Towards Interfaces in Natural Language

Machine Learning Meetup, Bratislava, 2018

Zuzana Nevěřilová

Konica Minolta Laboratory Europe, Brno, Czech Republic

October 23, 2018

Konica Minolta Laboratory Europe

London
HQ + Exploration Lab

Brno
Digital Workplace Lab

Munich
Digital Healthcare Lab

Rome
Smart Systems Lab

KONICA MINOLTA

Digital workplace: Cognitive Hub

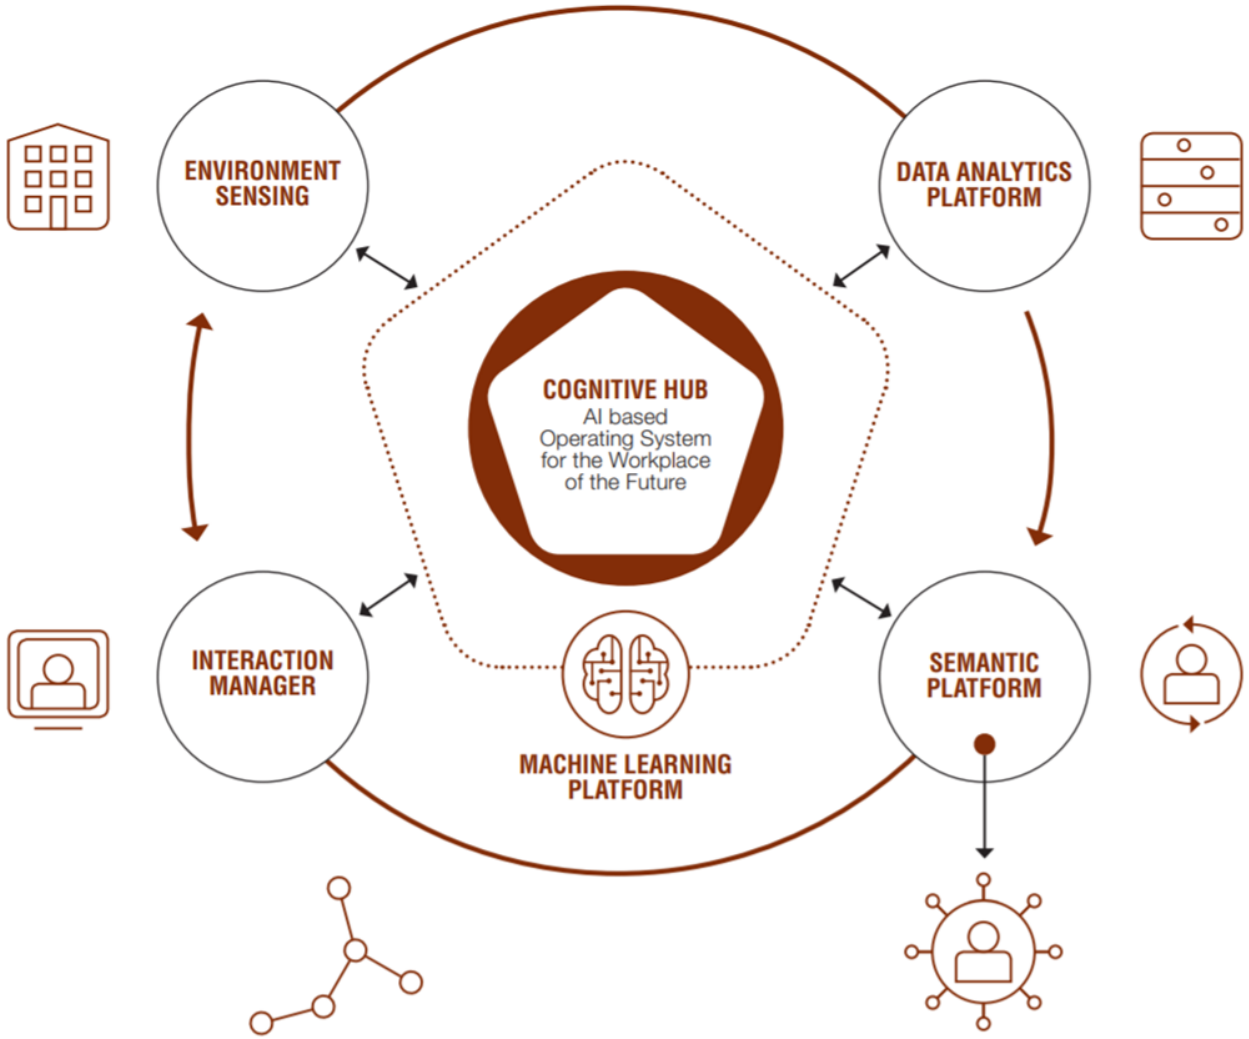

KONICA MINOLTA

What do office people do

Find the email I sent to Peter a week ago. I was about meeting in Copenhagen.

I cannot find the picture of Eiffel tower I shared with Julia a few weeks ago.

Find me a presentation of the new BMW.

I need to book a hotel in Brussels at the end of November.

Where are the charts of Adidas jacket sales?

Where is documentation to project Watchtower?

Find a cheap flight to Madrid in the second half of January.

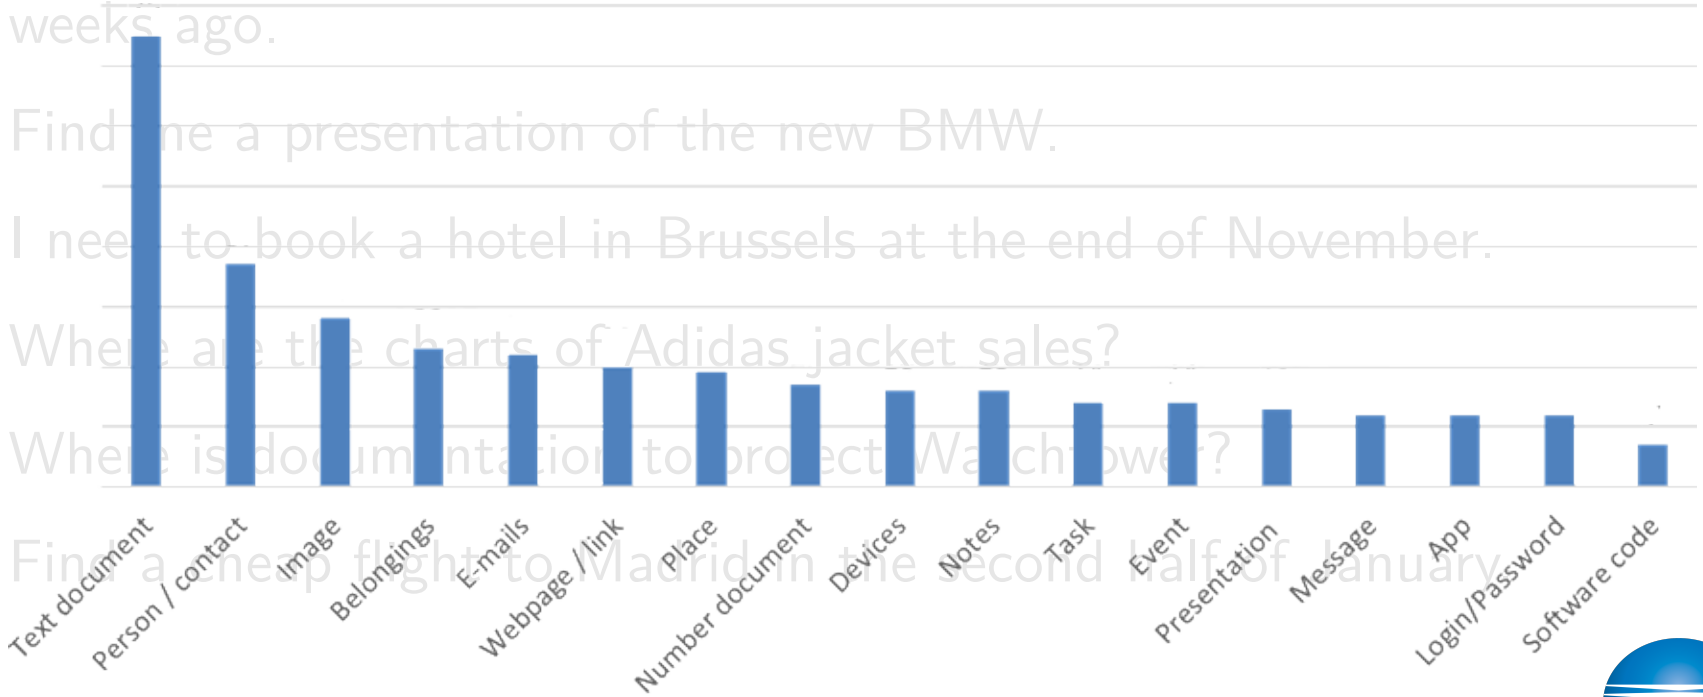

Office people search something very often

KONICA MINOLTA

1 Search in natural language

2 Recognition of user intent

3 Text analysis

4 Building the search query

5 Summary

KONICA MINOLTA

Search in natural language

- user says what she looks for
- the voice is transformed into text
- the text is understood
- the system gives feedback to the user
- the system provides search results

KONICA MINOLTA

Understanding natural language

I need the picture of Eiffel tower I shared with Julia a week ago.

find an image:
a file with MIME type `image/jpeg`

the content of the image has to be Eiffel tower:
it has to have tag `Eiffel tower`

the file was shared with Julia:
a person whose first name or last name is `Julia`

the file was shared within a time period:
2018/10/15–2018/10/17

KONICA MINOLTA

Understanding natural language

I need the picture of Eiffel tower I shared with Julia a week ago.

- recognize user's context (who, when, where)
- recognize user's intent (what)
- analyze the text
- build the search query

KONICA MINOLTA

Recognition of user intent

document: find a pdf file about Xiaomi cell phones
multimedia content: find lecture given by Mr. Smith
message: find email from Jane about the team building
event: when is my next presentation?
task: what my boss asked me to do yesterday?
person: find a Photoshop expert
organization: which university works with us on this project?
service: who can repair air condition in the Brno office?

Recognition of user intent

machine learning model with features:

- words
- part of speech
- word position (boost nouns at the beginning of the sentence)

successful in distinguishing intents in sentences such as:

find document with picture of a video camera
I need video about how to take pictures
where is a picture of a document

KONICA MINOLTA

Text analysis

semantic role labeling $\rightarrow$ predicate + arguments

recognition of entities in the arguments

KONICA MINOLTA

Semantic role labeling

I need the picture of Eiffel tower I shared with Julia a week ago.

- predicate: need(doer, object)
- predicate: share(doer, object, involved person, time)
- arg0 (doer): I
- arg1 (what is needed): picture of Eiffel tower
- arg1 (what is shared): picture of Eiffel tower
- arg2 (involved person): Julia
- time: a week ago

KONICA MINOLTA

Semantic role labeling as a BIO tagging problem

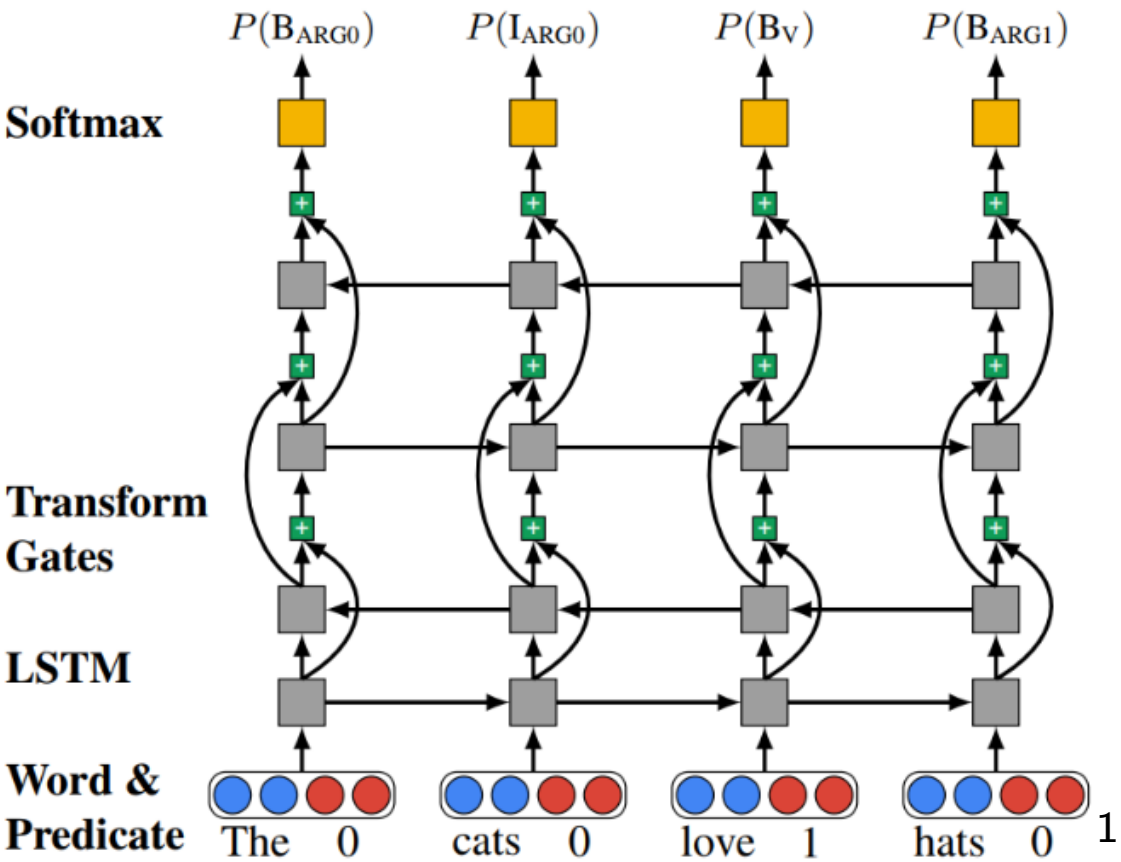

BIO tagging [Carpenter, 2009]

The	B_{ARG0}
cats	I_{ARG0}
love	B_V
hats	B_{ARG1}
.	O

[1] https://github.com/luheng/deep_srl

KONICA MINOLTA

Named entity recognition (NER)

I need the picture of Eiffel tower I shared with Julia a week ago.

Eiffel tower	LOCATION
Julia	PERSON
a week ago	DATE

state-of-the-art: Stanford NER[2]
based on
conditional random fields
[Sutton and McCallum, 2012]

- sequence boundaries (one word vs. multi-word expressions):
  - Eiffel tower
  - Eiffel
- correct entity type
  - European Union (LOCATION vs. ORGANIZATION)
  - Jackson (LOCATION vs. PERSON)
  - Ferrari (PERSON vs. PRODUCT)
  - Google (ORGANIZATION vs. PRODUCT)

[2] `https://nlp.stanford.edu/software/CRF-NER.shtml`

KONICA MINOLTA

Building the search query

interpret entities (named and unnamed)

interpret relations between entities

KONICA MINOLTA

Interpretation of the entities

I need the picture of Eiffel tower I shared with Julia a week ago.

a week ago

- depends on the time the query was said
- depends on the individual

2018/10/15–2018/10/17

picture = the search intent, i.e. a multimedia file

file of type picture

KONICA MINOLTA

Interpretation of the relations

I need the picture of Eiffel tower I shared with Julia a week ago.

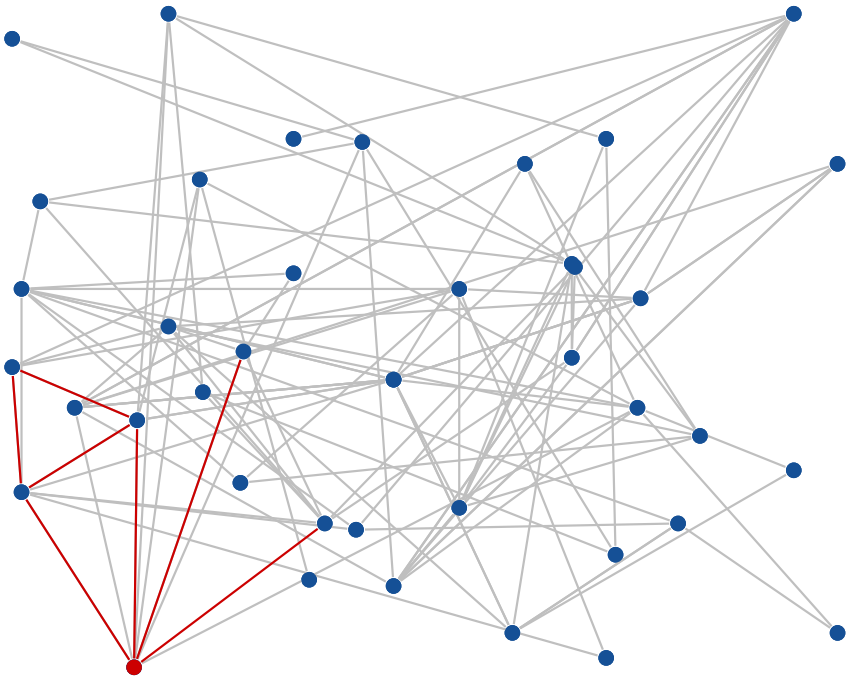

Cognitive Hub knowledge base

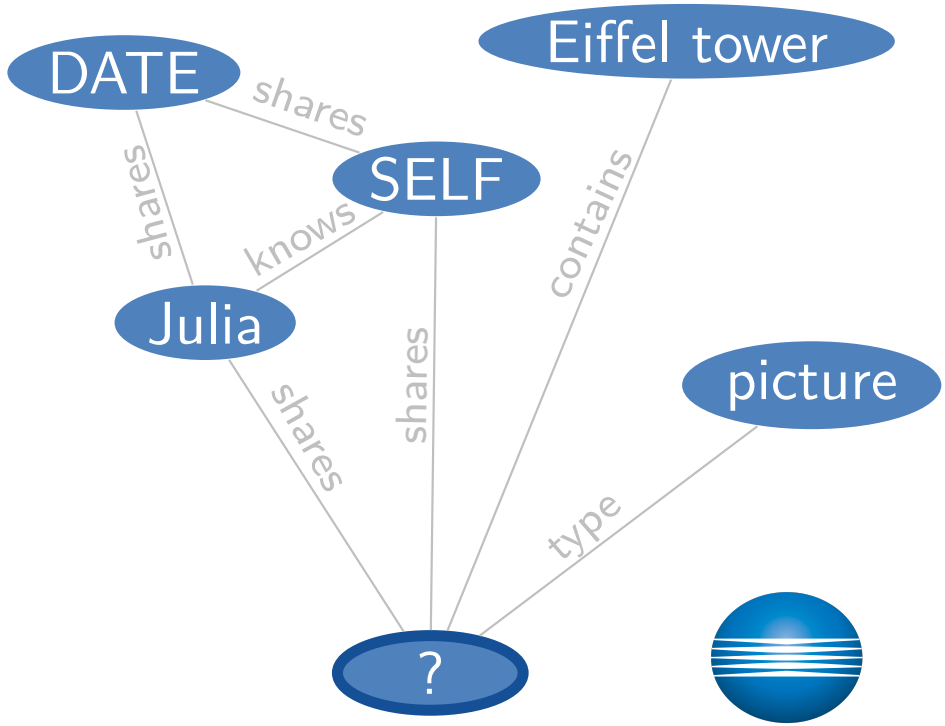

KONICA MINOLTA

Kimi the assistant

I need the picture of Eiffel tower I shared with Julia a week ago.

Search for	type	image
Contains	text	
	image	Eiffel tower
Relates to	people	Julia
	time	2018/10/15-2018/10/17
	source	

Summary

- the goal is to understand queries in natural language
- combination of ML + knowledge-based methods
- machine learning used for
  - recognition of user intent
  - semantic role labeling
  - named entity recognition
- knowledge-based methods used for
  - interpretation of the entities
  - interpretation of the relations
- search text is translated to graph query

What next? Natural language interfaces for other office tasks?


Carpenter, B. (2009).
Coding Chunkers as Taggers: IO, BIO, BMEWO, and BMEWO+.


He, L., Lee, K., Lewis, M., and Zettlemoyer, L. (2017).
Deep semantic role labeling: What works and what's next.
In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.


Sutton, C. and McCallum, A. (2012).
An introduction to conditional random fields.
Found. Trends Mach. Learn., 4(4):267–373.

KONICA MINOLTA