

Module 3

Chatbots – Natural Language Processing

"We no longer teach people how to communicate with systems, we are teaching systems to communicate with people."



About the Module

In this module, different types of chatbots are illustrated by applications. The students learn what chatbots are and how they work. Questions are asked, for example how programmers manage to make chatbots appear "human" or "intelligent" or why understanding human language is actually not that easy.

Objectives

Students will be able to...

...can explain what chatbots are and how they work

...can name applications that use chatbots

...describe why conversations with chatbots sometimes go wrong

...describe what the turing test is and are able to try it for themselves on a small scale

...describe an overview of what natural language processing is and how it works

Agenda

Time	Content
10 min	Analysis of failed chatbot conversations
10 min	Introduction to Chatbots
30 min	Testing the limits of a chatbot
30 min	Build Paper Chatbots (Click-Structure)
30 min	Build Paper NLP Chatbots

What are chatbots?

In this module, the students are thrown directly into the topic. At the beginning, conversations between a human user and a chatbot are shown, where the conversation did not succeed.

Through this first confrontation, the students are activated for this topic and can roughly imagine what can be understood by the term chatbots.

For most of these example-conversations, the chatbot's predefined word recognition is not sufficient enough for it to respond adequately to the user. The constant non-recognition frustrates users and chatbots are less well accepted by people.

With the help of these first few examples, the students can see how complex human language can be and how many things have to be taken into account when developing a chatbot.

Material

-  NLP - Introduction.pdf

Introduction

After this first short confrontation and discussion of the course of the conversation, the students should work on further questions in the plenum. The main thing here is personal experience and personal access to chatbots. The students should collect which chatbots they know and which they use themselves (do they use any at all?). The teacher can present chatbots from the different areas as input. Furthermore, it is about clarifying important questions about chatbots. What exactly are chatbots and what problems could arise with them?

To go deeper, further socially and ethically interesting questions can be worked on, which revolve around whether chatbots are a good alternative for lonely and depressed people and whether it should be obvious that the conversation partner is non-human.

What is a Chatbot?

The term "chatbot" is derived from "to chat" and "robot", i.e. a robot that can talk to you. It can answer questions automatically, without further human help, search for information on the Internet and much more. So that chatbots doesn't look so "artificial", companies sometimes use avatars, which are small images or animations, to give the chatbots an appearance. Many chatbots today use a variety of **AI** applications to make conversations run as smoothly as possible. This allows bots to store information and use it to teach themselves things. But still, if chatbots do not know what to do, they may pass the user on to a human colleague.

Which kind of chatbots are there?

Chatbots come in many different forms. Mainly chatbots are developed for reasons of efficiency, to get information quickly (e.g. when will the new kitchen cabinets be available from Ikea? Or can I cancel my order from Lieferando?). Some can wire money, others can schedule doctor appointments for you, or order food for you. And new features are constantly being added. Here are more examples of chatbots:

Commercial Bots

Chatbots replace human interactions, especially in customer service. This type of bot is intended to help customers when shopping online. Due to their 24/7 service and their quick answers, they are particularly suitable in this area. The bot helps to suggest suitable gift ideas or checks whether certain items are still in stock.

Therapeutic Bots

e.g. Woe-Bot or ELIZA. Woe-Bot, a distress chatbot, was designed to support undergraduate and graduate students who are at risk of developing depression or anxiety disorders. ELIZA, on the other hand, is considered one of the first chatbots (1966) in the history of computer science and imitates a psychotherapist

Social Bots

Other chatbots are designed for social interaction. Bots such as Mitsuku and Xiaoice are intended to satisfy human needs for entertainment and social exchange. It is precisely here that countless challenges stand out: In order to create real interpersonal connections, the conversations between humans and **AI** must be grammatically correct, topic-related, personalized, react to the mood (happy, sad, bored), reflect the tone (interrogative, compelling, declarative) and be potentially humorous, to name just a few requirements. **"When a person tells the chatbot about a sad event, such as a hip fracture of their neighbour, the chatbot should respond with compassion, sadness and maybe surprise, but definitely not joy."** Nevertheless, the developers also attach importance to the fact that the users are still aware that this is an artificial intelligence and not a human being. (Why could that matter?)

Social Media Bots (e.g. on Instagram and Twitter)

These bots enable fast growth of followers and likes by reacting to previously set parameters (e.g., specific hashtags, specific locations, followers from other accounts) with a comment or "like". Although this type of bot is very, very common, the interaction that is common with other chatbots does not take place. Especially on twitter there are some interesting social media bots. Well-known Twitter bots are, for example, @DearAssistant, which answers simple questions, @DeepDrumpf, which tries to imitate Donald Trump's language style in his messages, or @Pentametron, which tweets which. accidentally written in iambic pentameter retweeted.

Are chatbots dangerous?

In themselves, chatbots are harmless, even if that comes across differently in some science fiction films (comparing narrow and wide **AI**). However, with a large amount of stored data there is always a risk that someone will steal this information and use it for criminal purposes. Feeding in the wrong data can also result in racist or discriminatory bots, such as the Twitter bot *tay* (cf. bias errors in the ethics module).

About Social Bots

In particular, however, loneliness seems to be an increasingly present phenomenon today. More and more people prefer to communicate digitally than to talk in real life. Chatbots are mentioned as an extraordinary example. Bots like Mitsuku or Xiaoice talk to millions of users and are able to incorporate information from previous conversations. But communication with chatbots should not only be presented negatively. New technologies flowing into chatbots can help people out of loneliness. Especially in current times, when people find themselves in quarantine again and again, but also older people or those who don't know who to talk to about problems and worries, chatbots can help them find their way out of loneliness and against mental health problems such as depression or help anxiety disorders. But can these technologies actually replace human contact?

"Mitsuku doesn't pretend to be able to replace a real person, but she's always available if anyone needs her, instead of talking to the four walls," says Steve Worswick, the developer of Mitsuku. These chatbots differ from voice assistants like Siri and Alexa in that they respond to user input in a friendly and empathetic manner.

A well-known example is the Chinese chatbot Xiaoice, which has over 660 million users around the world and responds to messages in a sensitive, almost human way. Users of the bot let Xiaoice cheer them up, tell jokes or write all their worries away while the bot listens carefully.

Another extremely exciting project is the App Replika. Replika is an artificial intelligence trying to create a digital copy of your personality – a bot trying to chat like you. *"Replika writes with you. Asks questions. take care of you Wants to be the friend you've always wanted."*

Material

-  NLP – Introduction.pdf

References

1. Replika

2. Xiaoice Vs. Tay: Two A.I. Chatbots, Two Different Outcomes/

3. Woebot Health

Let's talk about Chatbots!

Why are chatbots used? What are their advantages? Discuss in groups!

Have you already come into contact with chatbots? If so, for what purpose?

What problems or dangers could there be with chatbots?

Possible further questions:

Chatbots are also used with lonely people or people with mental health problems. Do you think that makes sense or is well accepted by those affected? What problems could arise when lonely people interact more (or only) with chatbots?

With some conversations you don't know exactly whether you're talking to a human being or a chatbot. Should chatbots be clearly marked as such?

Testing The Linguistic Limits Of Chatbots

In this phase of the module, the students should talk directly to a chatbot and find out for themselves what chatbots are capable of and what their limitations are.

Mitsuku (nickname "Kuki") would be particularly interesting for this exercise. This application has already won the Loebner Prize five times and is therefore good at providing authentic answers.

In order to use Kuki, however, a free account must be created, which can also be used to save the conversation histories for future reference. Furthermore, Kuki is unfortunately only available in English and answers very evasively in other languages, which means that there is no point in the students doing complete testing in other languages (but of course it would be exciting to let the students try that too).

If it makes no sense to create accounts for the lesson or there are no resources available for testing with Mitsuku, a paper exercise can be done as an alternative, in which the students can evaluate excerpts from the test talks for the Loebner Prize themselves. The ratings of the jury can be shown later and the results of the students can be discussed.

About the Loebner Prize

The Loebner Prize is awarded to programs that pass the Turing test. There are three categories for this award: Bronze Medal: for the program that proves to be the most "human-like" (awarded annually), Silver Medal: the program passes the written Turing test, and Gold Medal: should the program pass the total Turing test. Pass the test in which multimedia content such as music, speech, images and videos must also be processed. So far, no program has passed the Turing test. So far, only the bronze medal has ever been awarded.

About the Turing Test



With the conversation between the students and Mitsuku, a very simplified version of the **Turing test** is carried out.

A test developed by Alan Turing in 1950 to determine how well a program can mimic a human's language. In the Turing test, a person must be able to determine several times

without error whether an answer to a question was given by a computer or by another human. If the person cannot do this, the computer has "passed" the test.

In our case, the students already know that they will be conversing with a chatbot and, from this perspective, will try to find out how humanely a chatbot can answer and make initial assumptions as to how chatbots can "understand" our human questions and answers.

Material

-  NLP - Loebner Prize.pdf
-  NLP - Worksheet Chatbots Introduction.pdf

References

1. <https://chat.kuki.ai/chat>

Talk With A Chatbot!

In this exercise you will take a closer look at the **chatbot Mitsuku (nickname Kuki)**: <https://chat.kuki.ai>

Stop the time!

How long does it take for you to realize that this isn't human? How did you recognize that this is a computer system?

- What happens when statements are repeated?
- What happens if you phrase the same question differently?
- How can you tell that you are talking to a machine?
- Does the chatbot always give logical answers to your questions? How does the chatbot react if it doesn't understand something?
- Do you have to write whole sentences or are single words enough for the chatbot? Does he interpret these individual words correctly?
- Provide precise examples/quotes from your conversations with the chatbot! What could be done better, so that it does appear more human?
- Should it even appear more human?

Material

-  NLP - Worksheet Talk to a Chatbot.pdf

Mitsuku and the Loebner Prize

In this shortened exercise, the students can work with statements that Mitsuku made in the Turing test by the Loebner Prize jury. This exercise can simply be printed out and filled out by the students, the jury's rating can finally be displayed with the PowerPoint or written separately on the blackboard.

Steve Worswick's Mitsuku is an apparently 20 year old chatbot and was subjected to the Turing test in a competition.

Mitsuku was asked 20 questions, some you can find in the table below.

Read the questions and their answers carefully.

In partner work, discuss how "human" the different ones are for you.

Answers are, and give points for each answer:

- 0 points - for a completely inappropriate, non-human answer
- 1 point - for a nondescript answer like "I don't know" or similar
- 2 points - for answers as one might expect them from humans

Material

-  NLP - Worksheet Mitsuku.pdf

Paper-Chatbots

In this last major work phase, the students will now try to design the theoretical basis for a chatbot themselves. The programming of an actual chatbot as such is not intended in this module, since it would go beyond the time frame and this is only about the basic understanding of how to actually build a chatbot.

In this module, two different approaches to constructing chatbots are examined in more detail. You can theoretically do both parts or, if time is short or you want to decide on a level of difficulty, you can do one of the two paper chatbot exercises.

Basically, we are dealing with two different construction options here:

1. Rule-based chatbots that offer fixed click structures so that the user cannot enter free text but "clicks through the conversation" using buttons.
2. Chatbots allow free text input and are therefore usually based on natural language processing

In these two exercises, the students will recognize that in order to be able to provide a fluent, authentic conversation, a large number of eventualities must be covered.

Creating a Flowchart-Based Chatbot

In this exercise, students develop a flowchart-based chatbot using poster paper, pens, and post-its. It would also be possible to work digitally with mind map tools or websites such as **Canva** or **Miro** .

Decision tree chatbots are the most unchatty kind of chatbots and are preprogrammed to follow a sequence, which can be very simple or complex. This chatbot works using pre-selected widgets with button options. It allows you to get creative with your chatbot's text and display options, but your user is expected to choose between these options that you define. Companies use them because they're cheaper to build, quicker to deploy and can still be useful, entertaining and educational.

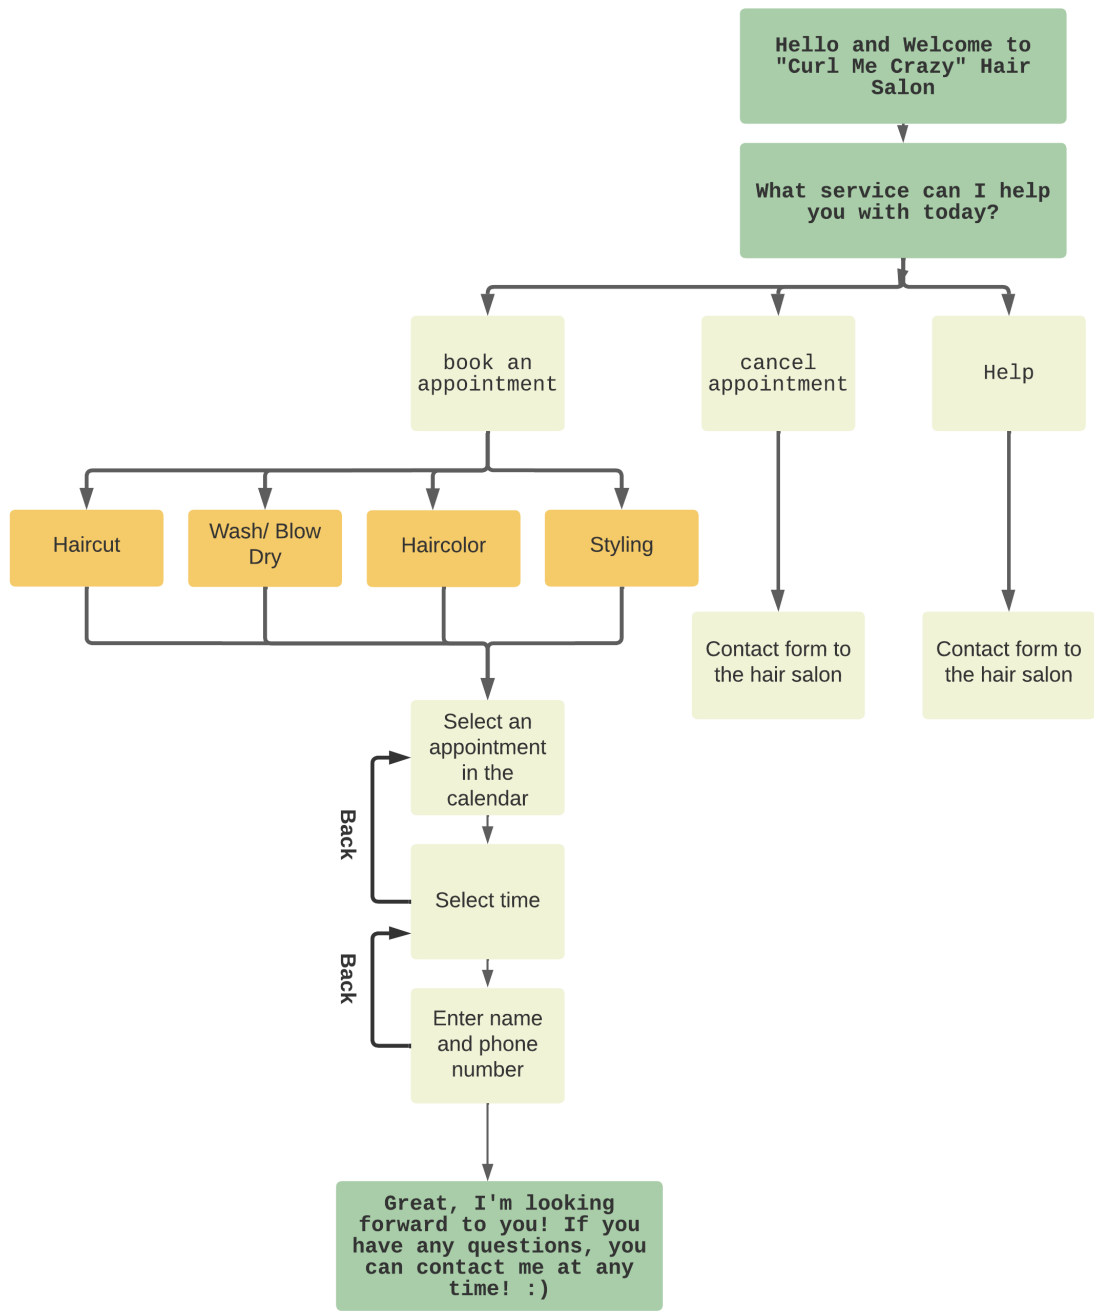
At the beginning of this exercise, the students should think about the purpose of the chatbot (should you be able to use it to book hairdresser appointments, ask about the weather or order pizza...).

- Look for a topic on which the chatbot should specialize. Do you want to be able to book a hairdressing appointment with it? Do you want to order pizza? Should he take on the role of therapist? Should he make initial diagnoses for sick people? Or do you want to calculate and manage your budget?
- Once you have decided on a type of chatbot, you must think about what tasks this chatbot should take on. For example, a chatbot for a hair salon is designed to make appointments. What information does the chatbot need to ask for so that an appointment can be booked? (Date of the appointment, personal data such as name and telephone number, which kind of hair-service). Furthermore, functions such as cancelling an appointment or contacting a person should also be considered.
- In the next step you must build a flowchart where every information will be asked. These kind of chatbots are just like interactive flowcharts.

How to Make a Flowchart?

1. Start with pen and paper and draw your flowchart. You can use sticky notes on a whiteboard or draw diagrams with a pencil on a big paper. This allows you to move or exchange individual points.



2. Make sure that all decision-making processes make sense, and you don't end up in a loop that the user can't get out of.
3. Avoid flowchart dead-ends! The user should reach a finishing node at the end.
4. Keep your flowcharts as simple as possible but still give your users enough choices that make sense.
5. Don't forget to include a closing phrase as well, so the user doesn't get stuck in a dead end. However, the final formula should be open enough so that the user knows that he can ask questions at any time. Examples can be: "I'm here if you need me. Don't be shy!" or "I'll be gone, but let me know if you need anything :)"



If you have more time when building your chatbot, you can also consider the personality of your chatbot. The character of the chatbot does not have to be highly complex. An authentic, humorous chatbot can strengthen the trust of users and is also more fun to use. Is your chatbot a quick-witted, middle-aged woman, or a small, shy boy, or a serious older man? Keep in mind which company your chatbot represents so that the essence of your chatbot fits your fictional brand.

When your chatbot is finished, exchange ideas with another group and play through the chatbot! Does everything make sense or are there problems?

Material

-  NLP - Chatbot with Click-Structure.pdf
-  NLP - Worksheet Flowchartchatbot.pdf

Creating a NLP-Based Chatbot

In this exercise, students develop a paper chatbot that reacts to keywords like **NLP** models and thereby gives correct answers. The students receive a table with possible input and output words and have to respond appropriately to the given input. The students will soon find out that a large field of possibilities has to be covered and that spare keys have to come into play so that a conversation can run smoothly.

How does a NLP-Chatbot work?

When using a Natural Language Processing (**NLP**) Chatbot you have to enter a verbal input using the keyboard or saying it out loud. The bot analyses your words and turns them into information. It consists of understanding the human language, **Natural Language Understanding (NLU)** and creating the language **Natural Language Generation (NLG)**.

A chatbot sees what letters you write but has **no knowledge of the meaning of the words**. The meaning of the phrase "I want to make an appointment" is perfectly clear to you, but to a computer it's just a list of letters. First, it means as much to him as "mis sdaijhw wek" does to you.

Let's say we have a chatbot that books and cancels appointments at a hair salon. The chatbot asks "How can I help you at our Hair Salon?" and the user replies "I want to book an appointment.". So, this response is the input that the chatbot receives. The chatbot now starts a **step-by-step process** (also called a pipeline) to recognize the **necessary structures** so that the program can give the correct answer.

Tokenization

The first structure that is recognized is the **beginnings and endings of the sentences and words**. It is always clear to us what the words are, but to the computer it is not. The separation into individual words is called **tokenization**. The bot needs to "know" that punctuation marks are not part of the word, so the word is "appointment" and not "appointment." He must also know that not every point is the end of a sentence and, if necessary, resolve abbreviations.

Lemmatization

The second thing is to find **the basic form of the words**. We humans do not have to determine the basic form of every word in order to understand it. However, it is very useful for a computer to reduce the number of different word forms. Because this means that the computer recognizes that it can treat "appointment" and "appointments" equally in many cases. For example, the sentence "Is it possible to book any appointments at the moment?" has the same basic forms as our example sentence for the most important words. Accordingly, the computer can recognize that

these 2 sentences should probably be treated the same. The determination of the basic form is called **lemmatization**.

Part of Speech (Pos) Tagging

The third structure is the parts of speech. As humans, we don't care about the part of speech either. However, this is helpful for the computer, since **nouns and verbs** are usually more important for the rough meaning of a sentence than other parts of speech. Determining the parts of speech also helps us in the next step. The determination of the parts of speech is also called **part of speech (Pos) tagging**.

Syntax Analysis

The **recognition** of subject, object and verb, as well as other **grammatical constructs**, is the second to last important structure. The computer performs a **syntax analysis** using a parser. He takes the three structures already created and creates a model that analyses **the dependencies of the words** within a sentence. In the simplest case we define subject, object and meaningful verb to cover the most important parts of the sentence

Semantic Analysis

The final structure is now about **the meaning of the input**. As I said, the computer does not know what an "appointment" is or what "to book" means. But to give a correct answer we only need the following information: Does he want to book or to cancel an appointment or does he want something else? (Intent) What kind of thing should be booked? (Parameter/Entity) This is the representation of the meaning for us. For this **semantic analysis**, the chatbot needs a list of possible things (e.g. appointment) about which we can give answers and examples for each intention (book, cancel).

The keywords are compared through this list and the chatbot gives the right answer that matches the keywords.

Language assistants such as Alexa and Google Assistant also work in this way. The speech recognition is carried out there first. This converts the microphone signal into a character string. After that, the same creation of structures is performed.

What happens when AIs book hairdressing appointments?

In 2018, Google's "Duplex" application was the first to authentically arrange appointments at the hairdresser's or in a restaurant. Duplex's voice sounds so convincingly human that many initially thought the videos of these conversations were fake.

Google boss Sundar Pichai explained that these recordings were real but only the very least conversations went so perfectly like these. Most conversations with the **AI** still fail miserably.

A video of one of the successful phone calls can be found in the references.

References

1. (DE) Wie funktioniert Natural Language Processing in Chatbots?
2. (ENG) Google Assistant making a phone call

Paper-Exercise for Creating and Testing a Chatbot

Have a conversation with the paper chatbot!



As you have heard before, it takes a lot of steps that we humans take for granted to understand a statement.

1. One student is the user, another student takes on the role of the chatbot
2. The student who plays the human user wants to make an appointment at the hairdresser's and writes his request on a piece of paper
3. The human chatbot now goes through this sentence and picks out the important keywords by comparing the words with its table.
4. The chatbot can answer using the keywords it finds on the table.
5. If the human chatbot does not find any keywords on the table, an appropriate answer must be selected.

Try to solve the following tasks:

- Try to book an appointment at the hairdresser with this table! Do you manage it or do you have problems? How do you have to complete the table so that you can have a smooth conversation?
- What happens if a negation is made by a customer? (e.g., I do NOT want any more appointments)
- Imagine a customer who has a complaint because his appointment was lost, or his hair was cut wrongly and is now upset. How could the bot know that the customer is angry? How should the chatbot react to insults?
- What would have to be done to make the chatbot more authentic? How do you think he could pass a turing test?

Material

-  NLP - Chatbot.pdf
-  NLP - Worksheet Chatbot.pdf


Conclusion

In this last step within this module, language as a whole will be discussed again. By exploring existing applications but also by developing their own paper chatbots, the students have developed an understanding of how complex language can be and developing a program that can respond to it.

The human language is a challenge for artificial intelligence for sure. *"With all its ambiguities, nuances and misunderstandings, it is probably the most complex system that humans have ever developed."* That's why no one has yet succeeded in constructing a machine that simulates a credible human interlocutor.

But, as the last picture of the module presentation shows, the construct of language often does not work even between people. Models such as Schulz von Thun's four-page model can also be embedded in lessons to show how complex language can be for people and that the expectations between sender and receiver can often vary greatly.





EUROPEAN UNION

