# Practical AI
# (a.k.a. Business Intelligence)

***Introduction***

Prof. Abraham Bernstein, Ph.D.

Dynamic and Distributed Information Systems

## What this class is all about

- What is classical AI?

  The study of the principles by which natural or artificial machines manipulate knowledge:
  - how knowledge is acquired
  - how goals are generated and achieved
  - how concepts are formed
  - how collaboration is achieved

- What this class is not about:
  - Understanding intelligence as a phenomenon

# Or said in the words of Allan Newell:

… Exactly what the computer provides is the ability *not to be rigid and unthinking but, rather, to behave conditionally.* That is what it means to apply knowledge to action: It means to let the action taken reflect knowledge of the situation, to be sometimes this way, sometimes that, as appropriate …

And dynamic behavior is central to businesses:

"The ultimate goal is *flexibility*…"
Jorma Ollila, former CEO of Nokia



# Success Story: Medical Expert Systems

- Mycin (1980)
  - Expert level performance in diagnosis of blood infections
- Today: 1,000's of systems
  - Everything from diagnosing cancer to designing dentures
  - Often outperform doctors in clinical trials
- Major hurdle today:
  - non-expert part
  - doctor/machine interaction

# Success Story: Data Mining

- Bases:
  - Statistics, Machine Learning, Data Bases
- Applications:
  - House price estimations
  - Credit scoring
  - Fraud detection (Credit cards, Cell phones)
  - Marketing (Cumulus)
  - Information Filtering (Spam)
  - Human Computer Interaction (Clippy & Co.)




courtesy JPL

# Success Story: Chess

- I could feel – I could smell – a new kind of intelligence across the table
  - Kasparov
- Examines 5 billion positions/second
- "Intelligent" behavior emerges from brute-force search

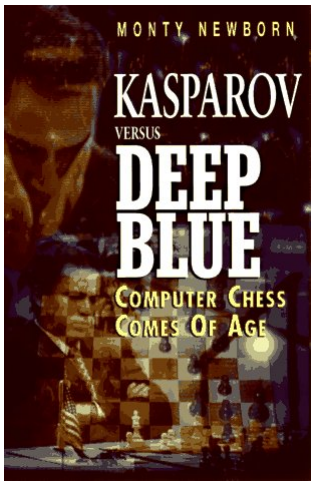




# Not Speed Alone…

- Speech Recognition
  - "Word spotting" feasible today
  - Continuous speech – rapid progress
  - Turns out that "low level" signal not as ambiguous as we once thought
- Translation / Interpretation / Question-answering
  - Very limited progress

  *The spirit is willing but the flesh is weak. (English)*

  *The vodka is good but the meat is rotten. (Russian)*

# Historic Perspective

- 1940's - 1960's: Artificial neural networks
  - McCulloch & Pitts 1943
- 1950's - 1960's: Symbolic information processing
  - General Problem Solver – Simon & Newell
  - "Weak methods" for search and learning
  - 1969 - Minsky's Perceptrons
- 1940's – 1970's: Control theory for adaptive (learning) systems
  - USSR – Cybernetics – Norbert Weiner
  - Japan – Fuzzy logic
- 1970's – 1980's: Expert systems
  - "Knowledge is power" – Ed Feigenbaum
  - Logical knowledge representation
  - AI Boom
- 1985 – 2000: A million flowers bloom
  - Resurgence of neural nets – backpropagation
  - Control theory + OR + Pavlovian conditioning = reinforcement learning
  - Probabilistic knowledge representation – Bayesian Nets – Judea Pearl
  - Statistical machine learning
- 2000's: Towards a grand unification
  - Unification of neural, statistical, and symbolic machine learning
  - Unification of logic and probabilistic KR
  - Autonomous systems



# This course focuses on

- AI techniques that have been employed in practice:
  - Smart Searching
  - Knowledge representation/reasoning
  - Probabilistic reasoning
  - Learning
  - Natural language processing (a tiny bit as examples)
- Practical Case studies of their use
  - Credit Scoring, Machine Translation, Fraud Detection (we will see)

# Course Mechanics

- Time/space: Tue 2–4 pm, BIN 2.A.01
- Lecturer: Prof. A. Bernstein, Ph.D.
  Dr. Panos Karras
- Assistants: Cathrin Weiss, Jiwen Li
- Requirements: 3 Assignments, Final Exam
- Assignments: eLearning, paper & pencil, programming
- Book: Russel, S., & Norvig, P. *Artificial Intelligence: A Modern Approach* **2e**
- Web–site: www.ifi.unizh.ch/ddis/

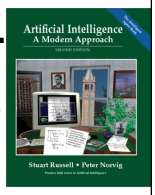




# Course Mechanics: Schedule

| *Date* | *Subject* | *Assignment* |
|---|---|---|
| 23.2. | Introduction | |
| | **Part 1: Intelligent Search - Problem Solving and Planning as Search, Search** | |
| 2.3. | Informed Search | A1 out |
| 9.3. | Constraint Satisfaction, Adversarial Search | |
| | **Part 2: Knowledge intensive processing** | |
| 16.3. | Logic review (Propositional Logic, First Order Logic) | A1 back, A2 out |
| 23.3. | Logical Programming | |
| 6.4. | Logical Programming, Knowledge Representation | |
| | **Part 3: Uncertainty, probability, learning and probabilistic reasoning** | |
| 12.4. | Modeling uncertainty - probability revisited | |
| 20.4. | Bayesian Belief Networks | A2 back, A3 out |
| 27.4. | Bayesian Belief Networks, Probabilistic Reasoning | |
| 4.5. | Reasoning over time - (hidden) Markov models | |
| 11.5. | Induction Part I: Decision Trees | |
| 18.5. | Induction Part II: Naive Bayes | A3 back |
| 25.5. | TBA | |
| 1.6. | Questions and answers - Wrap-up | |