Artificial Intelligence

Module 9: AI Applications and Ethics

Dr. Chandra Prakash
Assistant Professor
Department of Computer Science and Engineering

(Slides adapted from StuartJ. Russell, B Ravindran, Mausam, Dan Klein and Pieter Abbeel, Partha P Chakrabarti, Saikishor Jangiti

Module 9: AI Applications

- PART 9.1 : Computer Vision and Robotics
- PART 9.2 : Natural language understanding
- PART 9.3 : AI in Healthcare
- PART 9.4 : Ethics of AI

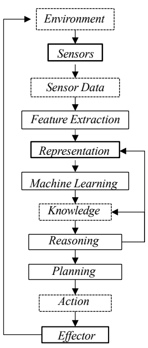

Open Question:

What can be learned from data?

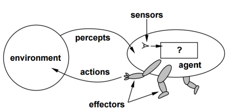

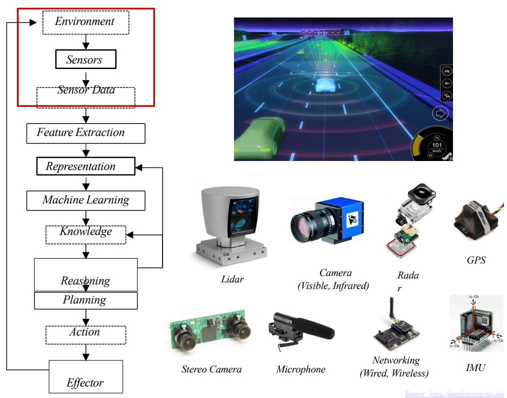

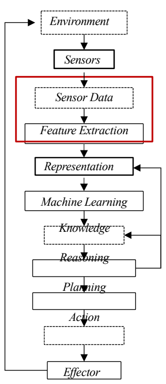

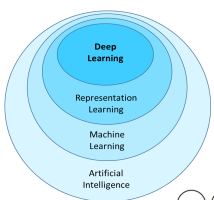

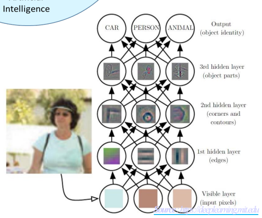

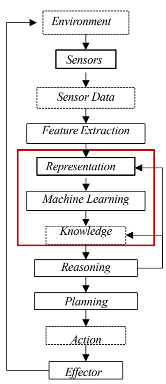

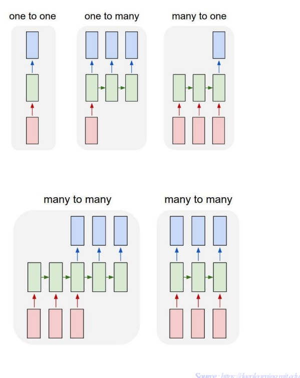

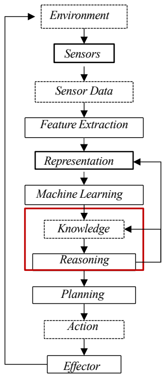

Image Recognition:
If it looks like a duck

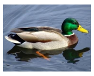

Audio Recognition:
Quacks like a duck

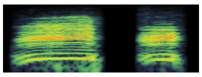

Activity Recognition:
Swims like a duck

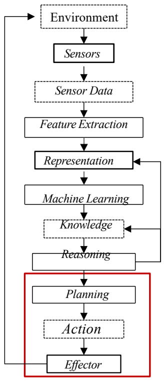

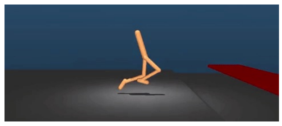

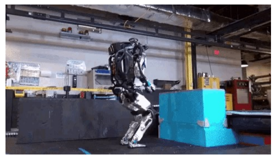

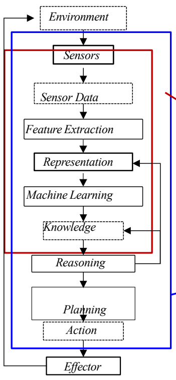

The promise of
Deep Learning

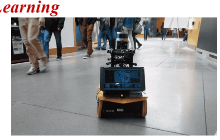

The promise of
Deep Reinforcement Learning

"In part because few real-world problems are as constrained as the games on which DeepMind has focused, DeepMind has yet to find any large-scale commercial application of deep reinforcement learning."

Aug 14, 2019 Wired : https://www.wired.com/story/deepminds-losses-future-artificial-intelligence/

Source : Simulation and Automated Deep Learning

To date, for most successful robots operating in the real world: Deep RL is not involved

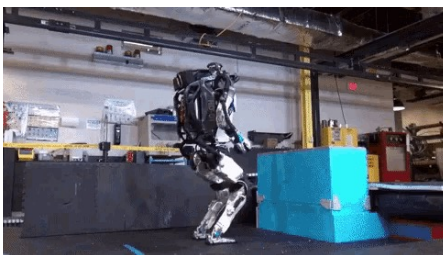

But… that's slowly changing:
Learning Control Dynamics + Learning to Drive

But… that's slowly changing:
Object detection using DRL

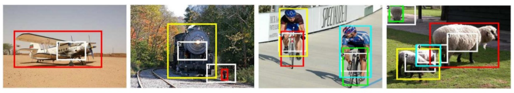

Deep Reinforcement Learning of Region Proposal Networks for Object Detection, 2018

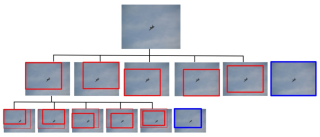

- Hierarchical Object Detection wi… Deep Reinforcement Learning

The outline of application domains of RL in healthcare

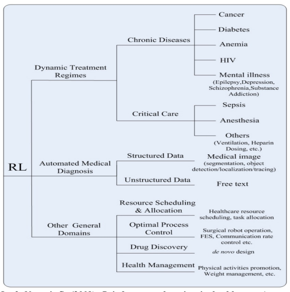

Source : Yu, C., Liu, J., & Nemati, S. (2019). Reinforcement learning in healthcare: A survey. arXiv preprint arXiv:1908.08796.

Other Deep Learning Modules

- **Oxford VGG Model**
- **Google Inception Model**
- **Microsoft ResNet Model**
- **Xception**: Deep Learning with Depthwise Separable Convolutions
  - **[Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1251-1258).]**
- **EfficientNet**: Rethinking Model Scaling for Convolutional Neural Networks
  - **[Tan, M., & Le, Q. V. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. arXiv preprint arXiv:1905.11946.]**
- **Google's word2vec Model**
- **Stanford's GloVe Model**

COMPUTER VISION TASKS

Single Object Recognition

Face Detection, Viola & Jones, 2001

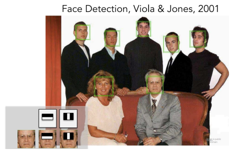

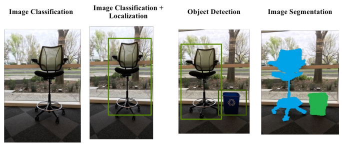

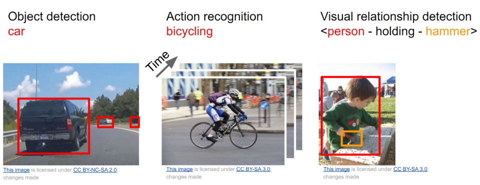

Beyond recognition: Segmentation, 2D/3D Generation

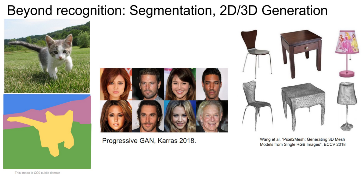

Progressive GAN, Karras 2018.

Wang et al, "Pixel2Mesh: Generating 3D Mesh Models from Single RGB Images", ECCV 2018

Supervised Learning in Natural Language Processing

- Machine translation

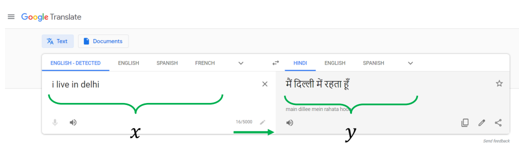

Caption Generation

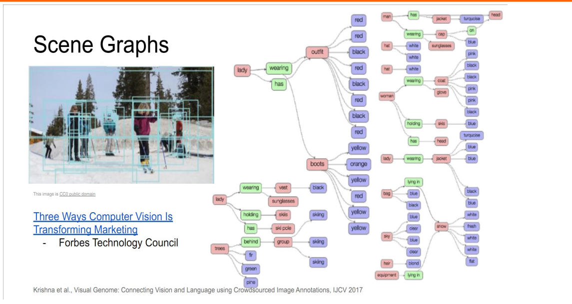

Image Captioning

This image is copyright-free United States government work
Example credit: Andrej Karpathy

Motion Analysis

Clustering Genes

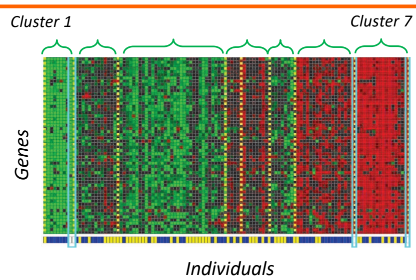

Identifying Regulatory Mechanisms using Individual Variation Reveals Key Role for Chromatin Modification. [Su-In Lee, Dana Pe'er, Aimee M. Dudley, George M. Church and Daphne Koller. '06]

New state-of-the-art:

Unsupervised learning examples

- In generative modeling , we want to learn a distribution over some dataset, such as natural images.
- We can evaluate a generative model by sampling from the model and seeing if it looks like the data.

Meet Sophia: World's First AI Humanoid Robot

https://www.youtube.com/watch?v=S5t6K9iwcdw

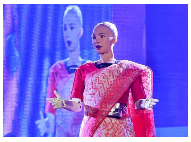

Working

Shotton, Jamie, et al. "Real-time human pose recognition in parts from single depth images." CVPR 2011. Ieee, 2011.

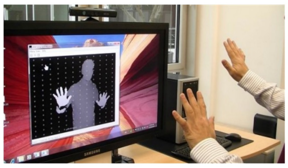

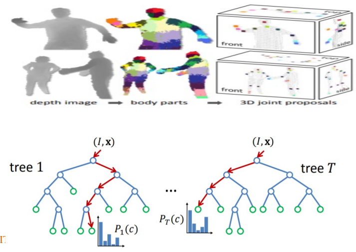

Natural Language Processing

Introduction

- Language is meant for Communicating about the world.
- By studying language,
  - Understand more about the world.
  - Build computational mode of language, powerful tool for communicating about the world.
  - Exploit knowledge about the world, in combination with linguistic facts

Speech is the representation of the mind and

writing is the representation of speech

Source: *https://thinkpalm.com/blogs/natural-language-processing-nlp-artificial-intelligence/*

Language and Intelligence

Turing Test

Computer *Human*

Human Judge

- Human Judge asks tele-typed questions to Computer and Human.
- Computer's job is to act like a human.
- Human's job is to convince Judge that he is not machine.
- Computer is judged "intelligent" if it can fool the judge
- Judgment of intelligence is linked to appropriate answers to questions from the system.

Natural Language Processing (NLP)

- The process of computer analysis of input provided in a human language (natural language), and conversion of this input into a useful form of representation.
- The field of NLP is primarily concerned with getting computers to perform useful and interesting tasks with human languages.
- The field of NLP is secondarily concerned with helping us come to a better understanding of human language.
- NLP includes :
  - Understanding : processing of mapping from an input form into a more immediately useful form.
  - Generation
  - Multilingual translation.

NLP

- Natural language processing
- NLP includes :
  - Understanding : processing of mapping from an input form into a more immediately useful form.
  - Generation
  - Multilingual translation.
- NLP problem can be divided into two tasks:
  - **Processing written text**
    - using lexical, syntactic and semantic knowledge of the language as well as the required real world information.
  - **Processing spoken language**
    - using all the information needed above plus additional knowledge about phonology as well as enough added information to handle the further ambiguities that arise in speech.

Brief History of NLP

- 1940s –1950s: Foundations
  - Development of formal language theory (Chomsky, Backus, Naur, Kleene)
  - Probabilities and information theory (Shannon)
- 1957 – 1970s:
  - Use of formal grammars as basis for natural language processing (Chomsky, Kaplan)
  - Use of logic and logic based programming (Minsky, Winograd, Colmerauer, Kay)
- 1970s – 1983:
  - Probabilistic methods for early speech recognition (Jelinek, Mercer)
  - Discourse modeling (Grosz, Sidner, Hobbs)
- 1983 – 1993:
  - Finite state models (morphology) (Kaplan, Kay)
- 1993 – present:
  - Strong integration of different techniques, different areas.

Forms of Natural Language

- The input/output of a NLP system can be:
  - written text
  - speech
- We will mostly concerned with written text (not speech).
- To process written text, we need:
  - lexical, syntactic, semantic knowledge about the language
  - discourse information, real world knowledge
- To process spoken language, we need everything required to process written text, plus the challenges of speech recognition and speech synthesis.

"

NLP Issues

- The Problem : **English sentences are incomplete descriptions of the information that they are intended to convey.**
  - **Some dogs are outside** is incomplete – it can mean
    - Three dogs are on the lawn.
    - Moti, Hira & Tomy are on the lawn.
  - The good side : Language allows speakers to be as vague or as precise as they like. It also allows speakers to leave out things that the hearers already know.
- The Problem : **The same expression means different things in different context.**
  - Where's the water? (When you are thirsty, it must be potable)
  - Where's the water? (Dealing with a leaky roof, it can be filthy)
  - The good side : **Language lets us communicate about an infinite world using a finite number of symbols.**
- The problem : **There are lots of ways to say the same thing :**
  - Mary was born on October 11.
  - Mary's birthday is October 11.
  - The good side : **Whey you know a lot, facts imply each other. Language is intended to be used by agents who know a lot.**

Ambiguity: Story : Comma(,) killed a man

- Ever wondered about how incorrect punctuation can create havoc?
  - A man was found guilty and awarded a death sentence. The day of hanging had come. Meanwhile his relatives went for the mercy of the then President and he cancelled his death sentence. A telegram was immediately sent to the jail stating-" Save him, not hang him".

But the postal department has sent the telegram which read " save him not, hang him". The jailer saw this message and hanged the person.

- रोको मत ,जाने दो
- रोको, मत जाने दो

Ambiguity

`I made her duck.`

- How many different interpretations does this sentence have?
- What are the reasons for the ambiguity?
- The categories of knowledge of language can be thought of as ambiguity resolving components.
- How can each ambiguous piece be resolved?
- Does speech input make the sentence even more ambiguous?
  - Yes – deciding word boundaries

Ambiguity (cont.)

- Some interpretations of : **`I made her duck.`**
  1. I cooked *duck* for her.
  2. I cooked *duck* belonging to her.
  3. I created a toy duck which she owns.
  4. I caused her to quickly lower her head or body.
  5. I used magic and turned her into a *duck*.
- `duck` – morphologically and syntactically ambiguous: noun or verb.
- `her` – syntactically ambiguous: dative or possessive.
- `make` – semantically ambiguous: cook or create.
- `make` – syntactically ambiguous:
  - Transitive – takes a direct object. => 2
  - Di-transitive – takes two objects. => 5
  - Takes a direct object and a verb. => 4

Why NL Understanding is hard?

- Natural language is extremely rich in form and structure, and **very ambiguous**.
  - How to represent meaning,
  - Which structures map to which meaning structures.
- One input can mean many different things. Ambiguity can be at different levels.
  - Lexical (word level) ambiguity -- different meanings of words
  - Syntactic ambiguity -- different ways to parse the sentence
  - Interpreting partial information -- how to interpret pronouns
  - Contextual information -- context of the sentence may affect the meaning of that sentence.
- Many input can mean the same thing.
- Interaction among components of the input is not clear.

Resolve Ambiguities

- Introduce *models* and *algorithms* to resolve ambiguities at different levels.
  - **part-of-speech tagging**
    - Deciding whether `duck` is verb or noun.
  - **word-sense disambiguation**
    - Deciding whether `make` is `create` or `cook`.
  - **lexical disambiguation**
    - Resolution of part-of-speech and word-sense ambiguities are two important kinds of lexical disambiguation.
  - **syntactic ambiguity**
    - `her duck` is an example of syntactic ambiguity, and can be addressed by probabilistic parsing.

Components of NLP

- **Natural Language Understanding**
  - Mapping the given input in the natural language into a useful representation.
  - Different level of analysis required:
    - ***Morphological analysis,***
    - ***Syntactic analysis,***
    - ***Semantic analysis,***
    - ***Discourse analysis,*** and
    - **Pragmatic Analysis:**
- **Natural Language Generation**
  - Producing output in the natural language from some internal representation.
  - Different level of synthesis required:

    deep planning (what to say), ***syntactic generation***
- NL Understanding is much harder than NL Generation. But, still both of them are hard.

Knowledge of Language

- **Phonology**
  - concerns how words are related to the sounds that realize them.
- **Morphology**
  - concerns how words are constructed from more basic meaning units called morphemes. A morpheme is the primitive unit of meaning in a language.
- **Syntax**
  - concerns how can be put together to form correct sentences and determines what structural role each word plays in the sentence and what phrases are subparts of other phrases.
- **Semantics**
  - concerns what words mean and how these meaning combine in sentences to form sentence meaning. The study of context-independent meaning.
- **Pragmatics**
  - concerns how sentences are used in different situations and how use affects the interpretation of the sentence.
- **Discourse**
  - concerns how the immediately preceding sentences affect the interpretation of the next sentence. For example, interpreting pronouns and interpreting the temporal aspects of the information.
- **World Knowledge**
  - includes general knowledge about the world. What each language user must know about the other's beliefs and goals.

Exercise : Find the Similarity

D1
- Respected Sir, I'm extremely sorry as I couldn't submit my TUT_4_5 Assignment in time. My assignment was completed just in time. But I couldn't submit as it stopped taking responses immediately after the deadline(11:59 pm , 19/09/2021). I'd be really grateful if you could please allow me to submit the assignment. I won't repeat this late-submission in future.

D2
- Respected Sir, I am extremely sorry that I could not submit the TUT 4&5 before the deadline. Initially the deadline was of 22 ' September. I don't have Team App in laptop due to some issue. I am using teams in browser. So i did get any notification and I was not aware about the prepone of the deadline. Sir I have already shown by code for question 1 & questions 2 in class. I have completed my assignment. Sir, could you please allow me to submit the TUT ? I will be very thankful to you. And I assure you that I won't repeat this mistake.

D3
- Respected Sir, I am extremely sorry that I could not submit the TUT 4&5 before the deadline. I completed my work around 1 AM at night. I was busy giving my internship exam . Sir, could you please allow me to submit the TUT ? I will be very thankful to you. And I assure you that I won't repeat this mistake. Thanks & Regards.

D4
- Respected Sir, I am extremely sorry that I could not submit the TUT 4&5 before the deadline. I completed my work around 2 AM at night. Because I was busy in Aptitude , Verbal , Computer Knowledge & Coding rounds of several internships at D2C platform for two days (Saturday and Sunday). Sir, could you please allow me to submit the TUT in the late submission? I will be very thankful to you sir. And I assure you I won't repeat this mistake ever again. Thanks & Regards.

Plagiarism or not ??

CASE Studies

- Delhi air pollution
- Machine learning predicts World Cup winner
- Making a neural synthesizer

Models to Represent Linguistic Knowledge

- Use formalisms (*models*) to represent the required linguistic knowledge.
  - **State Machines**
    - FSAs, FSTs, HMMs, ATNs, RTNs
  - **Formal Rule Systems**
    - Context Free Grammars, Unification Grammars, Probabilistic CFGs.
  - **Logic-based Formalisms**
    - first order predicate logic, some higher order logic.
  - **Models of Uncertainty**
    - Bayesian probability theory.

Steps in Natural Language Processing

- **Morphological Analysis:**
  - Individual words are analyzed into their components and non-word tokens such as punctuation are separated from the words.
- **Syntactic Analysis:**
  - Linear sequences of words are transformed into structures that show how the words relate to each other.
  - Based on Language Rules by analyzer
    - Cat this is.
    - Is this cat.
    - Is cat this.
    - This is cat.
- **Semantic Analysis:**
  - The structures created by the syntactic analyzer are assigned meanings.
  - Mapping is made between the syntactic structures and object in the task domain.
  - Structure for which no such mapping is possible may be rejected.

Steps in NLP (Cont…)

- **Discourse integration:**
  - The meaning of an individual sentence may depend on the sentences that precede it and may influence the meanings of the sentences that follow it.
    - John wanted it.
    - He is a nice guy.
- **Pragmatic Analysis:**
  - The structure representing what was said is reinterpreted to determine what was actually meant.
    - Do you know what time it is?

1. Morphological Analysis

- Individual words are analyzed into their components and non-word tokens such as punctuation are separated from the words.
- Analyzing words into their linguistic components (morphemes).
- Morphemes are the smallest meaningful units of language.

cars	car+PLU	
giving	give+PROG	
geliyordum	gel+PROG+PAST+1SG	- I was coming

- Ambiguity: More than one alternatives

flies	fly_{VERB}+PROG	
	fly_{NOUN}+PLU	
adamı	adam+ACC	- the man (accusative)
	adam+P1SG	- my man
	ada+P1SG+ACC	- my island (accusative)

1. Morphological Analysis (cont.)

- Tokenization
- Stop words
- Stemming
- Part-of-Speech (POS) Tagging
- Lexical Processing

Part-of-Speech (POS) Tagging

- Each word has a part-of-speech tag to describe its category.
- Part-of-speech tag of a word is one of major word groups (or its subgroups).
  - **open classes** -- noun, verb, adjective, adverb
  - **closed classes** -- prepositions, determiners, conjuctions, pronouns, particples
- POS Taggers try to find POS tags for the words.
- `duck` is a verb or noun? (morphological analyzer cannot make decision).
- A POS tagger may make that decision by looking the surrounding words.
  - `Duck!` (verb)
  - `Duck is delicious for dinner.` (noun)

2. Syntactic Analysis

- Syntactic analysis must exploit the results of morphological analysis to build a structural description of the sentence.
- The goal of this process, called **parsing**,
  - is to convert the **flat list** of words that forms the sentence into a structure that defines the units that are represented by that flat list.
  - A flat sentence has been converted into a hierarchical structure and that the structure correspond to meaning units when semantic analysis is performed.
- There are different parsing formalisms and algorithms.
- Most formalisms have two main components:
  - **grammar** -- a declarative representation describing the syntactic structure of sentences in the language.
  - **parser** -- an algorithm that analyzes the input and outputs its structural representation (its parse) consistent with the grammar specification.

2. Syntactic Processing

- CFGs are in the center of many of the parsing mechanisms. But they are complemented by some additional features that make the formalism more suitable to handle natural languages.
- **Parsing Tree**
  - Reference markers are shown in the parenthesis in the parse tree
  - Each one corresponds to some entity that has been mentioned in the sentence.
  - Reference markers are useful later since they provide a place in which to accumulate information about the entities as we get it.

Resolve Ambiguities using Parsing tree

`I made her duck`

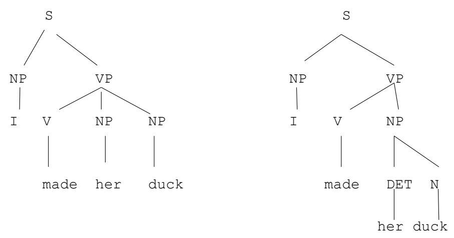

Syntactic Processing

- It plays an important role in natural language understanding systems for two reasons:
  - Semantic processing must operate on sentence constituents.
    - If there is no syntactic parsing step, then the semantics system must decide on its own constituents.
    - If parsing is done, on the other hand, it constrains the number of constituents that semantics can consider.
    - Syntactic parsing is computationally less expensive than is semantic processing. Thus it can play a significant role in reducing overall system complexity.
  - Although it is often possible to extract the meaning of a sentence without using grammatical facts, it is not always possible to do so. Consider the examples:
    - The satellite orbited Mars
    - Mars orbited the satellite
    - In the second sentence, syntactic facts demand an interpretation in which a planet revolves around a satellite, despite the apparent improbability of such a scenerio.

Syntactic Processing

- Almost all the systems that are actually used have two main components:
  - A declarative representation, called a grammar, of the syntactic facts about the language.
  - A procedure, called parser, that compares the grammar against input sentences to produce parsed structures.

Context-Free Grammar

- In a context free grammar the left hand side of a production rule is always a single nonterminal symbol.
- In a general grammar, it could be a string of terminal and/or nonterminal symbols.
- A context-free grammar (CFG) G is a quadruple (V, Σ, R, S) where
  - V: a set of non-terminal symbols
  - Σ: a set of terminals (V ∩ Σ = Ǿ)
  - R: a set of rules (R: V → (V U Σ)*)
  - S: a start symbol.

Context-Free Grammar

- The grammars are called *context free* because – since all rules only have a nonterminal on the left hand side – one can always replace that nonterminal symbol with what is on the right hand side of the rule.
- The *context* in which the symbol occurs is therefore not important.
- Example
  - V = {q, f,}
  - Σ = {0, 1}
  - R = {q → 11q, q → 00f,
f → 11f, f → ε }
  - S = q
  - (R= {q → 11q | 00f, f → 11f | ε })

Grammars and Parsers

- The most common way to represent grammars is as a set of production rules.
- A simple Context-free phrase structure grammar for English:
  - S → NP VP
  - NP → the NP1
  - NP → PRO
  - NP → PN
  - NP → NP1
  - NP1 → ADJS N
  - ADJS → ε | ADJ ADJS
  - VP → V
  - VP → V NP
  - N → file | printer
  - PN → Bill
  - PRO → I
  - ADJ → short | long | fast
  - V → printed | created | want
- First rule can be read as “ A sentence is composed of a noun phrase followed by Verb Phrase”; Vertical bar is OR ; ε represnts empty string.
- Symbols that are further expanded by rules are called nonterminal symbols.
- Symbols that correspond directly to strings that must be found in an input sentence are called terminal symbols.

Grammars and Parsers

- Grammar formalism many linguistic theories, which in turn provide the basis for many natural language understanding systems.
- Pure context free grammars are not effective for describing natural languages.
- NLPs have less in common with computer language processing systems such as compilers.
- Parsing process takes the rules of the grammar and compares them against the input sentence.
- The simplest structure to build is a **Parse Tree**, which simply records the rules and how they are matched.
- Every node of the parse tree corresponds either to an input word or to a nonterminal in our grammar.
- Each level in the parse tree corresponds to the application of one grammar rule.

Example : Parsing tree

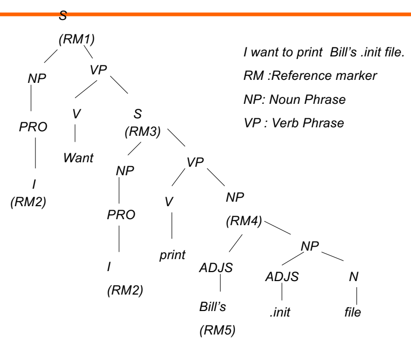

I want to print Bill's .init file.

RM :Reference marker

NP: Noun Phrase

VP : Verb Phrase

A Parse tree for a sentence

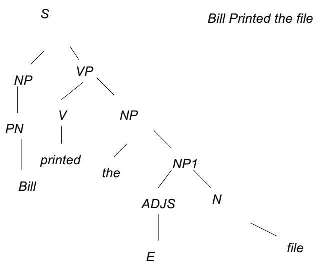

Bill Printed the file

A parse tree

- John ate the apple.

1. S -> NP VP
2. VP -> V NP
3. NP -> NAME
4. NP -> ART N
5. NAME -> John
6. V -> ate
7. ART-> the
8. N -> apple

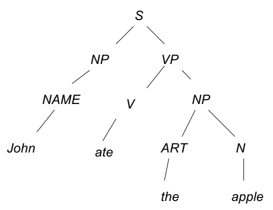

Different lists in grammar rules

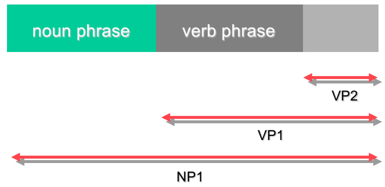


```
sentence(NP1-VP2):-
        noun_phrase(NP1-VP1),
        verb_phrase(VP1-VP2)
```

Exercise:

For each of the following sentences, draw a parse tree

- John wanted to go to the movie with Sally
- I heard the story listening to the radio.
- All books and magazines that deal with controversial topics have been removed from the shelves.

3. Semantic Analysis

- Assigning meanings to the structures created by syntactic analysis.
  - Mapping words and structures to particular domain objects in way consistent with our knowledge of the world.
- Semantic can play an import role in selecting among competing syntactic analyses and discarding illogical analyses.
  - `I robbed the bank` -- bank is a river bank or a financial institution
- Decide the formalisms which will be used in the meaning representation.
- Semantic analysis must do two important things:
  - It must map individual words into appropriate objects in the knowledge base or database
  - It must create the correct structures to correspond to the way the meanings of the individual words combine with each other.

3. Semantic Analysis

- Producing a syntactic parse of a sentence is only the first step toward understanding it.
- Produce a representation of the meaning of the sentence.
- Because understanding is a mapping process, we must first define the language into which we are trying to map.
- There is no single definitive language in which all sentence meaning can be described.
- **The choice of a target language** for any particular NL understanding program must depend on what is to be done with the meanings once they are constructed.

Lexical processing

- The first step in any semantic processing system is to look up the individual words in a dictionary (or lexicon) and extract their meanings.
- Many words have several meanings, and it may not be possible to choose the correct one just by looking at the word itself.
- The process of determining the correct meaning of an individual word is called **word sense disambiguation** or **lexical disambiguation**.
- It is done by associating, with each word in lexicon, information about the contexts in which each of the word's senses may appear.
- Sometimes only very straightforward info about each word sense is necessary. For example, baseball field interpretation of diamond could be marked as a LOCATION.
- Some useful semantic markers are :
  - PHYSICAL-OBJECT
  - ANIMATE-OBJECT
  - ABSTRACT-OBJECT

Lexical Processing

- The purpose of lexical processing is to determine meanings of individual words.
- Basic methods is to lookup in a database of meanings -- **lexicon**
- We should also identify non-words such as punctuation marks.
- Word-level ambiguity -- words may have several meanings, and the correct one cannot be chosen based solely on the word itself.
  - bank in English
  - yüz in Turkish
- Solution -- resolve the ambiguity on the spot by POS tagging (if possible) or pass-on the ambiguity to the other levels.

Sentence-Level Processing

- Several approaches to the problem of creating a semantic representation of a sentence have been developed, including the following:
  - **Semantic grammars**, which combine syntactic, semantic and pragmatic knowledge into a single set of rules in the form of grammar.
  - **Case grammars**, in which the structure that is built by the parser contains some semantic information, although further interpretation may also be necessary.
  - **Conceptual parsing** in which syntactic and semantic knowledge are combined into a single interpretation system that is driven by the semantic knowledge.
  - **Compositional semantic interpretation**, in which semantic processing is applied to the result of performing a syntactic parse

4. Discourse Integration

- Discourses are collection of coherent sentences (not arbitrary set of sentences)
- Discourses have also hierarchical structures (similar to sentences)
- **anaphora resolution** -- to resolve referring expression
  - `Mary bought a book for Kelly. She didn't like it.`
    - **She** refers to Mary or Kelly. -- possibly Kelly
    - **It** refers to what -- book.
  - `Mary had to lie for Kelly. She didn't like it.`
- Discourse structure may depend on application.
  - Monologue
  - Dialogue
  - Human-Computer Interaction

5. Pragmatic Analysis

- The final step toward effective understanding is to decide what to do as a results.
- One possible thing to do is to record what was said as a fact and be done with it.
- For some sentences, whose intended effect is clearly declarative, that is precisely correct thing to do.
- But for other sentences, including this one, the intended effect is different.
- We can discover this intended effect by applying a set of rules that characterize cooperative dialogues.
- The final step in pragmatic processing is to translate, from the knowledge based representation to a command to be executed by the system.
- The results of the understanding process is
- Lpr /wsmith/stuff.init
- "lpr" is th operating system's file print command.

Statistical NLP

- Corpora
- Counting the elements in a corpus
- N-Grams
- Smoothing

Spell Checking

- A spell checking is one of the basic tools required for language processing.
- It is used in a wide variety of computing environments including word processing, character or text recognition system, speech recognition and generation.
- IT involves:
  - Identifying words and non words
  - Suggesting the complex possible alternatives for its correction
- "Divya sar on box"
  - Sar --- sat

  "Divya at on box"

Some NLP Applications

- Machine Translation – Translation between two natural languages.
  - See the Babel Fish translations system on Alta Vista.
- Information Retrieval – Web search (uni-lingual or multi-lingual).
- Query Answering/Dialogue – Natural language interface with a database system, or a dialogue system.
- Report Generation – Generation of reports such as weather reports.
- Some Small Applications –
  - Grammar Checking, Spell Checking, Spell Corrector

Machine Translation

- Machine Translation -- converting a text in language A into the corresponding text in language B (or speech).
- Different Machine Translation architectures:
  - interlingua based systems
  - transfer based systems

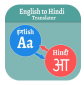

- How to acquire the required knowledge resources such as mapping rules and bi-lingual dictionary? By hand or acquire them automatically from corpora.
- Example Based Machine Translation acquires the required knowledge (some of it or all of it) from corpora.

Case Study : Text Summarization

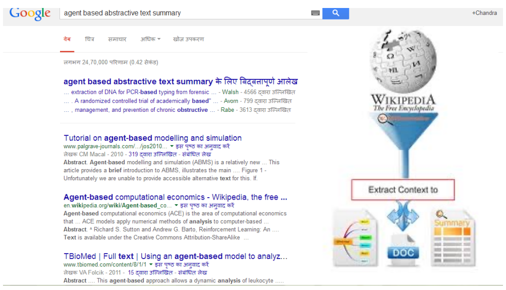

Case Study : Text Summarization

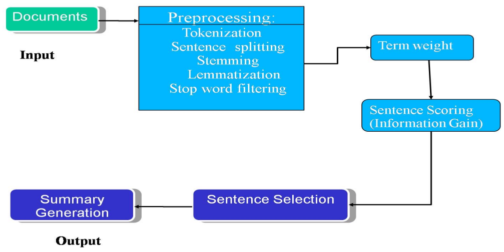

Chandra Prakash, Anupam Shukla "Automated summary generation from singe document using information gain " Springer, Contemporary Computing ,Communications in Computer and Information Science Volume 94, pp 152-159, 2010 .

Methodology proposed (HAMS)

- Input:
  - Document with text is fed into the system.
- Preprocessing:
  - Tokenization: Divides the character sequence into words
  - sentence splitting further divides sequences of words into sentences, and so on.
  - Stemming or Lemmatization
  - Stop word filtering Feature Extraction :
- Sentence Ranking: Machine Learning
- Human Feedback
- Output\ Result:
  - Generated Summary
  - an abstract.

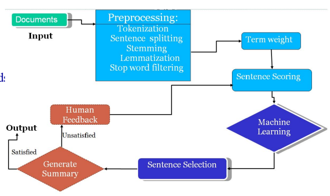

Methodology Steps..

Methodology for text summarization involves

- **Term Selection using Pre-Processing**
  - Tokenization or Segmentation
  - Stop word Filtering
  - Stemming or Lemmatization
- **Term weighting**
  - Term Frequency (TF):

    $W_i(T_j)=f_{ij}$

    where f_{ij} is the frequency of ***j*th** term in sentence *i*.

  - Inverse Sentence Frequency (ISF) : $W_i(T_j) = f_{ij} \times \log\left(\frac{N}{n_j}\right)$

    where *N*=no of sentences in the collection

    n_j =no of sentence where the term *j* appears.

Methodology Steps (cont…)

Weight of a Term is calculated as :

$(TW)i, j= (ISF)\ I,j$

Where $(TW)I,j$ is Term weight if ith sentence and jth Term.

- **Sentence Signature**
  - Sentences that indicate key concepts in a document.
  - Term Sentence Matrix

$$(TSM)=\begin{bmatrix} W11 & W12 & .... & W1n \\ W21 & W22 & .... & W2n \\ .... & .... & .... & .... \\ Wm1 & Wm2 & .... & Wmn \end{bmatrix}$$

  - Information Gain
    - Term Frequency Weight Score
    - Inverse Sentence Frequency score
    - Normalized Sentence length score

$$(NSL)_i = \frac{\text{No of Words occuring in the sentences}}{\text{No of words occuring in the longest sentence in the document}}$$

    - Sentence position score

$$(SPS)i = \frac{n-i+1}{n}$$

    - Numerical Data Score

$$(PNS)i = \frac{\text{No of numerical data in the sentences}}{\text{Length of the sentence}}$$

Methodology Steps (cont…)

- Information Gain is calculated as

Information Gain (IG) = $(TFW)i+ ISFS(Tj)i + (NSL)i +(SPS)i+ (PNS)i$

where i is the sentence and j is the term

- Term-Sentence matrix after IG :

$$(TSM)=\begin{bmatrix} IG(W11) & IG(W12) & .... & IG(W1n) \\ IG(W21) & IG(W22) & .... & IG(W2n) \\ .... & .... & .... & .... \\ IG(Wm1) & IG(Wm2) & .... & IG(Wmn) \end{bmatrix}$$

Result

Comparison of generated text summary for HAMS

Methods	Precision value (P)	Recall Value(R)	F-score
SAAR (user feedback)	90	85	**87.42**
IG summary	75	65	70.57
OTS	75	60	66.66
PS	75	60	66.66
ETSS	75	60	66.66

Compared with some available automated text summarizers Open Text summarizer (OTS), Pertinence Summarizer (PS), and Extractor Test Summarizer Software (ETSS)

Comparison of Recall, Precision Value and F-score for HAMS

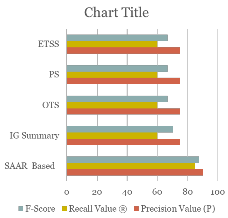

C. Prakash; A, Shukla, *Human Aided Text Summarizer "SAAR" using Reinforcement Learning*"Soft Computing and Machine Intelligence (ISCMI), 2014 International Conference on , vol., no., pp.83,87, 26-27 Sept. 2014

Exercise

- How Hindi based NLP system can be developed ??

Next Time : Ethics of AI

- Robots vs Humans
- Jobs
- Bias
- Fairness
- Accountability
- Transparency
- Privacy
- Ethical uses

References

- ***Artificial Intelligence** by Elaine Rich & Kevin Knight*, Third Ed, Tata McGraw Hill
- ***Artificial Intelligence and Expert System** by Patterson*
- http://www.cs.rmit.edu.au/AI-Search/Product/
- http://aima.cs.berkeley.edu/demos.html (for more demos)
- ***Artificial Intelligence and Expert System** by Patterson*
- Slides adapted from CS188 Instructor: Anca Dragan, University of California, Berkeley
- Slides adapted from CS60045 ARTIFICIAL INTELLIGENCE