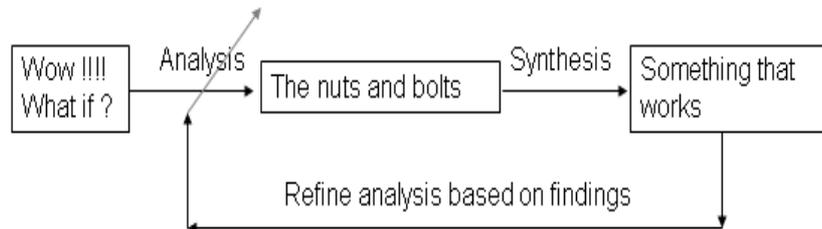


## Memory based Reasoning – Warm-up Lecture (02/01/2009)

(Note: these notes are intended to serve as a rough sketch of the lecture; these are NOT a substitute for standard textbook and/or reference literature)

- We got introduced to each other.
  - TAs for the course:
    - Raghunandan M.A. ([raghunandan.ma@gmail.com](mailto:raghunandan.ma@gmail.com))
    - Vipin B. S. ([vipin.bl@gmail.com](mailto:vipin.bl@gmail.com))
- We start with the question: why did we forget most of what we studied in high school/ higher secondary school? How can we make CS625 any different from those courses?
  - Three phases of learning:
    - Phase 1: we need to feel excited about knowing/learning something (piano sounds great -- won't it be great if I could learn playing piano? That way I could compose my own music and play it too...)
    - Phase 2: we go to a piano teacher. He takes us through the **analysis** phase. Unless we get to the common building blocks (the notes, the grammar, the nuts and bolts), we won't be able to appreciate how music can be composed from primitives.
    - Phase 3: We have a good handle of the analysis, and can go ahead and try our hand at composing new music using the building blocks, guided by the repertoire of music pieces we have mastered. This is the **synthesis** phase.



- Phase 2 can get boring at times, unless we keep the big picture right. Most kids find maths uninteresting in high school since they get trapped in Phase 2. Maths provides the analytic tools, but we need the Phase 1 to motivate kids to ask questions, to answer which we need these tools. Thus, Phase 1 connects learning to real world needs. We also need Phase 3 to show that the mathematical tools are fairly general and can be extrapolated to other problem settings as well. Phase 1 and 3 are a must to make learning fun. We forget things that are thrust upon us.
- Kindergarten kids fail to understand why they have to keep writing the letters a to z time and again, day in and day out. It is only much later that they see that words and sentences (which have meaning, unlike letters) can be composed of these letters.
- In the context of CS625:
  - Phase 1, two example motivations:
    - M1: How do humans organize and retrieve knowledge from their memories for problem solving? No doubt, cracking this is hard : “If the human brain would be so

simple that we could understand it, we would be so simple that we couldn't." – L. Watson. We are more interested in computational models than in purely psychological ones.

- M2: What if we could use the understanding obtained in M1 to actually make the web come alive? To be more precise, can we treat the web as a massive brain which exposes itself to us much like humans do, and interacts with us for problem solving? Some faculties would be:
  - Ability to organize/abstract knowledge (existing in structured/unstructured/multimedia forms) and index them for appropriate reminding
  - Associative memory
  - Proactively learning in the course of each interaction (through feedback/ expectation failures), asking questions to facilitate learning
- M1 is a science question, M2 is an engineering question. Unlike branches like civil or electrical engineering which are based on a solid scientific understanding (physical sciences), in Computing we talk of “Computer Sciences AND Engineering”. We have to develop computational models of nature (in our case human memory), as well as use these models to develop engineering artefacts (“intelligent” systems). CS625 will have both.
- Intelligent systems (Refer Chapter 1 of AI textbook by Russell and Norvig for a part of this section)
  - Two viewpoints :
    - AI systems should have “human-like” thinking/behaviour
    - AI systems should think/ behave rationally

	Thinks Rationally	Thinks like humans
	Acts Rationally	Acts like humans

Humans are not always rational: (Tversky and Kahneman , 1982 “Judgement under uncertainty: Heuristics and Biases)

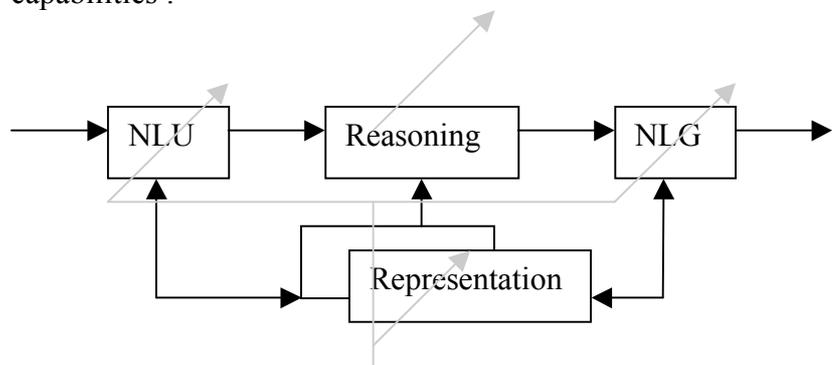
- A more fundamental question: is it always a good idea to emulate humans?
  - Humans would have never discovered wheels if they only looked at emulating how animals move around. On smooth surfaces wheels work better than many other

modes. On a computationally “smooth” surface (where certain tasks like enumeration/sorting are easier than vision, for example) maybe emulating human brains abilities may not always be the best idea?

- Aeronautical engg. texts do not define the goal of their field as making “machines that fly so exactly like pigeons that they can fool even other pigeons” (Russell, Norvig)

- Turing Test : “the computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written response came from a person or not” (Russell Norvig)

- To qualify the computer would need to have the capabilities :



Machine Learning

- NLU: natural language understanding
- NLG: natural language generation
- Reasoning
- Representation
- Machine Learning
- Total Turing Test: needs the following as well
  - Vision
  - Robotics
- Above 7 fields (6 if we club NLU and NLG into NLP) : almost cover all of AI, as it is today
- We briefly covered Models, Rules and Cases as knowledge representation schemes. How can we do diagnosis using a model, using a system of rules and using past experiences
  - Experts (doctors or helpdesk personnel for example) evolve in their lifetime – they may start with textbook knowledge (partial models), then switch over to a combination of models and rules (rules may be viewed as shortcuts in the model), and later almost wholly rely on memories (experiences).
  - Well defined mathematical models are rare in many real world scenarios – memory based reasoning is effective in such ill-defined domains.

- Also, even when a model (whole/partial) is available, it may not be efficient to solve problems using first principles each time a new problem comes in.  
(In later lectures, we will have occasions to contrast these modes of reasoning in more detail)
- We looked at a diagnosis problem: how do we arrive at the most likely hypothesis? Example: medical diagnosis. How is model-based diagnosis different from MBR-based diagnosis?
- Human Memories and Machine Memories
  - How is human memory different from an RDBMS ?
    - Ability to abstract/generalize/re-organize/compress: This is evident from the following three examples:
      - A person reading a newspaper narrates what he read very differently from what he actually read
      - If we go to a restaurant for the first time, we can describe our experience in good detail, but after around 20 odd trips, we are not reminded of individual instances unless they are unique in some way. Rather, we have a generalized structure derived from the instances, that helps us in creating expectations when we visit a restaurant next time. We learn when such expectations are violated (let's say all our 20 instances suggest we eat and then pay, and we go to a new restaurant which needs us to pay before we eat) – learning leads to revising the structures
      - A novice programmer, an average-skilled programmer and an expert programmer are each shown a factorial program for 10-30 secs (?) – then the program is taken away, and they are asked to reproduce what they have seen. The novice preserves the surface level features of the program, but may generate something that does not compile. The expert on the other hand, may write a program that looks quite different from the program he/she was shown, but does the same job (computes factorial correctly). Similar experiment for chess players. This shows that what we see can map onto quite different things in our memory based on our level of expertise. Experienced people develop more sophisticated representations of the world they reason on, which leads to more effective/efficient problem solving. (Reference for this experiment : Ways of Thinking, by Laszlo Mero)

- Content-addressable (we have not covered this in this lecture)
- Proactive (generating expectations) , the “racing mind” (Roger Schank) -- when there is apparently no problem to solve and we are at rest, our minds are still at work – there is evidence to suggest that our memories get reorganized so as to stay better prepared to handle new problems in future
- Our memories are not good at enumerating things; also remembering a 10 digit mobile number could be hard unless there are distinct patterns (chunking)
- That our memories are interestingly different can be seen from Constructive and Discriminative of words. We can discriminate between word meanings easily, but we are bad at giving constructive definitions of words (as in dictionaries)
- Many such differences can be traced to the different knowledge forms that we have – note that machines mostly rely on rational (explicable) knowledge only.
  - Types of knowledge (Roger Schank, Dynamic Memory Revisited)
    - Rational (easily explicable, why did someone choose one job against another?)
    - Emotional (why does someone love his mother? Hard to cook up 5 reasons)
    - Physical (how do I play badminton? Hard to teach anyone on the blackboard)
    - Nonconscious / subconscious (what processes were involved as I read and understood a story?)
  - Homework: Try and find as many fascinating facts about human memory as you can.
- Knowing a little more about how human memory works may help us build better search engines.
  - “AI researchers have long recognized that the more a system knows about a particular state of affairs, the longer it takes to retrieve the relevant information, and this presents a general problem where scaling up is concerned. Conversely, the more a human being knows about a situation or individual, the easier it is to retrieve other relevant information.”
    - Hubert L. Dreyfus
  - This course looks at how we can use MBR principles to build systems that can model concepts from texts, and use the same for effective retrieval. Better search engines?
- Course content (roughly):
  - Memory Based Reasoning : Conceptual Foundations, Early systems
  - Understanding Text : The Conceptual Dependency approach
  - Fundamentals of Machine Learning

- Case Based Reasoning (CBR) : Algorithms for retrieval/adaptation
- Applications of CBR – research issues
- Textual CBR : from words to concepts ...
- Textual CBR : algorithms for effective and efficient retrieval
- Evaluation (roughly):
  - Midsem : 20
  - EndSem : 30
  - Small assignment : 10
  - Big assignment (group) : 25
  - Reading Assignment : 15
- Class Times: Wednesdays: 3:30PM – 5:00 PM (To be revised?)  
Fridays: 3:00 PM – 4:30 PM