1. Types of conversational agents 1) Chatbot - Intended for casual conversation with a user 2) Task-oriented agents - Engage in discussion with a user to reach one or several specific goals often (but not always) in a specific domain.
2. Types of conversational agents based on dialog representation 1) Interpretable 2) Black Box (involving deep neural networks)
3. Types of chatbots 1) Pattern-based 2) Information-retrieval 3) Generative. The first two are interpretable and the third may be an interpretable or black box system.
4. Pattern-based - Usually based on AIML (Artificial Intelligence Markup Language). AIML is an XML-like language which is used to define template matching rules for chatbots. Each user input (pattern) is associated with an output (template). An example AIML specification is given on page 9. AIML also has the ability to make redirections from one pattern to another. So various patterns can be mapped to the same template. An example is given on page 10.
5. Information-retrieval - Match the user input to the most similar sentence in a dialog corpus (dataset of conversations) and return the corresponding response. Similarity between sentences is judged with sentence embeddings. A common embedding is TF-IDF (term frequency - inverse document frequency). TF vector is a vector of same length as vocabulary with the number of occurrences of each word in the sentence normalised by the number of words in the sentence. TF gives too much emphasis to common words like "we", "the", "a", etc. The IDF for a word "w" is defined in equation 2 on page 10. The IDF for less frequent words will be more than the IDF of more frequent words. We place the IDF of each word in a vector of the same length of the vocabulary. Now given a sentence "u" we can find "q(u)" which is the elements wise product of the sentence vector with the IDF vector. The sentence can then be compared with all other sentences using the normalized dot product (cosine similarity). The most similar sentence is found in this way. TF-IDF has the disadvantages that it does not consider order of words, synonyms and context.
6. The disadvantages of TF-IDF are solved by word embeddings. DNNs may be used to learn word embeddings.
7. Generative - They generate their responses by modelling a probability distribution over language (language models). This is currently done by training DNNs on large amounts of data. Seq2Seq models are used in generative chatbots. Transformers are another model used which use self-attention to consider context from other positions of the sentence.
8. Task-Oriented Agents - The classification into interpretable and black box is especially important for task-oriented agents as they are usually used for reaching concrete goals and may be deployed in high-stakes settings.
9. Figure 4 on page 15 shows the pipeline model for task-oriented agents. When using black box systems, the lines between the modules in the pipeline model are blurred. This is because there is a push towards requiring less and less human specification of how the processing should take place.
10. Interpretable task-oriented agents - 1) Finite state systems - The conversation follows a rigid and predefined structure and sequence. Dialog initiative is with the CA and the user can only answer the questions asked. 2) Frame-Based/Slot-Filling systems - There is mixed initiative. The system attempts to fill the slots in a frame through its conversation. The domain (for multi-domain systems) and intent also must be determined. Intent detection is implemented in the form of explicit, hand written rules or with semantic grammars, which define a set of rules which form a compact representation of many different variations of a sentence. This representation is usually insufficient and we require determination of dialog acts and dialog state tracking (ability to keep track of context). The methods used can maintain only a single, deterministic description of the dialog state.
11. Methods to model dialog state in a probabilistic and statistical way - 1) Markov Decision Process (MDP) - There are a set of states (S), a set of actions (A) and a matrix of transition probabilities. The matrix has the probability of moving from a particular state to another when performing a specific action. Any action and corresponding resulting state have an immediate reward associated with them. The dialog policy decides what action the agent takes in a given situation. The dialog policy is optimized using reinforcement learning to maximize reward.
12. 2) Partially Observable MDP (POMDP) - They take into account the inherent uncertainty of being in a state. It defines a belief state which defines a probability distribution over all possible dialog states.
13. Black box task oriented agents - They are similar to black box chatbots. They also use word embeddings and rely on recurrent, convolutional or transformer-based architectures. Memory networks or graph CNNs are also used. Black box systems can also be framed using the pipeline model with a different network for each module. But they are still connected and are dependent on each other so they are not completely independent. There are a methods to improve the interface between DNNs and the knowledge base. Normally the DNN for dialog policy choses the most appropriate information retrieval query. 1) Soft queries with multiple degrees of truth can be used to obtain richer results. 2) Querying of

large databases can be made faster by assimilating the knowledge base into network parameters during training.

14. Evaluation of CAs - Low level language processing metrics - precision, recall, GLUE (general language understanding evaluation)/SuperGLUE (they consist of different natural language understanding tasks and corresponding benchmarking datasets), BLEU, ROUGE
15. Interaction quality evaluation - 1) Conversational interaction - Human evaluation can be categorized into sensibleness (agent's answers make sense in the given context) and specificity (ability to provide specific and informative answers). Measuring these 2 parameters correlates well with measuring perplexity.
16. Other metrics like conversational user experience (overall user rating), coherence (number of sensible response over total number of responses), engagement (number of utterances), topical breadth (number of topics touched), topical depth (number of consecutive utterances of a topic) were used in the Alexa prize.