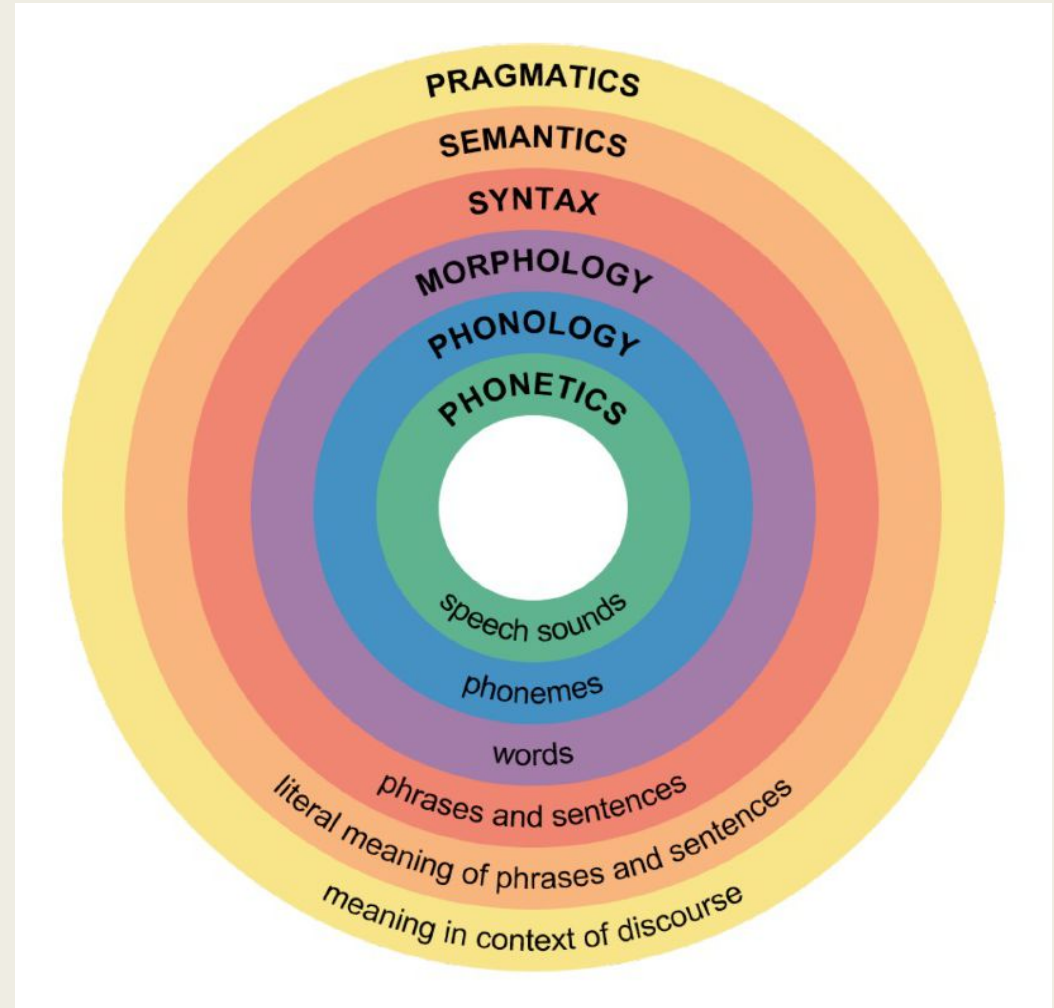
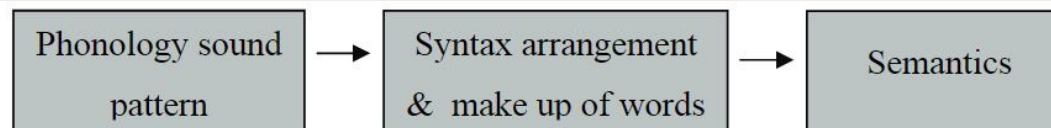


# SEMANTICS, PRAGMATICS & COMPUTATIONAL SEMANTICS



# Semantics

- A branch of linguistics for the study of meaning
  - Meaning can be studied on many levels: words, phrases, sentences, et cetera*
- **Compositionality**: how lexical items are combined to generate meanings
- **Lexicality**: the meanings of words
- **Contextuality**: interpretation of meaning in context; vagueness, connotation, presupposition, et cetera (link to **Pragmatics**)
- Interplay with other levels of language, such as **syntax** and **phonology**



**Three central components which make up a grammar: Phonology, Syntax and Semantics**

Alhaj, A. (2016) Understanding Semantics: A Textbook for Students of Linguistics and Translation. Anchor Academic Publishing.

# Pragmatics: Speech Acts

- Locutionary, illocutionary, and perlocutionary acts
- Speech Act Theory (Austin, 1962; Searle, 1968)
  - *Declaratives*: blessings, firings, baptisms, arrests, judicial acts such as sentencing
  - *Representatives*: assertions, statements, claims, hypotheses, descriptions, suggestions
  - *Expressives*: greetings, apologies, congratulations, condolences
  - *Directives*: commands, requests, challenges, invitations, orders, summons
  - *Commissives*: promises, oaths, pledges, threats, vows

# Computational Semantics

- A branch of Computational Linguistics; the study of computer-automated representations of meaning in natural language
- Manyfold applications, such as word sense disambiguation, semantic role labeling, coreference resolution
- **Lexical semantics**; study of the meanings of words
  - *Synonymy, polysemy, similarity, semantic fields, semantic frames, connotation, sentiment...*
- **Vector semantics**; word representations in vector space

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*.

# Vector Semantics

- Distributional hypothesis (Joos, 1950; Harris, 1954; Firth, 1957)
- Word vectors are generally called **word embeddings (usually dense)**
  - *Instead of needing the same words to appear in training and testing datasets → use similarity*
  - *Sparse vectors (e.g. tf-idf: the weight of a word in a document)*
  - *Dense vectors (e.g. word2vec)*
  - *These methods learn a static vector for each word/sentence/document*
  - *Cosine similarity: similarity between two vectors in space measured by the cosine of the angle between them*

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition.*

# Contextual Word Embeddings

- BERT (Devlin et al., 2019), GPT (Radford et al., 2019)
- Each word gets a different vector each time it occurs in a different context
- Bidirectional Encoder: looks at concordances on both sides of a word
- Transfer learning
  - *Pre-training: general domain*
  - *Fine-tuning: adapting to domain-specific tasks*

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*.

# Social life of semantics (and pragmatics)

- NLP models influence meaning construction and amplify existing constructs
- Amplification of **bias** (Suresh & Gutttag, 2021)

**Historical** → training data (text corpora) project real-world biases onto word embeddings

**Representation** → training data sampling under-represents some part of the population, fails to generalize well

**Measurement** → collecting or computing features and labels; a feature or label is a *proxy* to approximate some *construct* that is not directly encoded or observable

**Aggregation** → one-size-fits-all model is used for data in which there are underlying groups or types of examples that should be considered differently

**Learning** → model learning and optimization; accuracy measurement, e.g. cross-validation techniques

**Evaluation** → benchmark data used to evaluate a model does not represent the use population

**Deployment** → mismatch between the problem a model is intended to solve and the way in which it is actually used



# Chatbots & Dialogue Systems

Prof. Nitin Sawhney  
Dept. of Computer Science, Aalto University  
February 1, 2022



# Properties of Human Conversation



DIALOGUE  
STRUCTURE



TURN-TAKING



GROUNDING



END-POINT  
DETECTION

C<sub>1</sub>: ... I need to travel in May.  
A<sub>2</sub>: And, what day in May did you want to travel?  
C<sub>3</sub>: OK uh I need to be there for a meeting that's from the 12th to the 15th.  
A<sub>4</sub>: And you're flying into what city?  
C<sub>5</sub>: Seattle.  
A<sub>6</sub>: And what time would you like to leave Pittsburgh?  
C<sub>7</sub>: Uh hmm I don't think there's many options for non-stop.  
A<sub>8</sub>: Right. There's three non-stops today.  
C<sub>9</sub>: What are they?  
A<sub>10</sub>: The first one departs PGH at 10:00am arrives Seattle at 12:05 their time. The second flight departs PGH at 5:55pm, arrives Seattle at 8pm. And the last flight departs PGH at 8:15pm arrives Seattle at 10:28pm.  
C<sub>11</sub>: OK I'll take the 5ish flight on the night before on the 11th.  
A<sub>12</sub>: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air flight 115.  
C<sub>13</sub>: OK.  
A<sub>14</sub>: And you said returning on May 15th?  
C<sub>15</sub>: Uh, yeah, at the end of the day.  
A<sub>16</sub>: OK. There's #two non-stops ... #  
C<sub>17</sub>: #Act... actually #, what day of the week is the 15th?  
A<sub>18</sub>: It's a Friday.  
C<sub>19</sub>: Uh hmm. I would consider staying there an extra day til Sunday.  
A<sub>20</sub>: OK... OK. On Sunday I have ...

**Figure 24.1** Part of a phone conversation between a human travel agent (A) and human client (C). The passages framed by # in A<sub>16</sub> and C<sub>17</sub> indicate overlaps in speech.

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*.

# Properties of Human Conversation

- Speech Acts
- Adjacency pairs: A4, C5
- Sub-dialogues or side sequences: C17-A20
- Initiative (user, system, mixed)
- Inference and Implicature: A2, C3 (maxim of relevance)

C<sub>1</sub>: ... I need to travel in May.  
A<sub>2</sub>: And, what day in May did you want to travel?  
C<sub>3</sub>: OK uh I need to be there for a meeting that's from the 12th to the 15th.  
A<sub>4</sub>: And you're flying into what city?  
C<sub>5</sub>: Seattle.  
A<sub>6</sub>: And what time would you like to leave Pittsburgh?  
C<sub>7</sub>: Uh hmm I don't think there's many options for non-stop.  
A<sub>8</sub>: Right. There's three non-stops today.  
C<sub>9</sub>: What are they?  
A<sub>10</sub>: The first one departs PGH at 10:00am arrives Seattle at 12:05 their time. The second flight departs PGH at 5:55pm, arrives Seattle at 8pm. And the last flight departs PGH at 8:15pm arrives Seattle at 10:28pm.  
C<sub>11</sub>: OK I'll take the 5ish flight on the night before on the 11th.  
A<sub>12</sub>: On the 11th? OK. Departing at 5:55pm arrives Seattle at 8pm, U.S. Air flight 115.  
C<sub>13</sub>: OK.  
A<sub>14</sub>: And you said returning on May 15th?  
C<sub>15</sub>: Uh, yeah, at the end of the day.  
A<sub>16</sub>: OK. There's #two non-stops ... #  
C<sub>17</sub>: #Act... actually #, what day of the week is the 15th?  
A<sub>18</sub>: It's a Friday.  
C<sub>19</sub>: Uh hmm. I would consider staying there an extra day til Sunday.  
A<sub>20</sub>: OK... OK. On Sunday I have ...

**Figure 24.1** Part of a phone conversation between a human travel agent (A) and human client (C). The passages framed by # in A<sub>16</sub> and C<sub>17</sub> indicate overlaps in speech.

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition.*

# Rule-based Chatbots: *passing the Turing Test?*

- Rule-based chatbots programmed using rules and regular expressions (Eliza, ALICE, PARRY)
- Designed to simulate conversations with psychologists for therapeutic advice
- Uses simple pattern/transform rules and ranked keywords (ELIZA)
- Modeling mental states with affect variables for fear, anger, trust, anxiety etc. (in PARRY)
- PARRY: first rule-based chatbot to pass the Turing test! *Psychiatrists couldn't distinguish text transcripts of interviews from real people!*
- Framework of pattern/action architecture still widely used in modern chatbots today (ALICE)

WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU?  
My father  
YOUR FATHER  
You are like my father in some ways.  
WHAT RESEMBLANCE DO YOU SEE  
You are not very aggressive but I think you don't want me to notice that.  
WHAT MAKES YOU THINK I AM NOT AGGRESSIVE  
You don't argue with me.  
WHY DO YOU THINK I DON'T ARGUE WITH YOU  
You are afraid of me.  
DOES IT PLEASE YOU TO BELIEVE I'M AFRAID OF YOU  
My father is afraid of everybody.  
WHAT ELSE COMES TO MIND WHEN YOU THINK OF YOUR FATHER  
Bullies.  
DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR BOYFRIEND MADE YOU COME HERE

**Figure 24.4** Sample ELIZA dialogue from Weizenbaum (1966).

```
function ELIZA GENERATOR(user sentence) returns response  
Find the word w in sentence that has the highest keyword rank  
if w exists  
    Choose the highest ranked rule r for w that matches sentence  
    response ← Apply the transform in r to sentence  
    if w = 'my'  
        future ← Apply a transformation from the 'memory' rule list to sentence  
        Push future onto memory queue  
    else (no keyword applies)  
        either  
            response ← Apply the transform for the NONE keyword to sentence  
        or  
            response ← Pop the oldest response from the memory queue  
return(response)
```

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*.

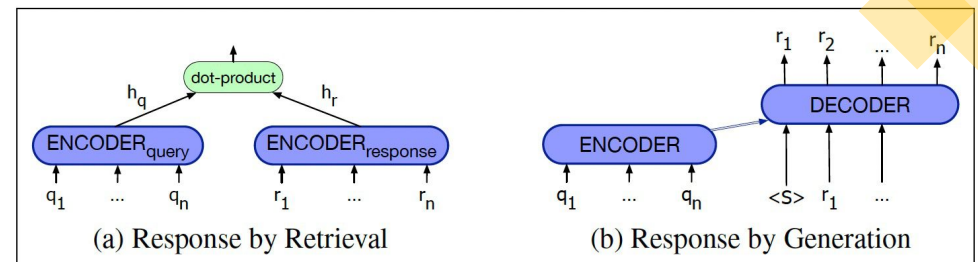
# Corpus-based Chatbots

## Training the Systems:

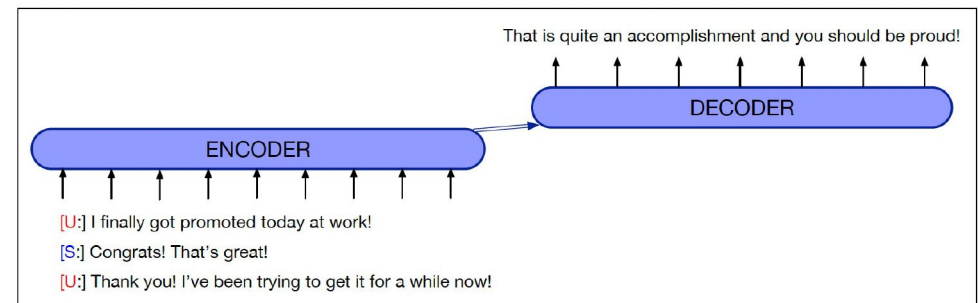
- Mining human-to-human conversations from thousands of textual transcripts of natural spoken corpora e.g. recorded telephone calls, movie dialogues or crowd-sourced datasets.
- Pre-train on large datasets from Twitter, social media, Wikipedia, books and news articles to acquire facts & narratives to support dialogues.
- Ongoing human responses expand the corpus.

## Response-Generation Methods:

- Response-generation systems using information retrieval methods and language generation algorithms for each turn of the dialogue.
- Training encoders for user queries and potential responses. Neural methods for capturing longer context and sentiment in conversations.
- Response production as encoder-decoder task – transducing from users' prior response to answer.
- Hybrid architectures combining rules and NLP models (GPT-2) to classify user intent and generate dialogues (Chirpy Cardinal).



**Figure 24.6** Two architectures for generating responses for a neural chatbot. In response by retrieval (a) we choose a response by finding the turn in the corpus whose encoding has the highest dot-product with the user's turn. In response by generation (b) we use an encoder-decoder to generate the response.



**Figure 24.7** Example of encoder-decoder for dialogue response generation; the encoder sees the entire dialogue context.

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*.



# Frame-based Dialogue Systems

- Supporting task-based dialogues with constrained domains like travel planning, calendar appointment or setting an alarm.
- **Frames** are knowledge structures representing intentions the system can extract from user interactions.
- Establish a **domain ontology of frames** with a collection of slots and semantic values.
- Goal: understand user intent using multiple frames and questions to complete slots and take action.
- How: identify domain, use production rules to invoke frames, determine user intent, and fill slots.
- Use semantic grammars, NLP and templates to generate prompts and actions.

Slots: City | Date | Time | Customer Name

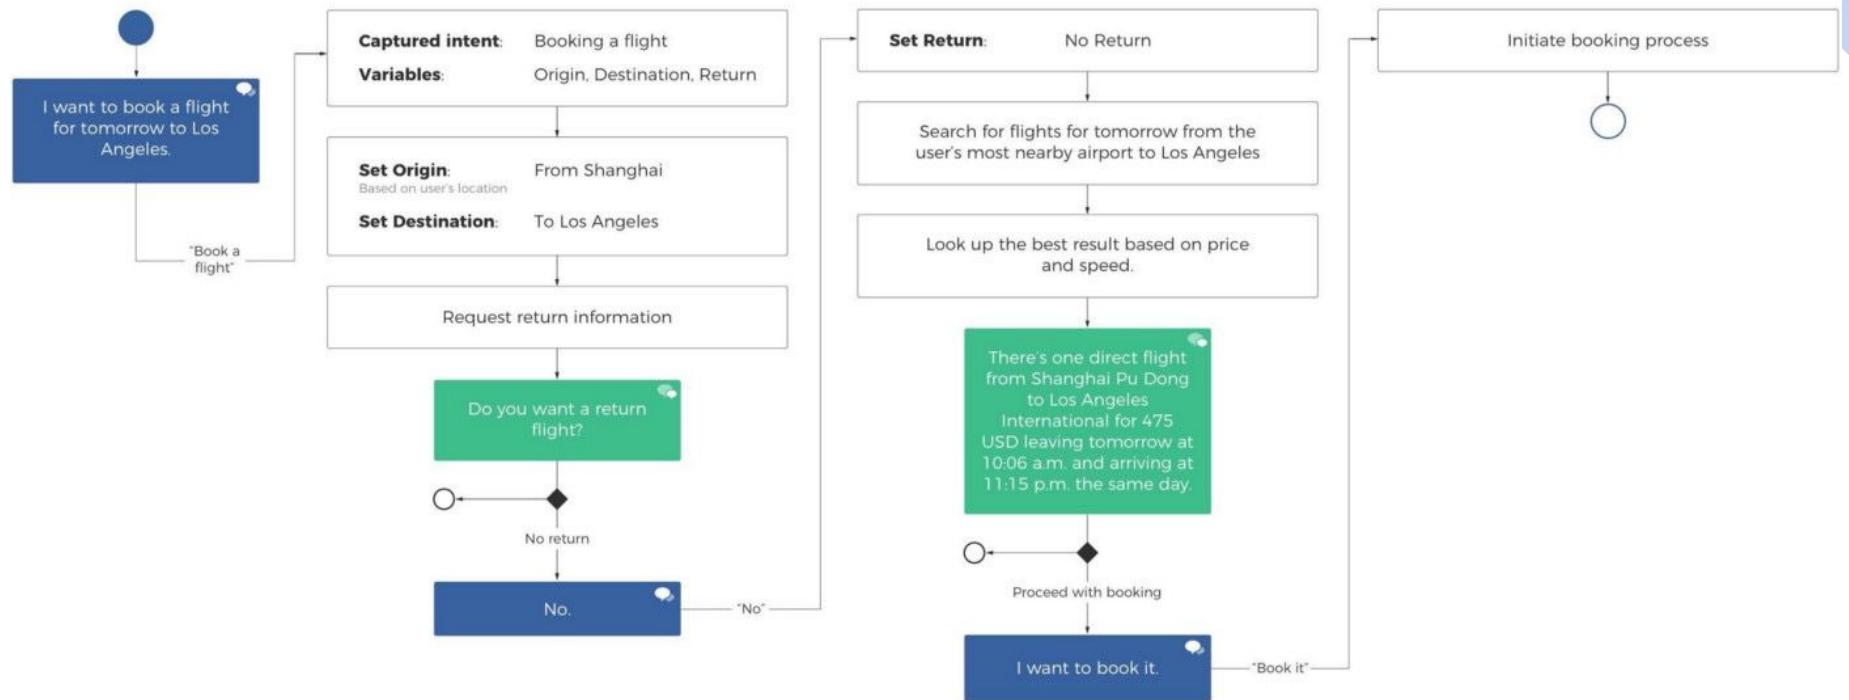
GUS: Hello. My name is GUS. I can help you plan a simple trip by air. Where do you want to go ?  
Client: I want to go to San Diego on May 28  
GUS: What time do you want to leave?  
Client: I must be in San Diego before 10 am  
GUS: Would you like P.S.A. flight 102 that arrives at 9.15 am ?  
Client: What is the next flight ?  
GUS: Air California flight 310 that leaves at 8.30am and arrives at 10.00am  
Client: I'll take the first one  
GUS: In what name should I make a reservation, please?  
Client: Dan Bobrow  
GUS: I have confirmed the following flight: P.S.A. flight 102 on Wednesday May 28 from San Jose to San Diego that leaves at 7.30 am and arrives at 9.15 am.  
GUS: What date do you want to return on ?  
Client: On Friday in the evening.  
GUS: Would you like the flight that leaves at 7.45 pm ?  
Client: That's fine.  
GUS: I have confirmed the following flight: P.S.A. flight 307 on Friday May 30 from San Diego to San Jose that leaves at 7.45 pm and arrives at 9.30 pm Thank you for calling. Goodbye

**Figure 24.9** The travel domain: A transcript of an actual dialogue with the GUS system of Bobrow et al. (1977). P.S.A. and Air California were airlines of that period.

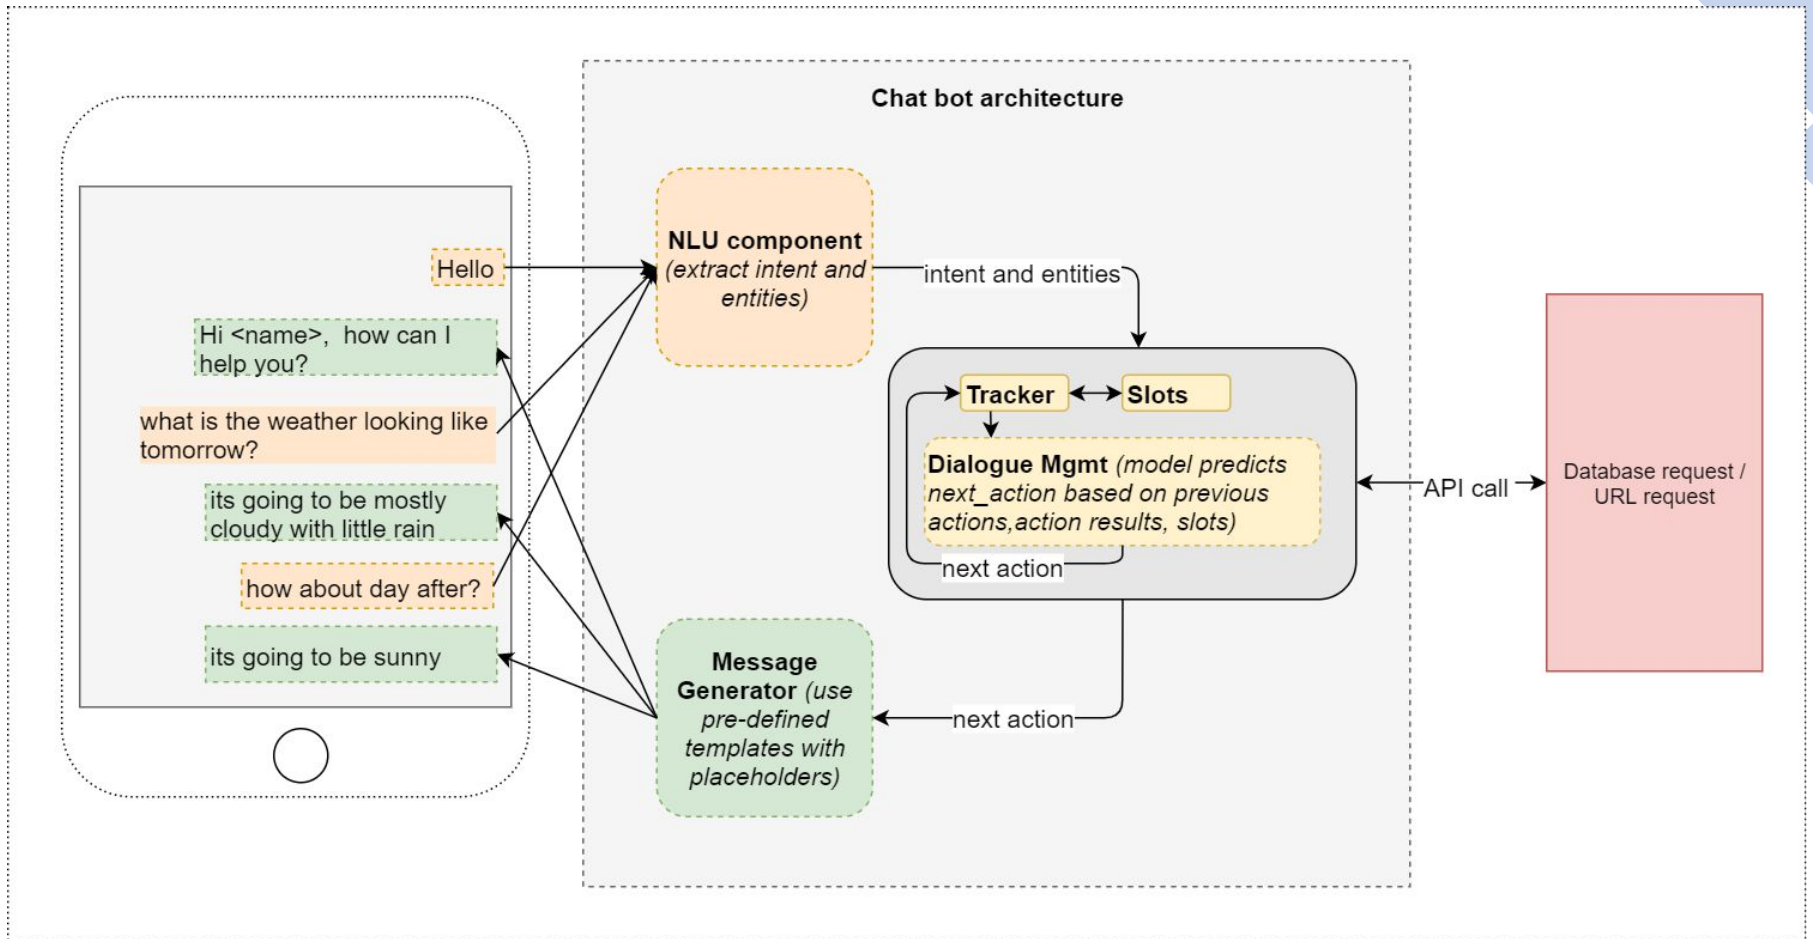
Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*.

# Chatbots & Voice Assistants

## Sample Dialog Flow

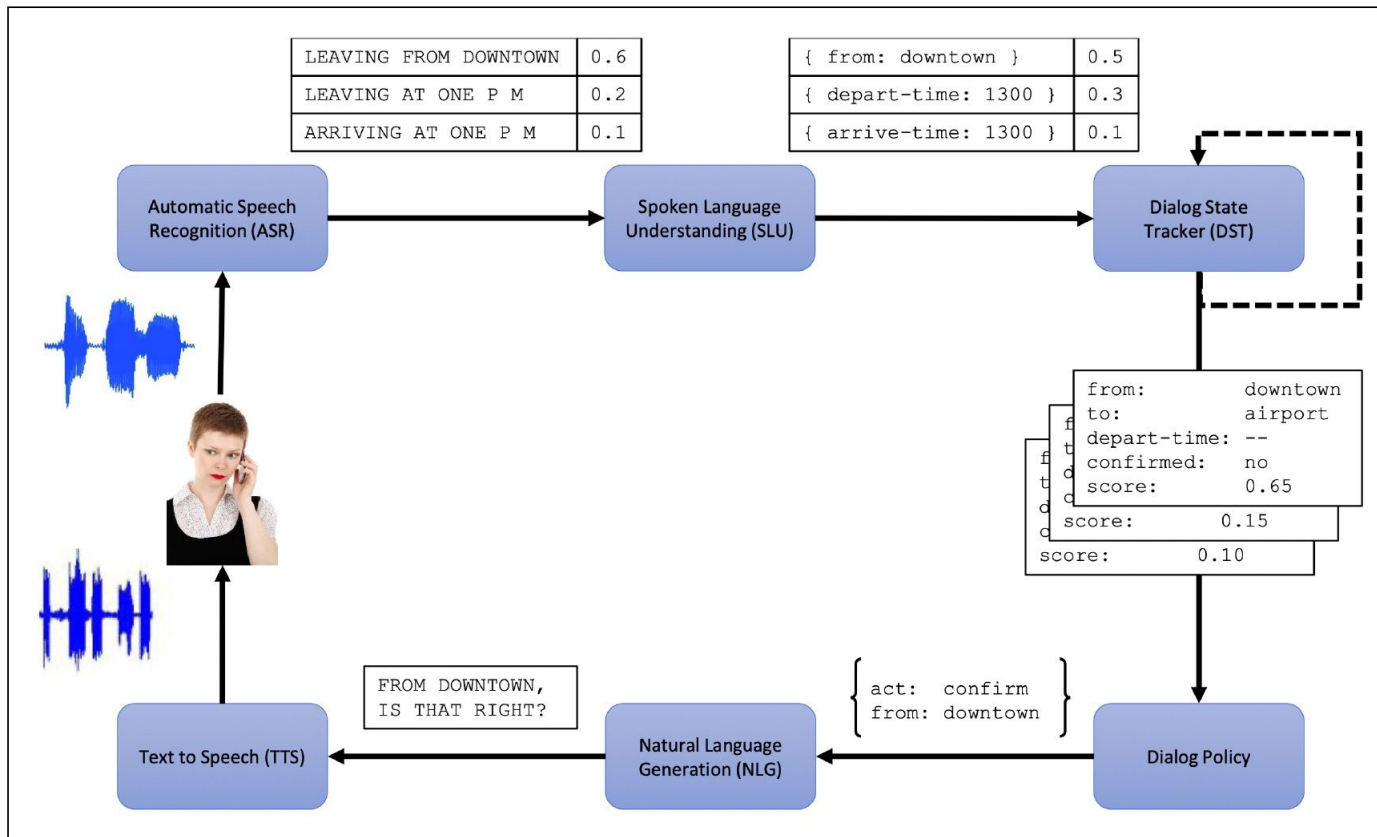


A sample dialog flow illustrating the intent, slot and overall conversation. Created by Frederik Goossens.



Conversational AI Chatbot Architecture, Ravindra Kompella

# Dialogue State Architecture



**Figure 24.12** Architecture of a dialogue-state system for task-oriented dialogue from Williams et al. (2016a).



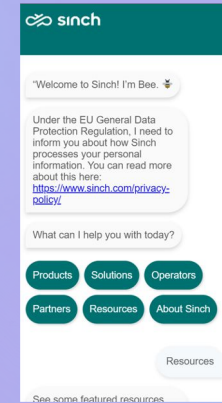
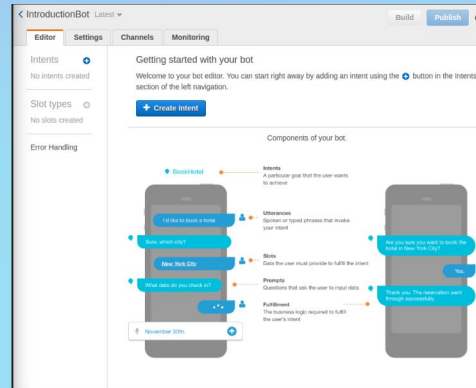
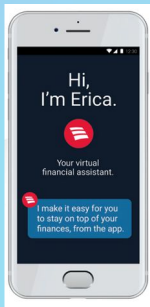


# Conversation Design



# Conversational Design

- Making experiences with interfaces more natural to what a human to human interaction could be
- Usually employed as a help tool to guide the user with tasks on a certain platform
- The closer it is to human conversation means less work for the user to understand.



# The Properties

Intent	Entity	Example utterance
CheckWeather	{ "type": "location", "entity": "Seattle" } { "type": "builtin.datetimeV2.date", "entity": "tomorrow", "resolution": "2018-05-23" }	What's the weather like in Seattle tomorrow?
CheckWeather	{ "type": "date_range", "entity": "this weekend" }	Show me the forecast for this weekend

## Building blocks

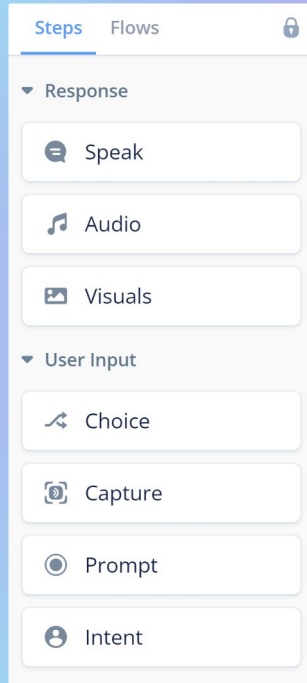
Intent

Entity/ Slots

Utterances

Prompting

Actions and Choices



## Attributes to consider

Who is the audience?

User persona

What do you want the user to feel?

What tone is appropriate?

Bot personas/ personalities

Use cases



# Scripting



Karen McCullah, Kirsten M. Smith



# Problems

Lack of language-specific knowledge

Ex: Spanish vs English syntax, Finnish

Lack of culturally-specific knowledge

Ex: Starting convo with Brazilians vs Americans

Bias

Ex: All Germans speak the same dialect

American English	British English
Take Out	Take Away
Sneakers	Trainers
Check	Bill
Trunk	Boot
Appetizer	Starter
Pitcher	Jug
Elevator	Lift
Flashlight	Torch
Cookie	Biscuit
Sedan	Saloon
Shopping Cart	Trolley
Windshield	Windscreen





# Considerations

Learning basic psychology to better understand how people may interact with an interface

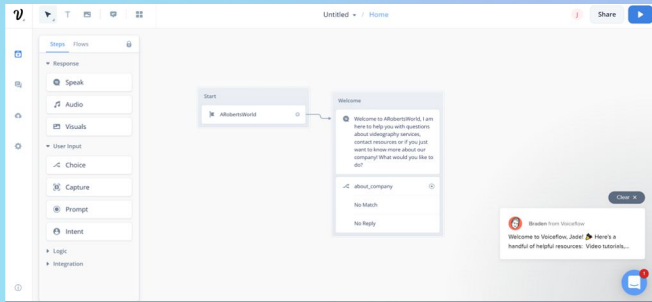
Learning linguistics and/or have copywriting experience to understand (every word matters)

Learning UI/UX design to understand design principles

Understanding your own biases and belief system

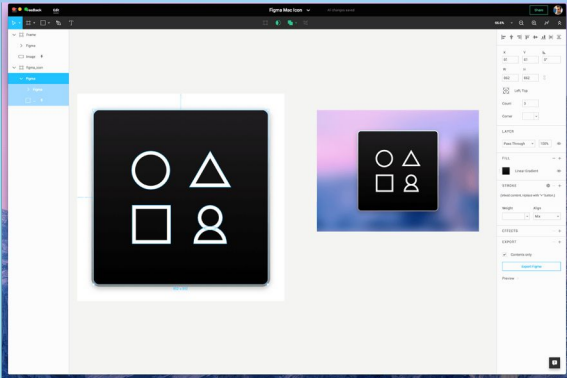
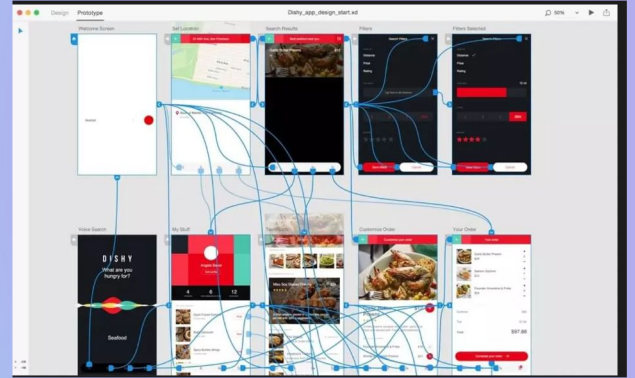


# Conversation Design and Bot Creation Tools





# Design Tools



**FIGMA**

**SKETCH**

**ADOBE XD**

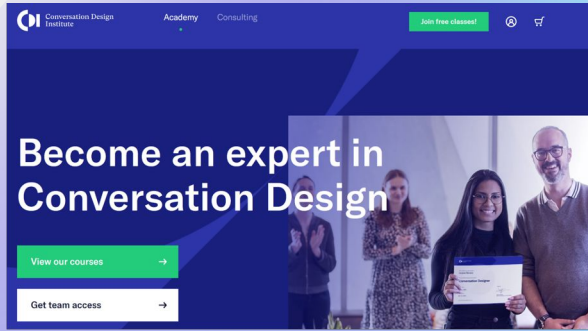


vs

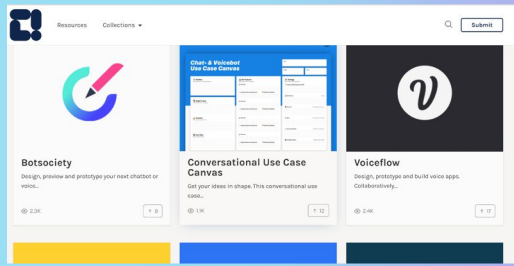


vs

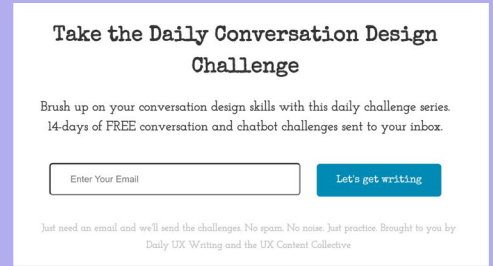




<https://www.conversationdesigninstitute.com/>

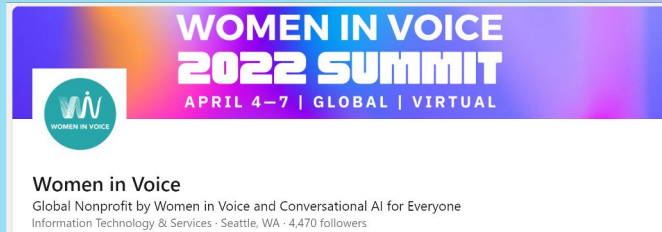


<https://cui.tools/resources/>



<https://dailyuxwriting.com/conversation-design-challenge>

<https://womeninvoicel.org/>



<https://convocollective.org/>

