

Artificial Intelligence and Subfields

Ming-Hwa Wang, Ph.D
Adjunct, Department of Computer Engineering
Santa Clara University
February 4, 2017

Artificial Intelligence

- ***Human intelligence:***

Use of intuition, common sense, judgment, creativity, goal-directedness, plausible reasoning, knowledge and beliefs

- ***Artificial intelligence:***

A field of study that encompasses computational techniques for performing tasks that apparently require intelligence when performed by human

- ***Turing test***

- ***Applications:***

Computer vision, speech and audio processing, natural language processing, robotics, bioinformatics and chemistry, video games, search engines, online advertising and finance

- ***Challenges:***

Solving the tasks that are easy for people to perform intuitively but hard for people to describe formally, e.g., speech understanding, face recognition

Traditional AI

- ***Fundamental issues:***

Knowledge representation, search, perception and inference

- ***Classic techniques:***

- Symbolic, rule-based, or algorithm-based: emphasize the use of prior knowledge, a system with manually created rules has limited success

- Always have exceptions for rules

- Almost all AI problems are NP-hard $O(2^n)$

- ***Artificial neural network (ANN):***

- Reverse engineer the computational principles behind the brain and duplicate its functionality

- The second best solution for almost any problem

- ***Early Failure of AI***

- Moore's Law and parallel/distributed computing

Recent AI Progress

- **Big progress fields:**
Machine learning, data mining, natural language processing, Computer Vision, etc.
- **Empirical, data-driven statistical approaches:**
Achieved encouraging results (statistical revolution), emphasize the data
- **AI Deep Learning:**
The hierarchy of concepts enables the computer to learn complicated concepts by building them out of simpler ones
- **Big Data:**
A supervised deep learning algorithm will generally achieve acceptable performance with around 5,000 labeled examples per category and will match or exceed human performance when trained with a dataset containing at least 10 million labeled examples
- **Changed Mindset:**
Not exact/perfect solution

AI Subfields

- Machine Learning
- Data Mining
- Information Retrieval and Semantic Web
- Speech Recognition and Natural Language Processing
- Image Processing/Recognition and Computer Vision
- Robotics
- Search
- Knowledge Representation and Knowledge Database
- Logic Reasoning and Probabilistic Reasoning
- Expert Systems

Expert Systems

- Domain-specific complex application

Logic Reasoning & Probabilistic Reasoning

- **Logic reasoning:**
Propositional calculus → predicate calculus → first-order predicate calculus
- **Logic programming:**
 - Functional Programming Language: Lisp, Scheme
 - Very High Level Programming Language: Prolog
- **Automated reasoning:**
 - Inference rule (generate new clauses from existing ones),
 - Subsumption (remove unnecessary clauses),
 - Demodulation (replace/rewrite/substitute clauses),
 - Theorem proving and qualification
- **Probability:** additive law, multiplicative law, Bayes' rule
- **Probability inference networks and fuzzy logic**
- **Dempster-Shafter Calculus:**
For manipulate degrees of belief, which doesn't require $B(A) + B(\neg A) = 1$

Knowledge Representation & Knowledge Database

- ***Production system***: pattern-action-oriented knowledge
- ***Inclusion or concept hierarchies***: object-oriented
- ***Mathematic logics***: first-order predicate calculus
- ***Frames of context***
- ***Semantic networks or semantic net***
- ***Constraint schemata or generalized constraints***:
Predicate calculus augmented with procedural information
- ***Relational database and NoSQL***
 - Key-value NoSQL (e.g., Redis)
 - Column-family NoSQL (e.g., Google's BigTable, HBase, Cassandra)
 - Document NoSQL (e.g., MongoDB, CouchDB)
 - Graph NoSQL (e.g., Neo4j)
- ***Challenges***:
 - Close world assumption
 - Expensive knowledge acquisition

Search

- ***Depth-first search***
Using backtracking
- ***Breadth-first search***
Using queue
- ***Heuristic search methods***
With cost function and priority queue or heap
- ***Approximation***
Close but not exact solution with deviation guarantee
- ***Randomize or Probabilistic methods***
- ***Planning***
Makes problem solving more tractable by making the operations conditional

Robotics:

- ***Sensor for perception:***

- Proprioception
- Force sensing
- Tactile sensing or touch sensing
- Sonar
- Camera and computer Vision

- ***Effector for action:***

- Actuator: converts software commands into physical motion
- Locomotion: change the position of the robot within its environment
- Manipulation: move other objects in the environment, change the shape or other physical properties of objects
- End effector: tool

- ***Autonomous Systems:***

Imaging processing/recognition and computer vision

- ***Vision:***

- The richest of the five sensing modalities of human, it occupies approximately $\frac{1}{4}$ of the brain
- Simultaneously with spatial, geometrical relations, and symbolic, semantic structures

- ***Manipulation:***

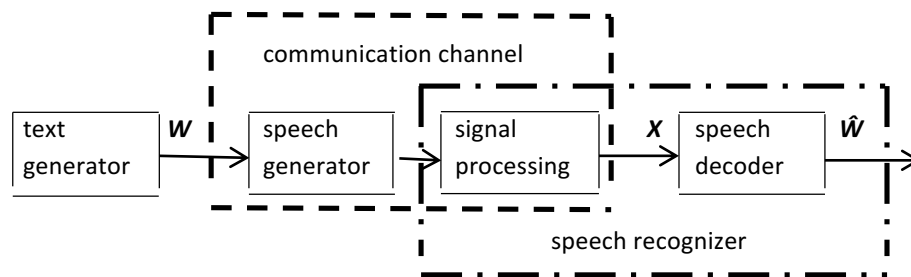
Elementary manipulation, extraction of meaningful structures, edge detection and segmentation into regions, analyzing shape, representing and determining 3D structure from 2D images, special heuristics for handling blocks world scenes

- ***Computer vision:***

The perception model, machine recognition of patterns, picture processing, automatic image analysis, scene analysis, image understanding

Speech & Natural Language Processing

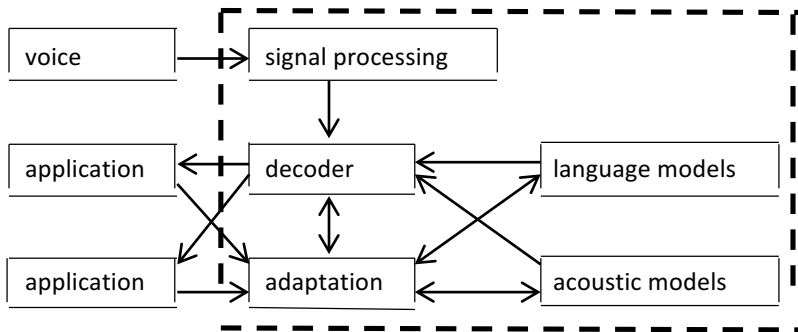
- Speech recognition and synthesis, stemming and lemmatization, syntax and parsing, semantic analysis and knowledge representation, formal language theory, statistical methods, probabilistic models, hidden Markov models, computational linguistic, machine translation, spoken language understanding, question answering, conversational agents, machine translation, summarization, and human-machine interaction
- source-channel mathematical model for a speech recognition system



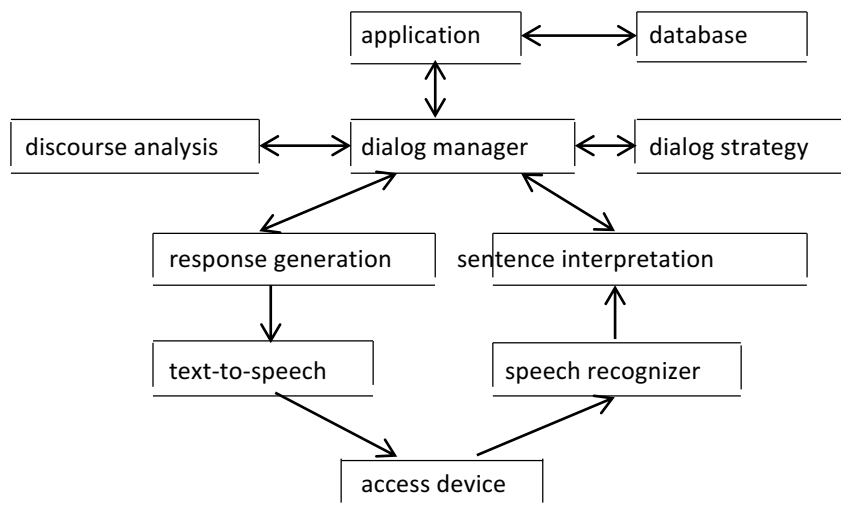
where W is a source word sequence, X is W 's acoustic signal, and \hat{W} is the recognized words sequence

Speech & Natural Language Processing

- basic system architecture of a speech recognition system

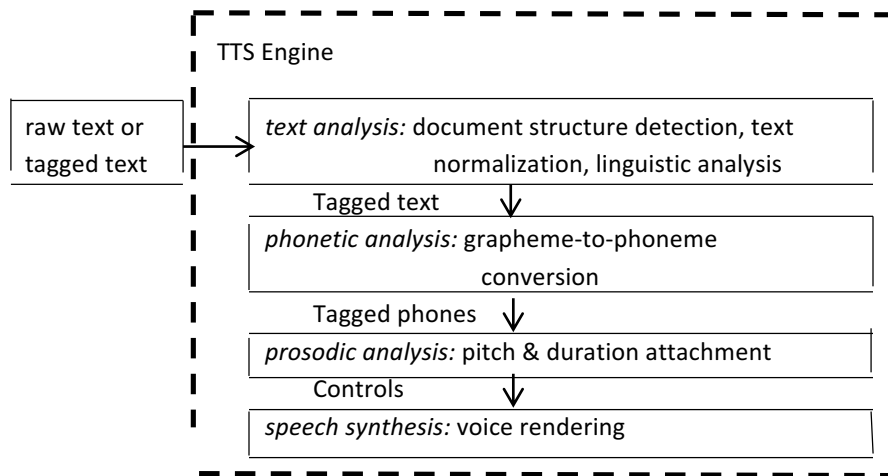


- basic system architecture of a spoken language understanding system



Speech & Natural Language Processing

- basic system architecture of a text-to-speech system



Information Retrieval & Semantic Web

- ***IR:***

Process data offline and store it in inverted index to provide good search performance

- Ranking
- Tools: Lucene, Solr, ElasticSearch, etc.

- ***Semantic Web:***

Semantic-based searching instead of keyword-based searching

- Resource Description Framework (RDF):
- Triples (subject, predicate, and object) and hypergraph
- RDF → RDFS (RDF Scheme) → Web Ontology Language (OWL)
- RDF Query Language SPARQL

Data Mining

- ***Finding similar items:***

Shingling, min hash, and LSH (locality sensitive hashing)

- ***Frequent itemsets:***

Market-basket model, A-Priori algorithm and variations

- ***Mining data stream:***

Sampling, filtering, counting, estimating moments

- ***Link analysis:***

PageRank, topic-sensitive PageRank, link spam, hubs and authority

- ***Text mining and social mining***

Machine Learning

- Recommendation (collaborative filtering), clustering, classification
- Supervised/predictive learning, unsupervised/descriptive learning, semi-supervised learning, reinforcement learning
- ML project lifecycle:
Understanding, data collection/cleaning, feature selection, model selection, training/tuning, evaluation, deployment, review
- Error-based learning: linear regression, gradient descent, multinomial logistic regression, modeling non-linear relationship), artificial neural networks
- Information-based learning: decision tree learning, entropy and information gain, model ensembles like boosting and bagging
- Similarity-based learning: feature space and distance metrics, nearest neighbor algorithm, *k-d* tree
- Probability-based learning: naïve Bayes, Bayesian networks
- No Free Lunch Theorem

Deep Learning

- Using deeper models can reduce the number of units required to represent the desired function and can reduce the amount of generalization error
- 2D matrix algebra → 3D tensor algebra
- Artificial neural network (ANN), convolutional neural network (CNN), recurrent neural network (RNN), deep neural network
- Convex and non-convex optimization
- Performance enhancement with GPU