

# Abstracts hands-on session

Please select one option at registration on Monday.

## Option 1

Title: Trying & Checklisting: Bias and Knowledge in NLP systems

Tutor: Gosse Minnema (RuG)

### Abstract:

In this workshop, you will get your first hands-on experience with state-of-the-art systems from different areas of NLP (from machine translation to chatbots and text understanding). The goal is to learn how to approach these systems in an analytical way: modern models often give impressive results, but also have limitations in what they can do, as well as biases. In the first part of the workshop, you will do experiments with several "plug-and-play" systems available online, to see how well they work and if you can spot problems. In the second part, we'll take a more systematic approach and use the CheckList system (Ribeiro et al., ACL 2020) to debug a sentiment analysis model.

### Requirements:

Participants should have general awareness of what language technology is about. This hands-on session is especially well suited for LCT students who will now be starting their Y1. No or only very basic programming experience is needed. (In the second part, we will use a Colab notebook, but only minimal code writing/editing is required). Participants are expected to bring their own laptop. In case you cannot bring your own laptop, you will work in a group with others.

## Option 2

Title: Coreference resolution with SOTA models

Tutors: Tatiana Anikina, Natalia Skachkova (DFKI)

### Abstract:

Coreference resolution is a task of linking all expressions (mentions) in the data that refer to the same entity. It is important for many natural language processing (NLP) tasks, such as information extraction, summarization, question answering, etc. Our coreference resolution workshop offers an opportunity to make yourself familiar with coref-hoi (Xu and Choi, 2020) - one of the SOTA span-based coreference resolution model. We will show how to preprocess data, use an already pre-trained model, evaluate and visualize the results. The participants will also learn the inner workings of the span-based approach, and how mention representations can be modified with additional features. The workshop will be practice-oriented, and assumes the participants performing all the steps on their laptops.

### Requirements:

It is expected that the participants have basic knowledge of word embeddings, Pytorch and transformers.

In order to be able to perform the tasks in Google Colab, a Google account and a laptop are needed. In case you cannot bring your own laptop, you will work in a group with others.

## Option 3

Title: Language-2-Pose Generation

Tutor: Cennet Oguz (DFKI)

### Abstract:

“How can we animate 3D-characters from a movie script or move robots by simply telling them what we would like them to do?” “How unstructured and complex can we make a sentence and still generate plausible movements from it?” These are questions that need to be answered in the long-run, as the field is still in its infancy. Therefore, we here present Language-2-Pose session for generating compositional actions, which handles complex input sentences. We will practice data preprocessing, use a pre-trained model, and evaluate and visualize the results. The participants will also experiment with language analysis by using a BERT model. The workshop will be practice-oriented.

### Requirements:

It is expected that the participants have basic knowledge of word embeddings and Pytorch. The experiments require a GPU. Participants who have a GPU laptop can experiment on their laptops, others will follow an interactive presentation by the tutor.

## Option 4

Title: Recognizing taxonomic relations using language models

Tutors: Esther Ploeger, Wessel G. Poelman (RuG)

### Abstract:

The aim of this workshop is to experiment with data generation techniques and to evaluate models and their output. In particular, to gain a better understanding of the ability of language models to recognize taxonomic relations between two words. The main task is a binary classification problem where a model has to predict if a given sentence is valid with regards to the taxonomic word relation in them, for example:

- I like trees, and in particular birches. 1 (valid)
- I like oaks, and in particular trees. 0 (not valid)

You will get access to pre-trained models that have been trained on data from various automatic data generation techniques. Your task is to perform error analyses on these models and the data. This could include many aspects: looking at possible shortcomings of the training data, 'probing' the models to see what their strengths and weaknesses are, evaluating the model with your own test set, etc. You can be creative here! Step by step details to get started will be provided in a Python notebook. At the end, you should present your findings (interesting output, suggestions for improvements, what worked and what did not work etc.).

## Requirements:

- Python programming knowledge (we are working with notebooks in Google Colab, so basically no setup needed, just a Google account)
- Some knowledge of data processing (going through csv's, storing results, we will include some examples as well)
- Some knowledge of NLP techniques. We are using transformer-based models, no deep knowledge is needed, but the basics are handy. If needed, we can include a brief intro of this as well.

In case you cannot bring your own laptop, you will work in a group with others.