape, but Becca calls it a cursed tape. After several odd occurrences, and an almost atomic wedgie, they both die. Meanwhile, in a farm outside Washington, D.C., widowed farmer Tom Logan (Charlie Sheen) and his clumsy brother George (Simon Rex) discover a crop circle, saying "Attack Here!", after noticing the dog's strange activity (such as driving a tractor or getting stoned).

Cindy Campbell (Anna Faris), now a blonde news reporter, announces the crop circles on the

***Scary Movie 3***

Theatrical release poster

| | |
|---|---|
| **Directed by** | David Zucker |
| **Produced by** | Robert K. Weiss<br>David Zucker |
| **Written by** | Craig Mazin<br>Pat Proft |
| **Based on** | Characters<br>by Shawn Wayans<br>Marlon Wayans<br>Buddy Johnson |

# Example BBC

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

# Example BBC

**Background**

box office $ 110,000,082 awards bmi film & tv awards 2004 james l. venable mtv movie + tv awards 2004 best cameo taglines reload for a third shot great trilogies come in threes . this time it 's personal . you 're warned . 2003 . similar movies scary movie 4 scary movie 2 scary movie a haunted house scary movie 5

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

# Example BBC

**Background**

box office $ 110,000,082 awards bmi film & tv awards 2004 james l. venable <u>mtv movie + tv awards 2004 best cameo</u> taglines reload for a third shot great trilogies come in threes . this time it 's personal . you 're warned . 2003 . similar movies <u>scary movie 4</u> scary movie 2 scary movie a haunted house scary movie 5

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

**S2SA :** i think it made $ [UNK]

# Example BBC

**Background**

box office $ 110,000,082 awards bmi film & tv awards 2004 james l. venable mtv movie + tv awards 2004 best cameo taglines reload for a third shot great trilogies come in threes . this time it 's personal . you 're warned . 2003 . similar movies scary movie 4 scary movie 2 scary movie a haunted house scary movie 5

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

**S2SA :** i think it made $ [UNK]

**GTTP :** you should check out scary movie , 4

# Example BBC

**Background**

box office $ 110,000,082 awards bmi film & tv awards 2004 james l. venable mtv movie + tv awards 2004 best cameo taglines reload for a third shot great trilogies come in threes . this time it 's personal . you 're warned . 2003 . similar movies scary movie 4 scary movie 2 scary movie a haunted house scary movie 5

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

**S2SA :** i think it made $ [UNK]

**GTTP :** you should check out scary movie , 4

**BiDAF :** mtv movie + tv awards 2004 best cameo

# Example BBC

**Background**

box office $ 110,000,082 awards bmi film & tv awards 2004 james l. venable mtv movie + tv awards 2004 best cameo taglines reload for a third shot great trilogies come in threes . this time it 's personal . you 're warned . 2003 . similar movies scary movie 4 scary movie 2 scary movie a haunted house scary movie 5

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

**S2SA :** i think it made $ [UNK]

**GTTP :** you should check out scary movie , 4

**BiDAF :** mtv movie + tv awards 2004 best cameo

**Golden :** you are wrong . mtv movie + tv awards 2004 best cameo

# BBC

- Key challenge in BBC is **knowledge selection**: finding appropriate background knowledge (e.g., a text fragment about a movie plot, or some basic facts about a movie) based on which the next response is to be generated

- Two families of approaches

    - Extraction-based

    - Generation-based

# Extraction-based BBC

- Responses produced by extraction-based methods are directly copied from background sentences
- Learn two pointers to extract spans from background material as possible responses
- Generated responses are often not natural due to their extractive nature

**Background**

box office $ 110,000,082 awards bmi film & tv awards 2004 james l. venable mtv movie + tv awards 2004 best cameo taglines reload for a third shot great trilogies come in threes . this time it 's personal . you 're warned . 2003 . similar movies scary movie 4 scary movie 2 scary movie a haunted house scary movie 5

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

**S2SA :** i think it made $ [UNK]

**GTTP :** you should check out scary movie . 4

**BiDAF :** mtv movie + tv awards 2004 best cameo

**Golden :** you are wrong . mtv movie + tv awards 2004 best cameo

# Generation-based BBC

- Influence generation process using the background material
- Sequence-to-sequence models often have a hard time using the background model
- Natural responses but they often break semantic units

**Background**

box office $ 110,000,082 awards bmi film & tv awards 2004 james l. venable mtv movie + tv awards 2004 best cameo taglines reload for a third shot great trilogies come in threes . this time it 's personal . you 're warned . 2003 . similar movies scary movie 4 scary movie 2 scary movie a haunted house scary movie 5

**Conversation**

**human 1 :** was it worth money ?

**human 2 :** cheesy and trashy , but very entertaining i bet it did n't win any awards ?

**S2SA :** i think it made $ [UNK]

**GTTP :** you should check out scary movie , 4

**BIDAF :** mtv movie + tv awards 2004 best cameo

**Golden :** you are wrong . mtv movie + tv awards 2004 best cameo

# Advances in generation-based BBC

- Often based on **encoder-decoder** architecture
- Often adopt a **local** perspective, generating one token at a time based on current decoding state
- Recently proposed fixes
    - Teach decoder to select semantic units (instead of individual words)
    - Use structured knowledge from a knowledge graph (in addition to text) in which the background text is grounded

# Advances in generation-based BBC

- Equip encoder-decoder with a **global** perspective

P. Ren et al., 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. AAAI.

# Advances in generation-based BBC

- Equip encoder-decoder with a **global** perspective
  - *Background & context encoders*
    - Encode background knowledge and conversational context

P. Ren et al., 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. AAAI.

# Advances in generation-based BBC

- Equip encoder-decoder with a **global** perspective
  - *Background & context encoders*
    - Encode background knowledge and conversational context
  - *Local knowledge selection (LKS)*
    - Outputs response token from vocabulary or from background, based on GKS

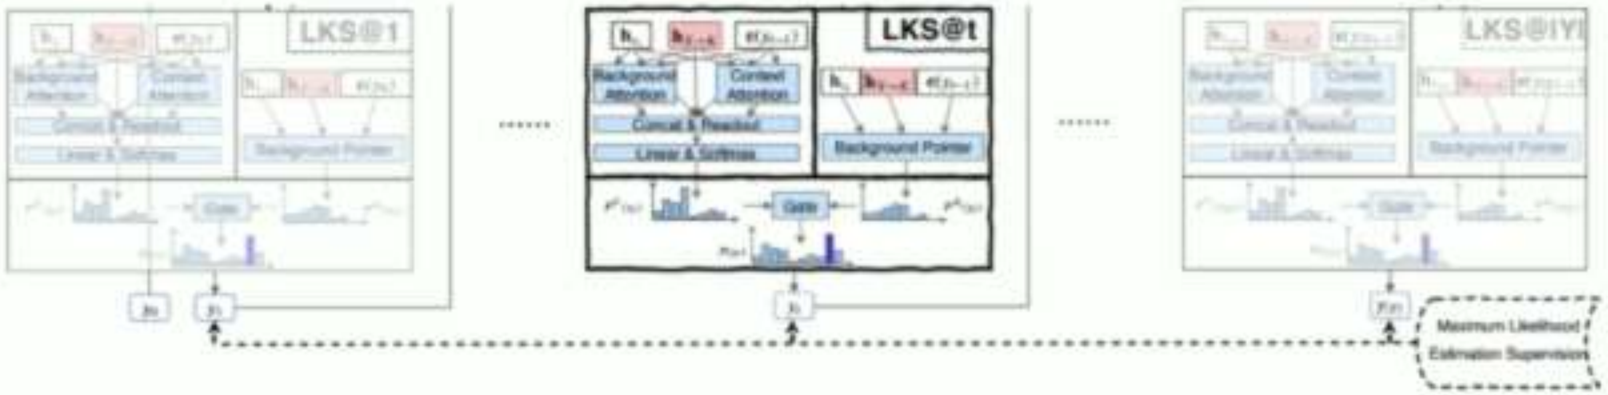


P. Ren et al., 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. AAAI.

# Advances in generation-based BBC

- Equip encoder-decoder with a **global** perspective
  - *Background & context encoders*
    - Encode background knowledge and conversational context
  - *Local knowledge selection (LKS)*
    - Outputs response token from vocabulary or from background, based on GKS
  - *Global knowledge selection (GKS)*
    - Evaluates matching between background and context, and decides what to talk about next

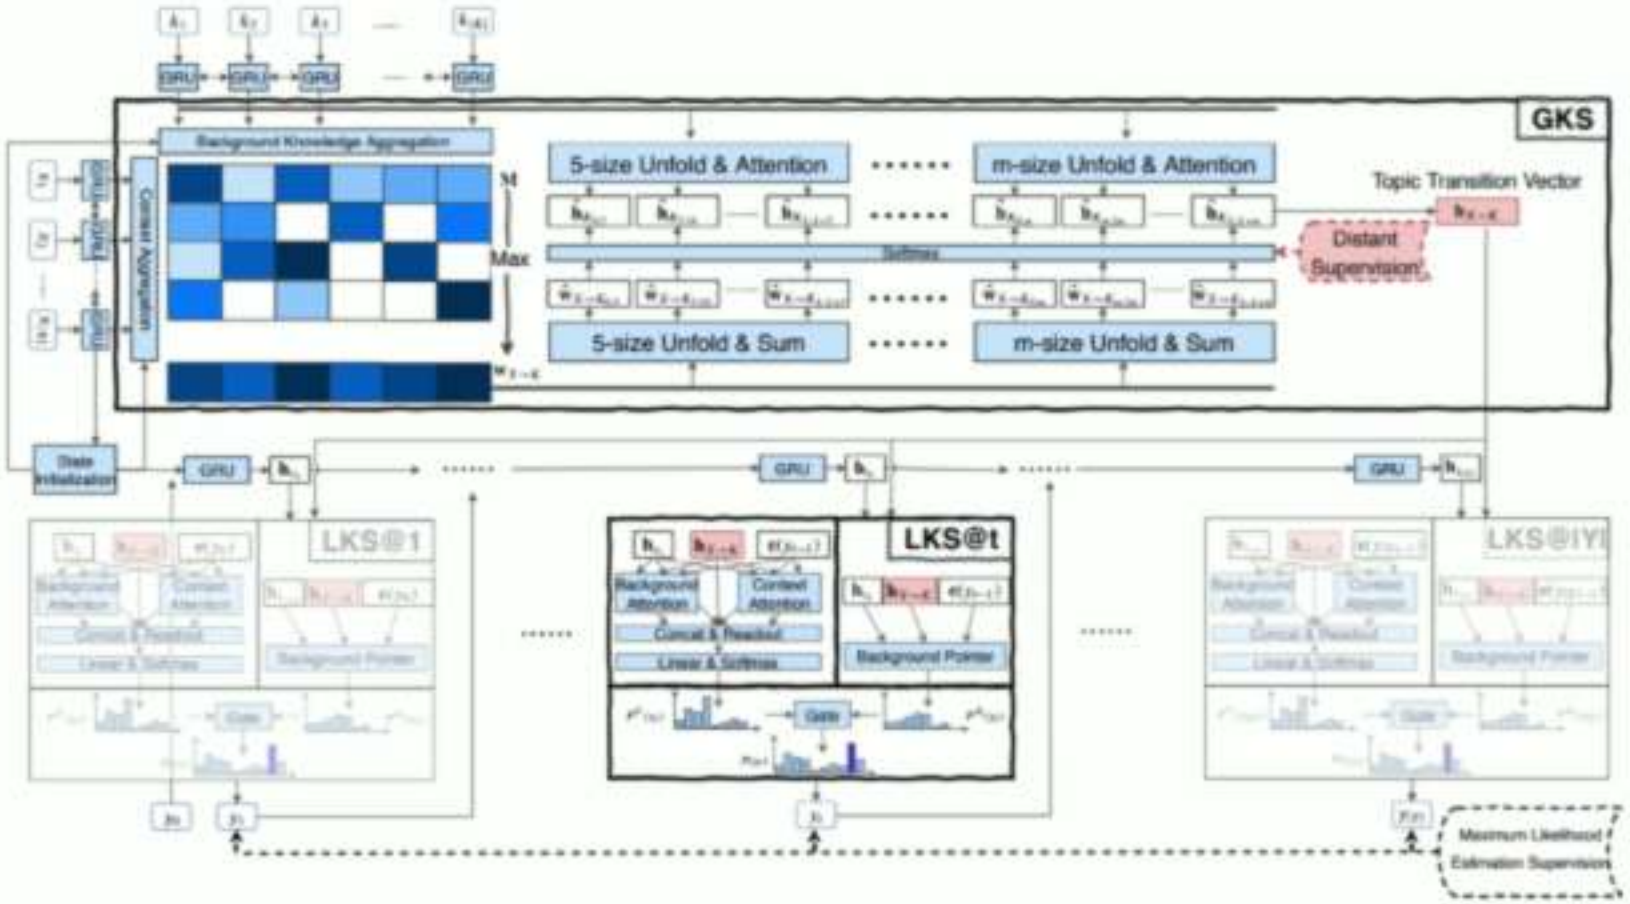


P. Ren et al., 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. AAAI.

# Global knowledge selection

- Given knowledge and conversational context, encode them into latent representations
- GKS module evaluates matching matrix between these representations
- Based on matching matrix, GKS decides on “what to talk about next” by selecting continuous spans from the background K to form a “topic transition vector”

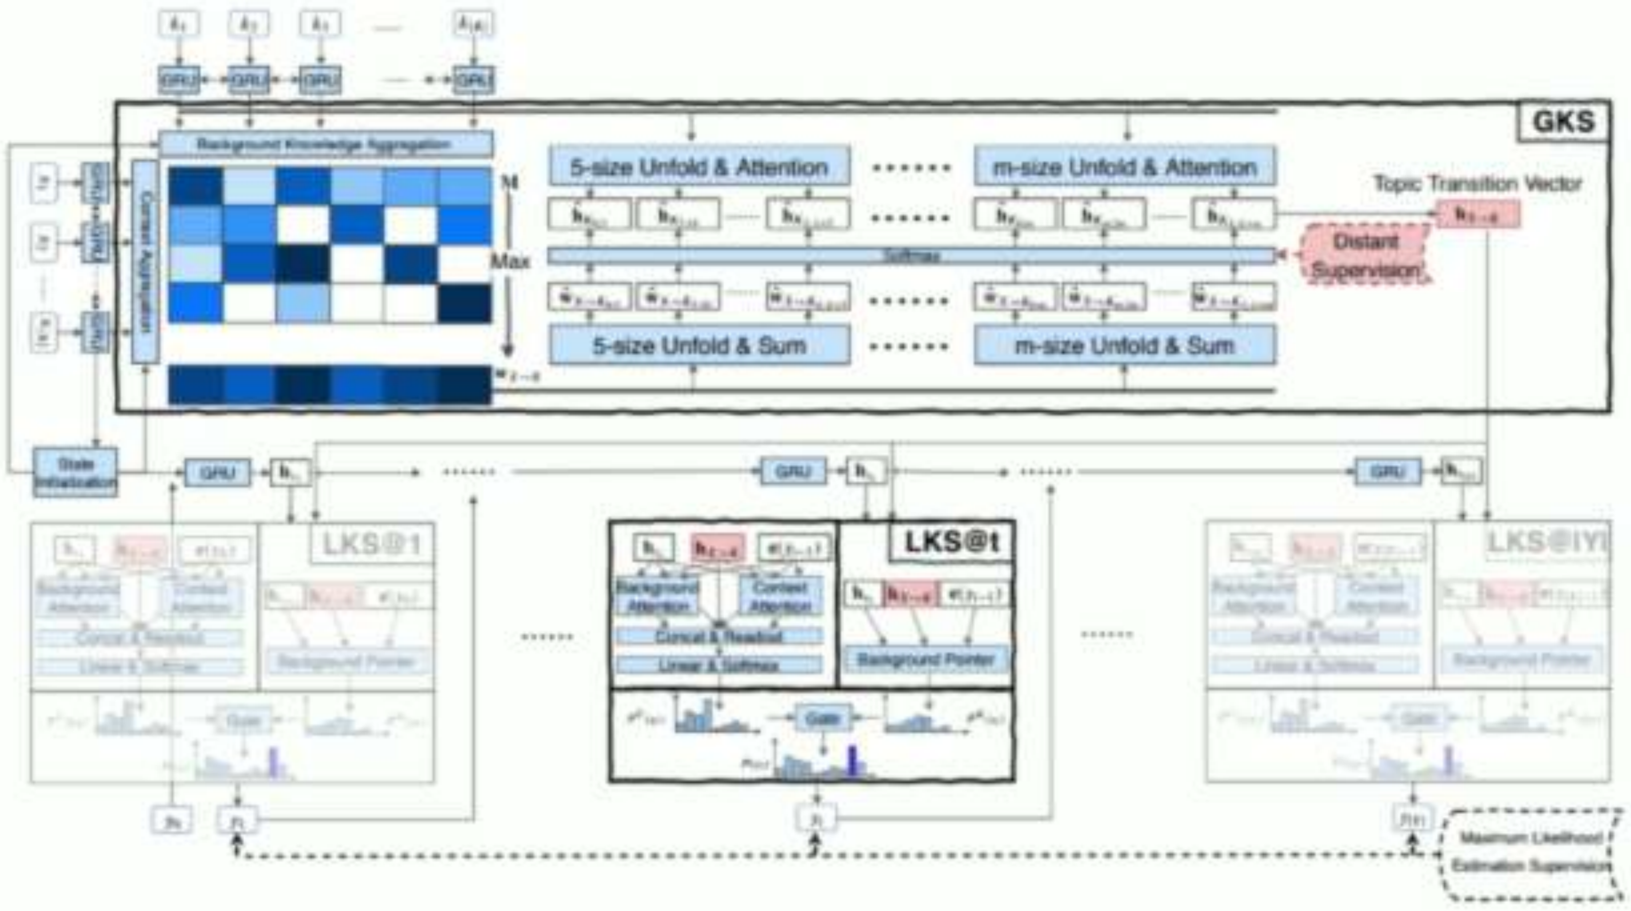


.

# Local knowledge selection

- At each decoding time step, LKS outputs response token by either generating from vocabulary or selecting from background K under the guidance of the topic transition vector produced by the GKS module

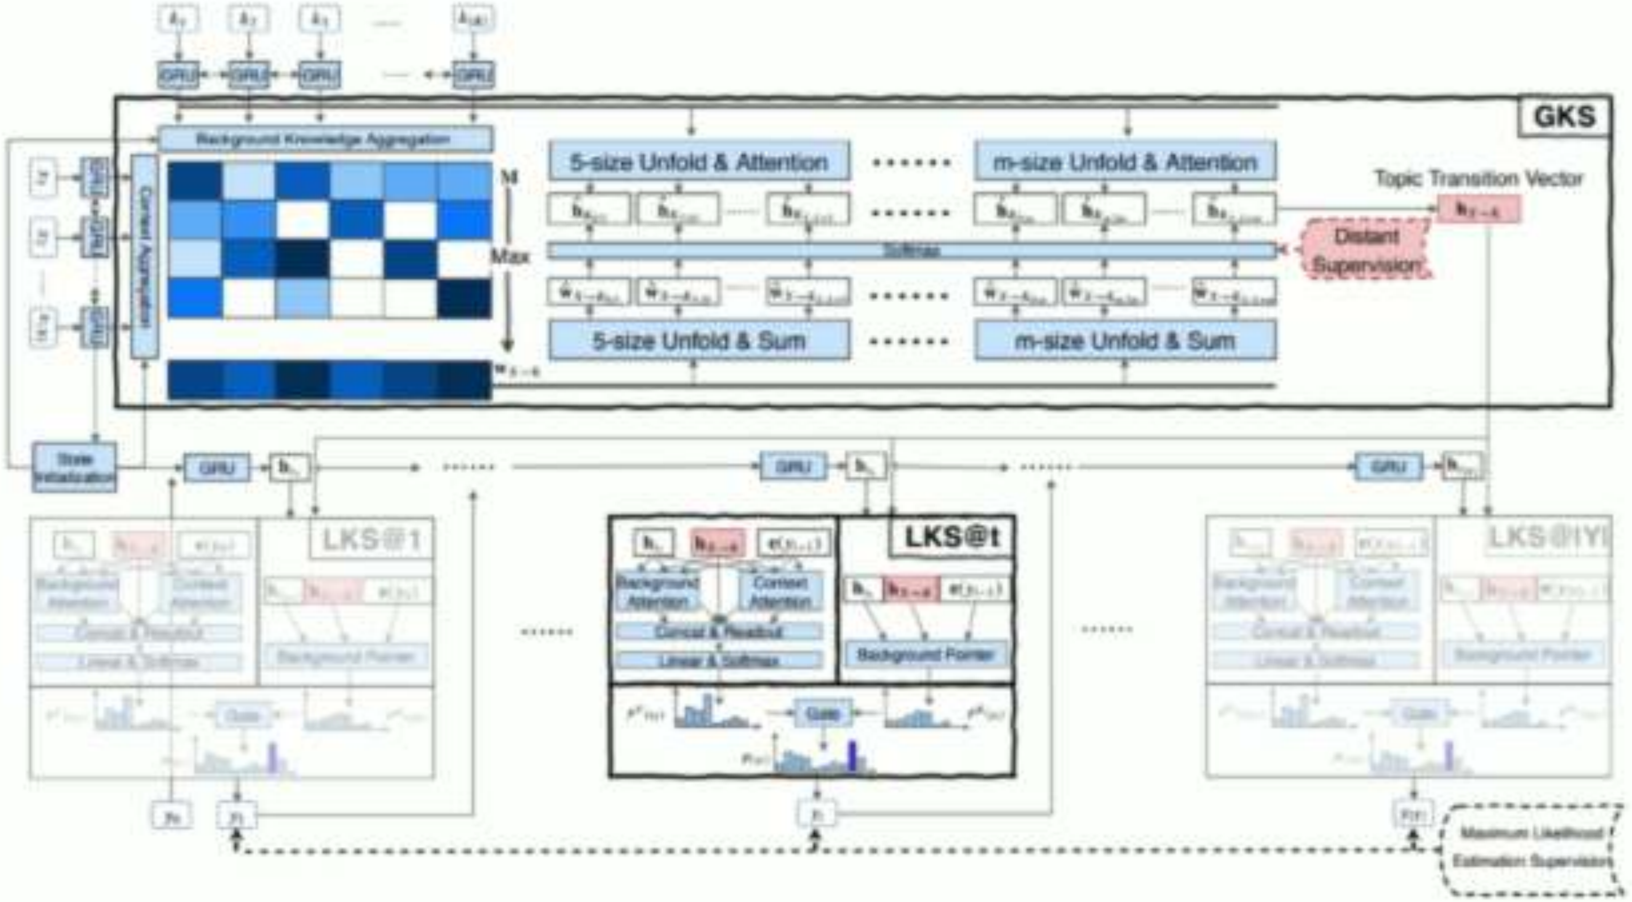


- Loss function a combination of three loss functions
  - Maximum Likelihood Estimation loss, Distant Supervision loss, Maximum Causal Entropy loss

# Advances in generation-based BBC

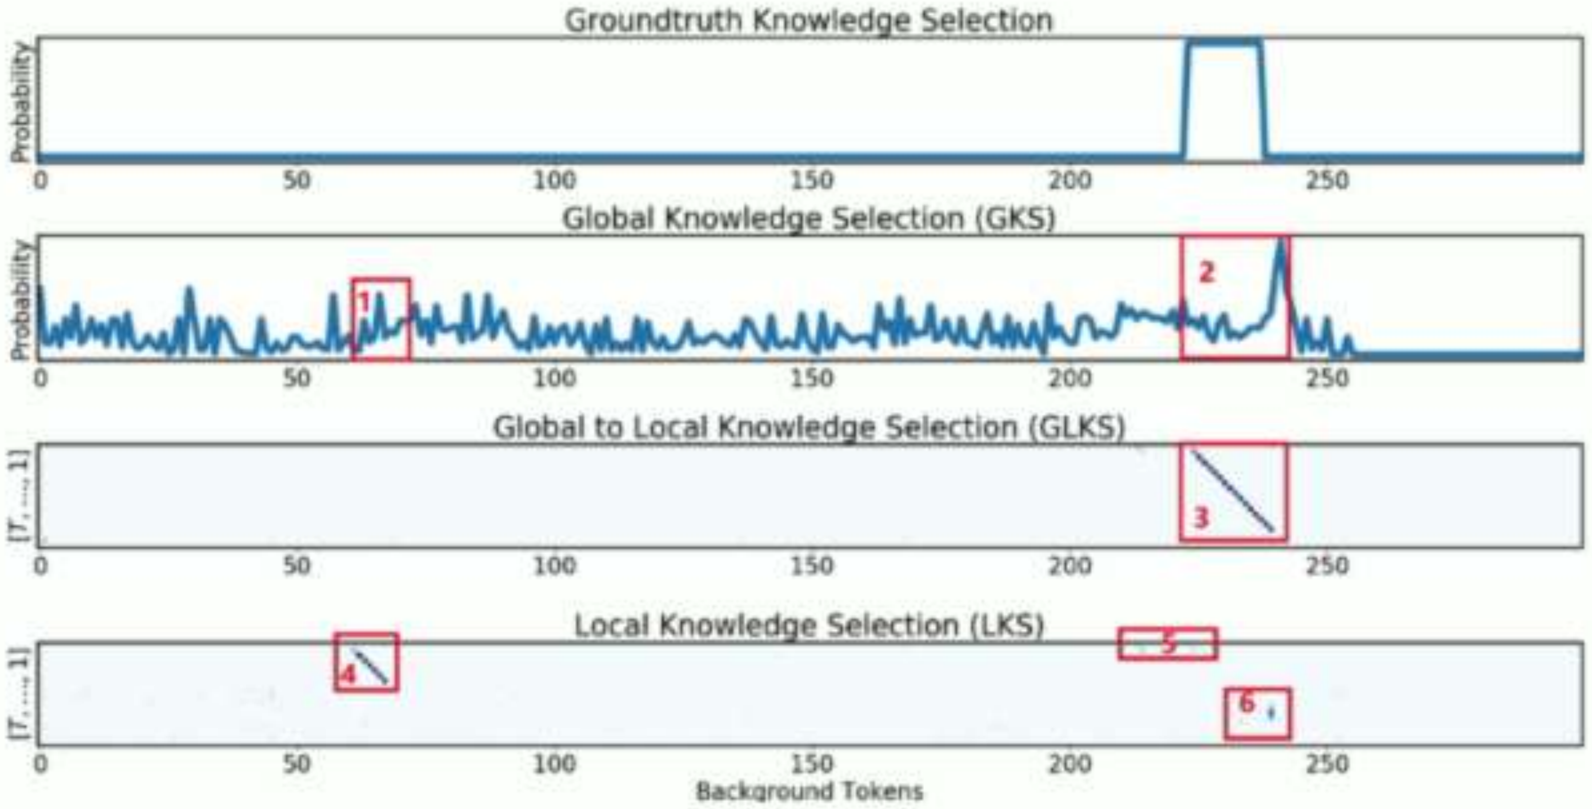

# Advances in generation-based BBC

**H1**: i loved all the tricks , and traps kevin created .

**H2**: me too , i loved when using a tape recorder , he tapes a message and slows down his voice , placing a hotel reservation .

**H1**: that was too funny , the hotel staff did n't believe him though .

*Dialogue context*

# Advances in generation-based BBC

**Backgound**: ... later that evening , he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee ... home alone 2 is a carbon copy , but it 's also much better and more complex a movie than the first ... regardless it 's a classic and i watch the first two movies every year ...

*Background*

**H1**: i loved all the tricks , and traps kevin created .

**H2**: me too , i loved when using a tape recorder , he tapes a message and slows down his voice , placing a hotel reservation .

**H1**: that was too funny , the hotel staff did n't believe him though .

*Dialogue context*

# Advances in generation-based BBC

| | | |
|---|---|---|
| | **Backgound**: ... later that evening , he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee ... home alone 2 is a carbon copy , but it 's also much better and more complex a movie than the first ... regardless it 's a classic and i watch the first two movies every year ... | *Background* |
| | **H1**: i loved all the tricks , and traps kevin created .<br>**H2**: me too , i loved when using a tape recorder , he tapes a message and slows down his voice , placing a hotel reservation .<br>**H1**: that was too funny , the hotel staff did n't believe him though . | *Dialogue context* |
| GTTP | it 's a classic and i watch the first two movies every year . | *Generation-based* |

# Advances in generation-based BBC

| | | |
|---|---|---|
| | **Backgound**: ... later that evening , he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee ... home alone 2 is a carbon copy , but it 's also much better and more complex a movie than the first ... regardless it 's a classic and i watch the first two movies every year ... | *Background* |
| | **H1**: i loved all the tricks , and traps kevin created .<br>**H2**: me too , i loved when using a tape recorder , he tapes a message and slows down his voice , placing a hotel reservation .<br>**H1**: that was too funny , the hotel staff did n't believe him though . | *Dialogue context* |
| GTTP | it 's a classic and i watch the first two movies every year . | *Generation-based* |
| RefNet | that it was so sad when he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee . | *Generation-based* |

# Advances in generation-based BBC

| | | |
|---|---|---|
| | **Backgound**: ... later that evening , he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee ... home alone 2 is a carbon copy , but it 's also much better and more complex a movie than the first ... regardless it 's a classic and i watch the first two movies every year ... | *Background* |
| | **H1**: i loved all the tricks , and traps kevin created .<br>**H2**: me too , i loved when using a tape recorder , he tapes a message and slows down his voice , placing a hotel reservation .<br>**H1**: that was too funny , the hotel staff did n't believe him though . | *Dialogue context* |
| GTTP | it 's a classic and i watch the first two movies every year . | *Generation-based* |
| RefNet | that it was so sad when he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee . | *Generation-based* |
| LKS | i know , it was a carbon copy , but it 's also much better and more complex a movie than the first . | *Generation-based, local only* |

# Advances in generation-based BBC

| | | |
|---|---|---|
| | **Backgound**: ... later that evening , he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee ... home alone 2 is a carbon copy , but it 's also much better and more complex a movie than the first ... regardless it 's a classic and i watch the first two movies every year ... | *Background* |
| | **H1**: i loved all the tricks , and traps kevin created .<br>**H2**: me too , i loved when using a tape recorder , he tapes a message and slows down his voice , placing a hotel reservation .<br>**H1**: that was too funny , the hotel staff did n't believe him though . | *Dialogue context* |
| GTTP | it 's a classic and i watch the first two movies every year . | *Generation-based* |
| RefNet | that it was so sad when he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee . | *Generation-based* |
| LKS | i know , it was a carbon copy , but it 's also much better and more complex a movie than the first . | *Generation-based, local only* |
| GLKS | so true , later that evening , he intends to access kevin 's room , but kevin fools him into thinking that he has walked in on his father , causing the concierge to flee . | *Generation-based, global + local* |

# Advances in generation-based BBC

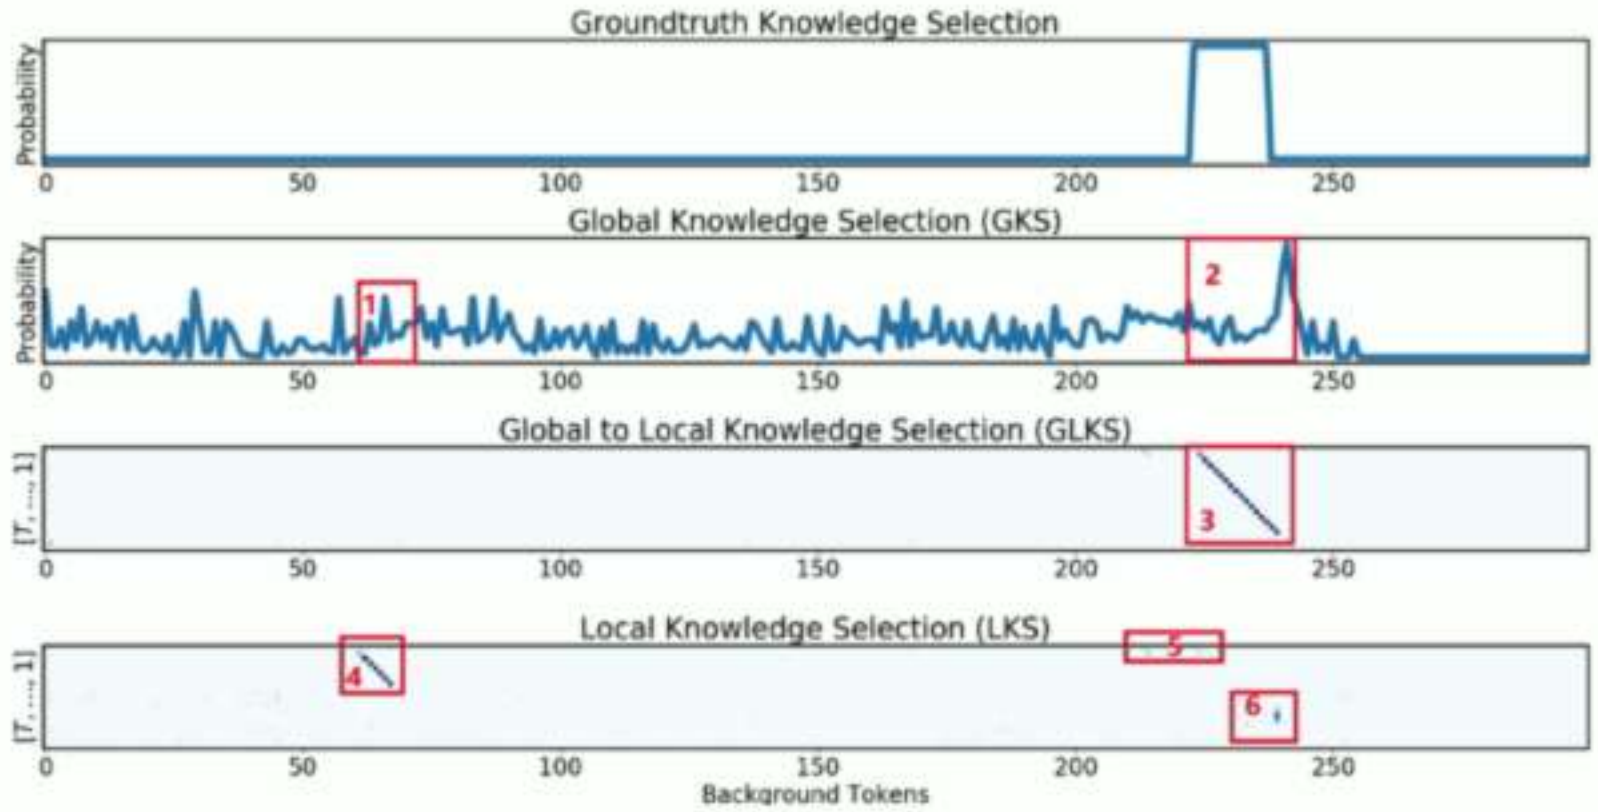

# Recent results on BBC

- **Holl-E dataset (Moghe et al., 2018)**
  - Built for movie chats in which each response is explicitly generated by copying and/or modifying sentences from the background
  - Background consists of plots, comments and reviews about movies collected from different websites

| | ROUGE-1 | | ROUGE-2 | | ROUGE-L | |
|---|---|---|---|---|---|---|
| | SR | MR | SR | MR | SR | MR |
| **no background** | | | | | | |
| S2S | 27.15 | 30.91 | 09.56 | 11.85 | 21.48 | 24.81 |
| HRED | 24.55 | 25.38 | 07.61 | 08.35 | 18.87 | 19.67 |
| **oracle background** | | | | | | |
| S2SA | 27.97 | 32.65 | 14.50 | 18.22 | 23.23 | 27.55 |
| GTTP | 29.82 | 35.08 | 17.33 | 22.00 | 25.08 | 30.06 |
| CaKe | 42.82 | 48.65 | 30.37 | 36.54 | 37.48 | 43.21 |
| RefNet | 42.87 | 49.64 | 30.73 | 38.15 | 37.11 | 43.77 |
| GLKS | **43.75*** | **50.67*** | **31.54*** | **39.20*** | **38.69*** | **45.64*** |
| **256-word background** | | | | | | |
| S2SA | 26.36 | 30.76 | 13.36 | 16.69 | 21.96 | 25.99 |
| GTTP | 30.77 | 36.06 | 18.72 | 23.70 | 25.67 | 30.69 |
| CaKe | 41.26 | 45.81 | 29.43 | 34.00 | 36.01 | 40.79 |
| RefNet | 41.33 | 47.00 | 31.08 | 36.50 | 36.17 | 41.72 |
| AKGCM | – | – | 31.87 | – | 37.09 | – |
| GLKS | **44.52*** | **50.06*** | **33.05*** | 38.87* | **39.63*** | **45.12*** |

| | Improved GTTP | | RefNet | | GLKS | |
|---|---|---|---|---|---|---|
| | ≥1 | ≥2 | ≥1 | ≥2 | ≥1 | ≥2 |
| N | 307 | 115 | 391 | 213 | **424** | **226** |
| I | 271 | 89 | **411** | **244** | 401 | 199 |
| A | 318 | 111 | 371 | 180 | **406** | **219** |
| H | 332 | 123 | 394 | 225 | **436** | **263** |

$\geq n$ means that at least $n$ MTurk workers think it is a good response w.r.t. *Naturalness* (**N**), *Informativeness* (**I**), *Appropriateness* (**A**) and *Humanness* (**H**).

*Out of a sample of 500*

P. Ren et al., 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. AAAI.

# What’s next?

# Background-based conversations

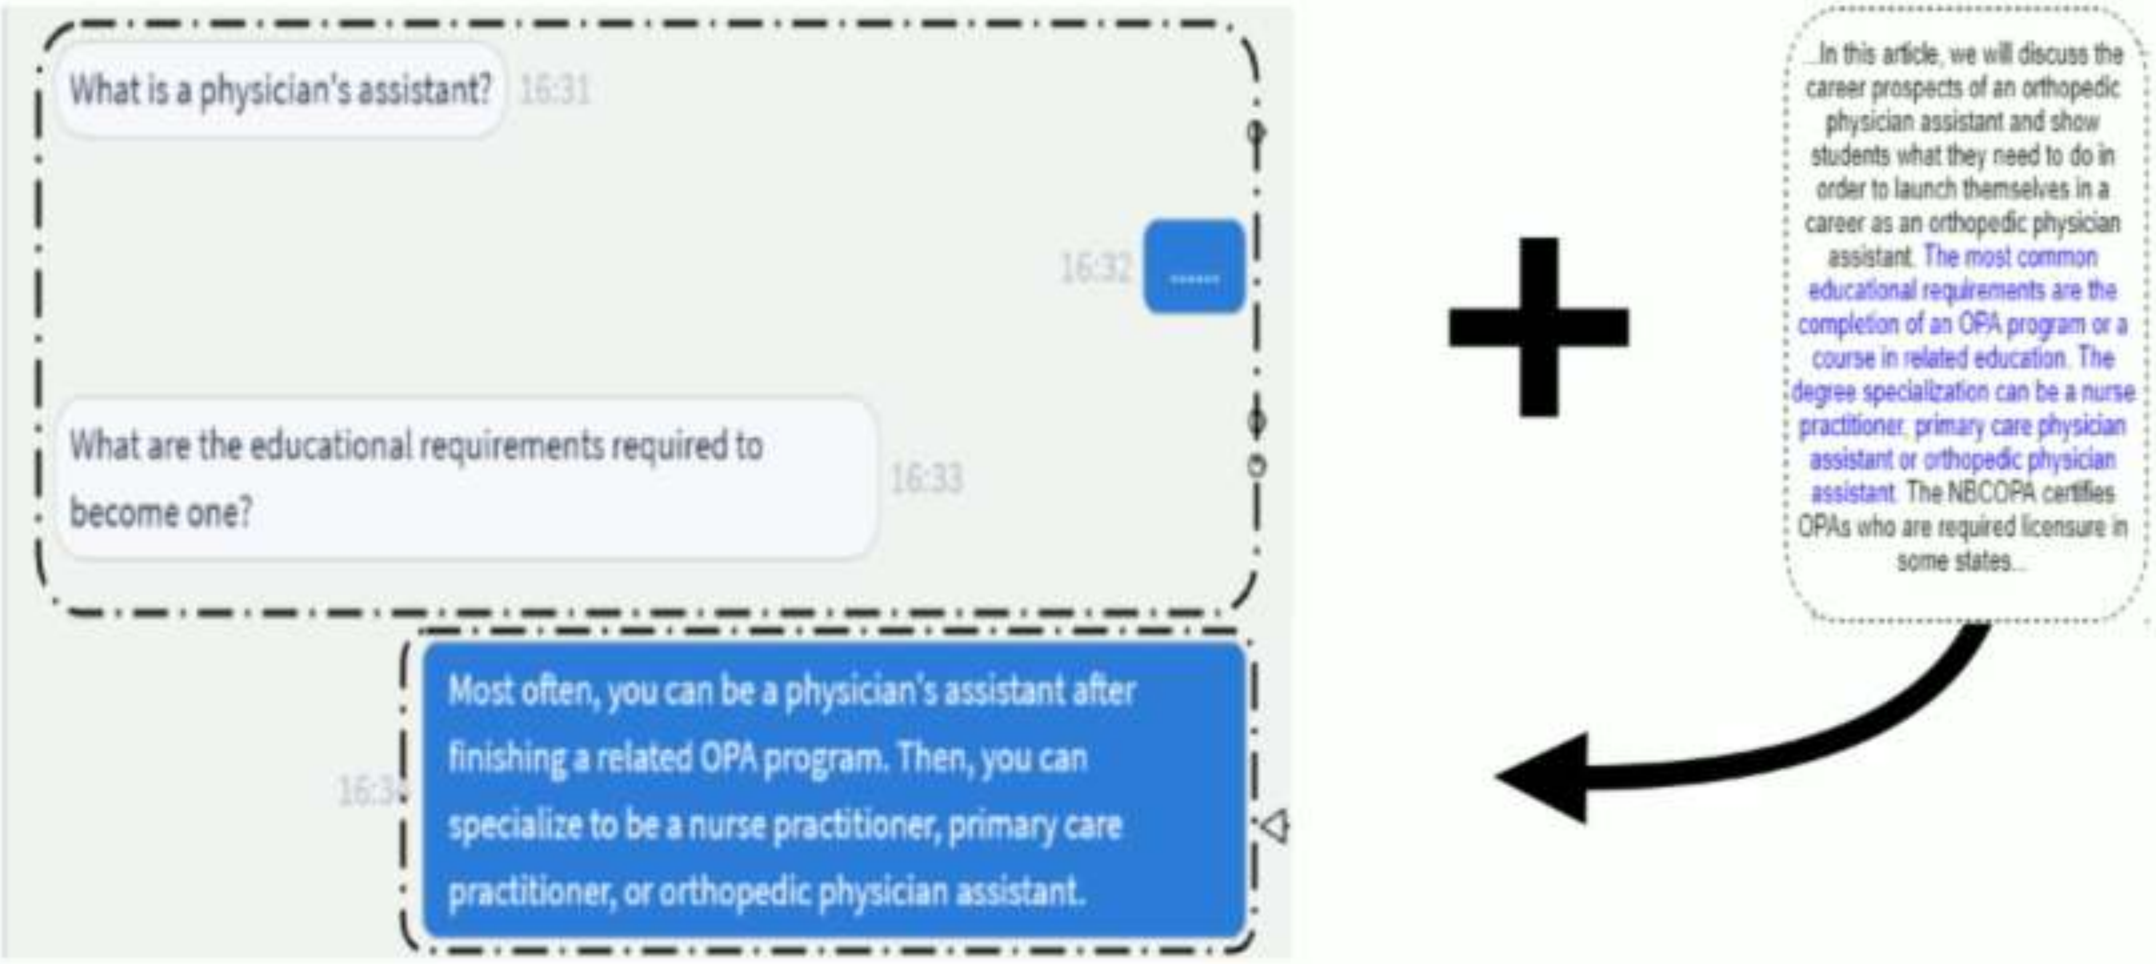

# SERP-based conversations

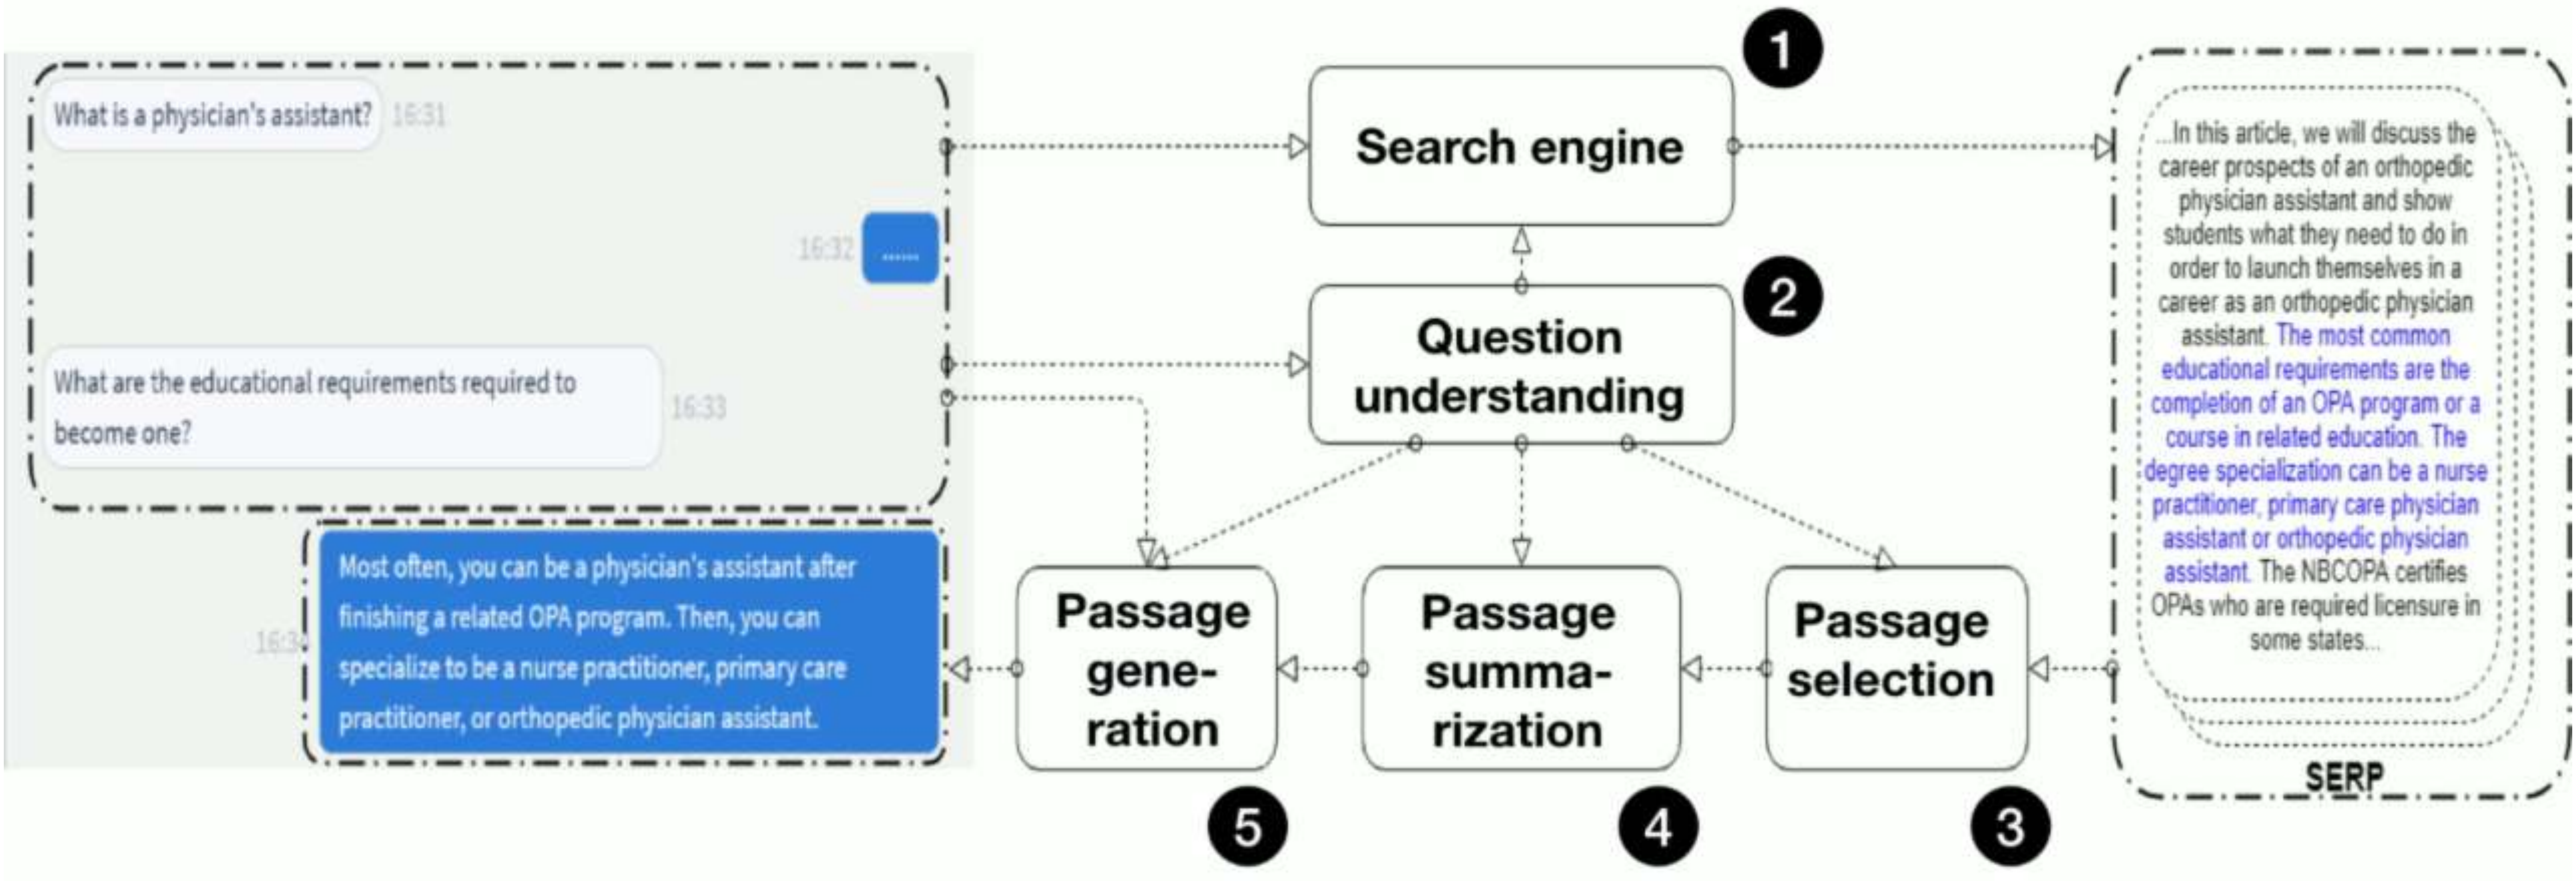

# ❶ Query formulation

- Learning reformulations so as to obtain the best results
  - Optimize for generating a response
  - Dynamically update formulations as conversation unfolds and user issues questions, answers and interactions
  - Query completion meets RL meets BERT


D. Odijk et al. 2015. Dynamic query modeling for related content finding. SIGIR

# SERP-based conversations

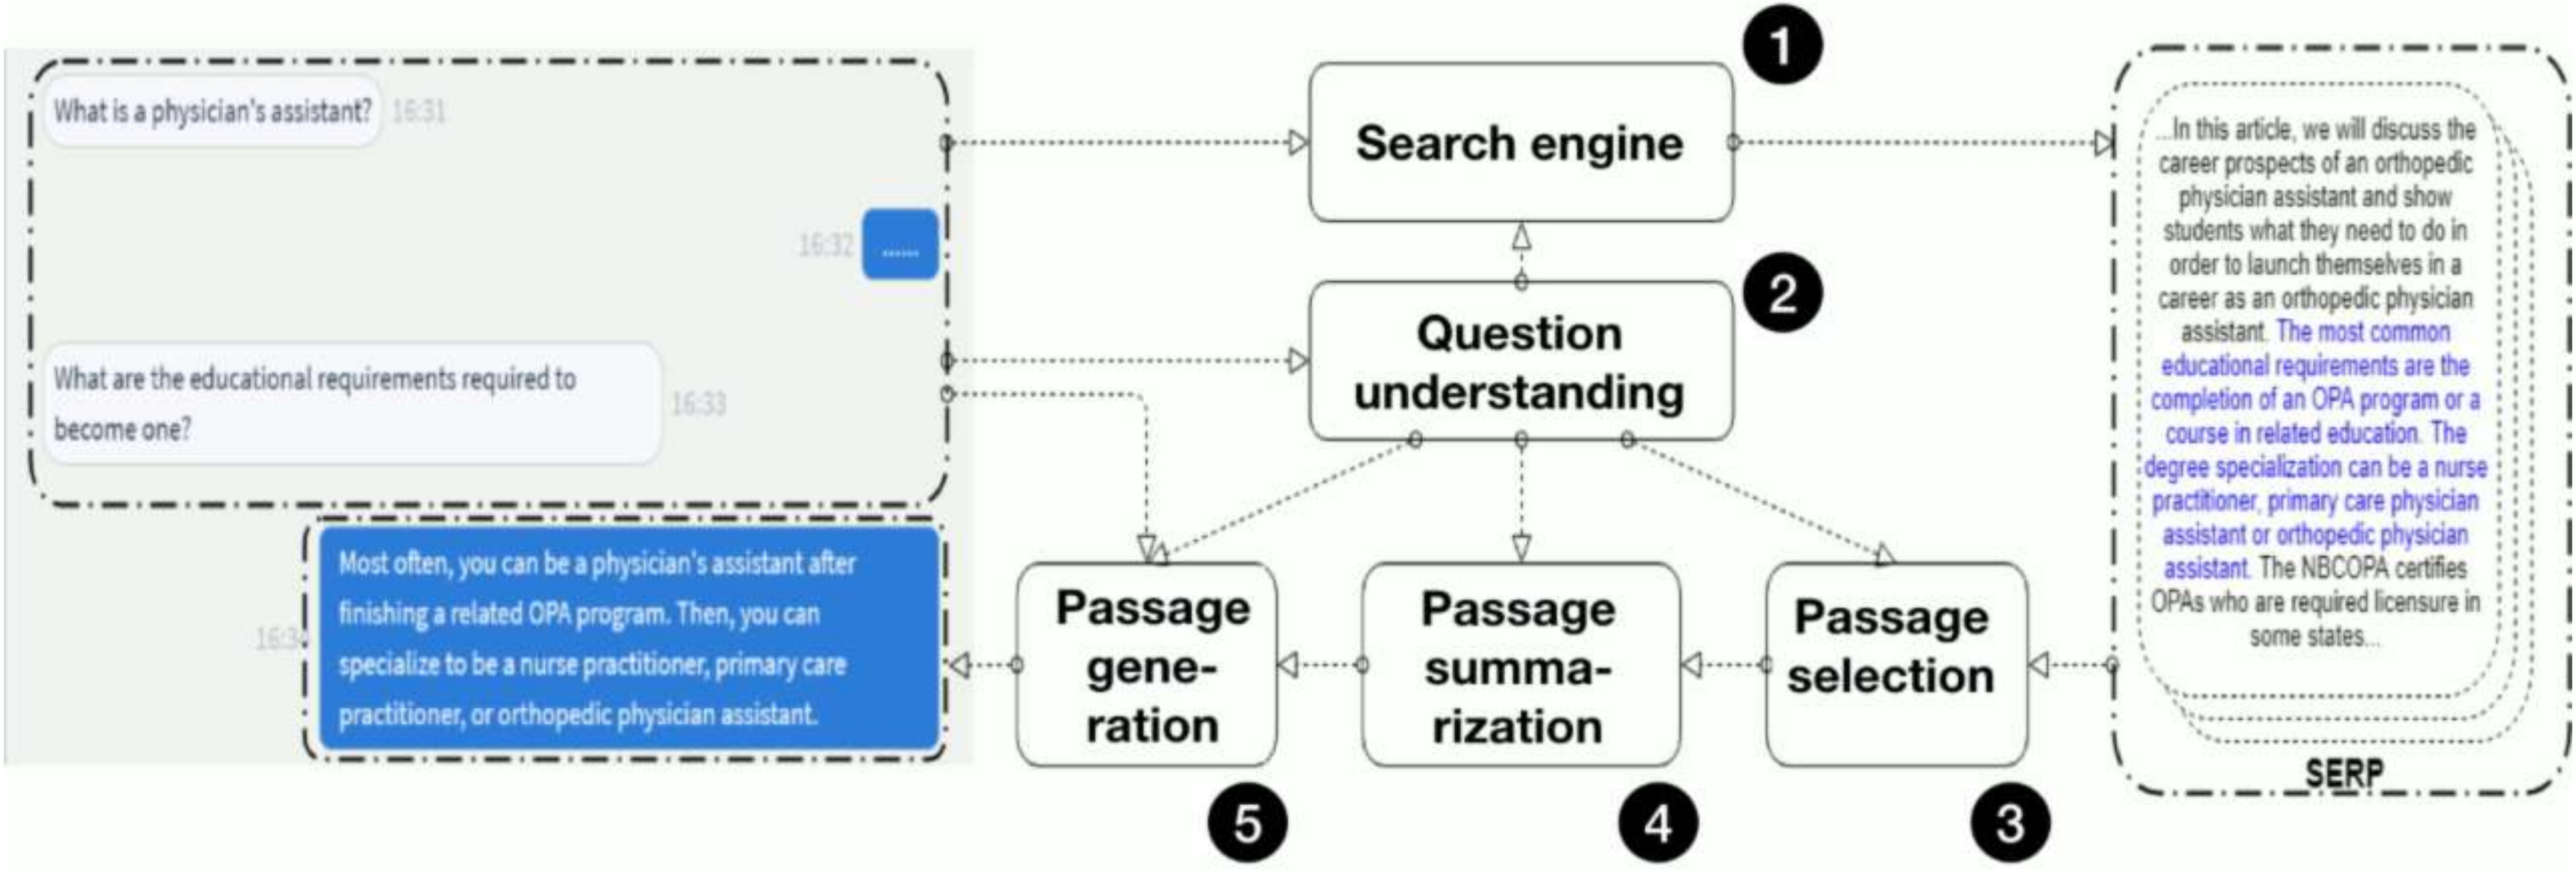

❶ # Query formulation

- Learning reformulations so as to obtain the best results
  - Optimize for generating a response
  - Dynamically update formulations as conversation unfolds and user issues questions, answers and interactions
  - Query completion meets RL meets BERT


Dynamic Query Modeling for Related Content Finding

Daan Odijk[1] d.odijk@uva.nl
Edgar Meij[2] emeij@yahoo-inc.com
Isaac Sijaranamual[1] i.b.sijaranamual@uva.nl
Maarten de Rijke[1] derijke@uva.nl
[1] University of Amsterdam, Amsterdam, The Netherlands
[2] Yahoo Labs, London, United Kingdom

ABSTRACT

Categories and Subject Descriptors

Keywords

1. INTRODUCTION

D. Odijk et al. 2015. Dynamic query modeling for related content finding. SIGIR

# ❷ Question understanding

- Inferring intent and shifts in intent
- Recognizing entities and relations
- Learning to ask in mixed-initiative setting

Preference Elicitation as an Optimization Problem

Anna Sepliarskaia
University of Amsterdam
Amsterdam, The Netherlands

Julia Kiseleva
University of Amsterdam
Amsterdam, The Netherlands

Filip Radlinski
Google
London, UK

Maarten de Rijke
University of Amsterdam
Amsterdam, The Netherlands

ABSTRACT

1 INTRODUCTION

CCS CONCEPTS

KEYWORDS

A. Sepliarskaia et al., 2018. Preference Elicitation as an Optimization Problem. In *RecSys*.

# ❸ Passage selection

- Single passage **may not be enough** to answer question or generate response

- Taking **broader context** into account
  - Consider low-ranked documents that are not immediately relevant, combining information from multiple documents, and reasoning over multiple facts from these documents to infer the answer/response


Learning to Transform, Combine, and Reason
in Open-Domain Question Answering

Mostafa Dehghani
University of Amsterdam

Hosein Azarbonyad
University of Amsterdam

Jaap Kamps
University of Amsterdam

Maarten de Rijke
University of Amsterdam

ABSTRACT

CCS CONCEPTS

KEYWORDS

1 INTRODUCTION

M. Dehghani et al., 2019. Learning to Transform, Combine, and Reason in Open-Domain Question Answering. In *WSDM*.

# ❹ Passage summarization

- **Multi-document** summarization
- Summarize multiple **multi-modal** search results
  - Dealing with multiple intents
  - Text vs. image/video vs. knowledge cards vs. …
  - Structured vs. unstructured
  - Static blogs/articles vs. live news/reports
- Optimize for "ideal summary" vs for "successful interaction"

5

# Response generation

- **Repetition** problem
- **Appropriateness** problem
  - Classifying responses into appropriate vs not-appropriate
- **No single correct response**
  - User inputs are not always questions with definitive answers or unique responses
- What to **talk about next**?
  - Support exploration, serendipity, …

Improving Neural Response Diversity with Frequency-Aware Cross-Entropy Loss

Shaojie Jiang
University of Amsterdam
Amsterdam, The Netherlands

Pengjie Ren
University of Amsterdam
Amsterdam, The Netherlands

Christof Monz
University of Amsterdam
Amsterdam, The Netherlands

Maarten de Rijke
University of Amsterdam
Amsterdam, The Netherlands

ABSTRACT

CCS CONCEPTS

KEYWORDS

1 INTRODUCTION

S. Jiang et al. 2019. Improving Neural Response Diversity with Frequency-Aware Cross-Entropy Loss. WWW

5

# Response generation

- **Why did you so that?**
  - Explainable for developers
    - Failure analysis
    - Identifying influential (online) training instances (Sharchilev et al., ICML 2018)
    - Reasoning path on knowledge graph as explanations (Moon et al., ACL 2019, Liu et al., arXiv 2019)
  - Explainable for users
    - Response/answer explanation

**Human: which is your favorite character in this ?**
**Bot: my favorite character was obviously the main character because through his perseverance he was able to escape a dangerous situation .**

# Wrap-up

# Wrap-up

- SERP-grounded conversations
    - Bring the richness of SERPs to a conversational setting
- General idea, recent advances, challenges and ambitions
- **Work in progress, a lot remains to be done**

# References


- Belkin Anomalous states of knowledge as a basis for information retrieval. Canadian J. Information Science, 1980
- Dehghani et al. Learning to Transform, Combine, and Reason in Open-Domain Question Answering, WSDM 2019
- Dinan et al. Wizard of Wikipedia: Knowledge-Powered Conversational agents. ICLR 2019
- Ghazvininejad et al. A Knowledge-Grounded Neural Conversation Model. AAAI 2018
- Kiseleva and de Rijke. Evaluating Personal Assistants on Mobile Devices. CAIR 2017
- Jiang et al. Improving Neural Response Diversity with Frequency-Aware Cross-Entropy Loss. WWW 2019
- Li et al. Incremental Transformer with Deliberation Decoder for Document Grounded Conversations. ACL 2019
- Lian et al. Learning to Select Knowledge for Response Generation in Dialog Systems. arXiv 2019
- Liu et al. Knowledge Aware Conversation Generation with Explainable Reasoning on Augmented Graphs. arXiv 2019
- Meng et al. RefNet: A Reference-aware Network for Background Based Conversation. AAAI 2020
- Moghe et al. Towards Exploiting Background Knowledge for Building Conversation Systems. EMNLP 2018
- Moon et al. OpenDialKG: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs. ACL 2019
- Odijk et al. Dynamic Query Modeling for Related Content Finding. SIGIR 2015
- Qin et al. Conversing by Reading: Contentful Neural Conversation with On-demand Machine Reading. ACL 2019
- Radlinski and Craswell. A Theoretical Framework for Conversational Search. CHIIR 2017
- Ren et al. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. AAAI 2020
- Rose and Levinson. Understanding User Goals in Web Search. WWW 2004
- See et al. Get To The Point: Summarization with Pointer-Generator Networks. ACL 2017
- Seo et al. Bidirectional Attention Flow for Machine Comprehension. ICLR 2017
- Sharchilev et al. Finding Influential Training Samples for Gradient Boosted Decision Trees. ICML 2018
- Sepliarskaia et al. Preference Elicitation as an Optimization Problem. RecSys 2018
- Serban et al. Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models. AAAI 2016
- Serban et al. A Hierarchical Latent Variable Encoder-Decoder Model for Generating Dialogues. AAAI 2017
- Vakulenko et al. Conversational Exploratory Search via Interactive Storytelling. SCAI 2017
- Vakulenko et al. Measuring Semantic Coherence of a Conversation. ISWC 2018
- Vakulenko et al. Conversational Browsing: Dialog-based Access to Structured Information Sources. *Under review*, 2019
- Zhang et al. Personalizing Dialogue Agents: I have a dog, do you have pets too? ACL 2018
- Zhang et al. Improving Background Based Conversation with Context-aware Knowledge Pre-selection. SCAI 2019.
- Zhou et al. Commonsense Knowledge Aware Conversation Generation with Graph Attention. IJCAI 2018
- Zhou et al. A Dataset for Document Grounded Conversations. EMNLP 2018