AI & IP

Fabio G. Cozman Universidade de São Paulo - Brazil

Hot Summer in Al

- Many applications; sharp interest.
- Old "winters" are forgotten...



- ► Still, much debate...
 - What is AI? Is it possible?
 - Is it risky?

A bit of AI

- ► AI: devices that think/act in ways that seem to be intelligent.
 - Perhaps similarly to humans.
 - Perhaps following some rationality standards.
- ► A few distinctions:

General try to be "really" smart, may pass Turing test, etc. **Pragmatic** try to solve specific problems, may not be human-like.

Neats look for elegant solutions, typically with mathematical basis.

Scruffies want to build complex systems that work well.

- Search (and optimization).
- Planning.
- Knowledge representation and reasoning.
- Machine learning.
- ► Vision, language processing, sensing, robotics.

Key ingredients?

Knowledge Representation and Reasoning Logic, probabilities ...

Decision Making Planning, Negotiation ...

Machine Learning

Statistical, neural, ...

AI, now: data, computing, insights

Knowledge Representation and Reasoning

Decision Making

Machine Learning Statistical Neural, Evolutionary

Problem solving by search: problem had to be well represented.

- General Problem Solver: famous effort.
 - Separated declarative knowledge from search.
 - A production system with if-then rules.



AI and probability: Not always close

► McCarthy & Hayes (1969):

... a formalism that required numerical probabilities would be epistemologically inadequate.

Why? probabilities are not available, not easy to get, not easy to handle.



After the winter: Expert systems

- ► Many systems: PROSPECTOR, CADUCEUS, etc.
- ► The famous MYCIN rule-system (shell E-MYCIN):

```
(defrule 52
if (site culture is blood )
  (gram organism is neg )
  (morphl organism is rod )
  (burn patient is serious)
  then 0.4
  (identity organism is pseudomonas )
```

Explosion of approaches

- Various logic-based formalisms to capture epistemic states.
- Confidence factors.
- Dempster-Shafer theory belief functions.
- Possibility theory.
- Probabilistic logic.
- Probability intervals.
- ▶ ...and probabilities (Bayesian networks).

Bayesian networks



Probabilistic planning: MDPs

- One has states, actions, and transitions with rewards.
- Transitions are probabilistic (maybe Bayesian networks).
- Probabilities may not be precisely specified.



By Shipra Agrawal, IEOR8100, Lecture 1 - Reinforcement Learning

The rise of data-driven machine learning

- ▶ Our world is now awash with (big) data.
- Extracting patterns and following them turns out to be a winning strategy.

A. Halevy, P. Norvig, F. Pereira, The unreasonable effectiveness of data, IEEE Intelligent Systems, 24(2):8–12, 2009.

Winograd scheme



The trophy would not fit in the suitcase because it was too [small/large].

Classification: k-neareast neighbors, naive Bayes, tree classifiers, random forests, SVMs.
 Regression: linear regressors, splines.
 Representation learning: PCA, ICA.
 Clustering: k-means, topic models.

Multi-Layer Perceptron



Deep learning: the hot stuff

Huge models, learning from big data.



- There has been discussion on "epistemic vs. aleatoric" uncertainty in machine learning,
 - ▶ where *aleatoric* refers to random disturbances,
 - ▶ and *epistemic* refers to lack of modeling knowledge.
- Recent proposals based on credal sets and the like.

Is AI risky? Utopia e Dystopia

Utopia:

- Increase in productivity, with fair distribution of results.
- Better medicine, better public services, ...
- Humans can focus on interesting activities.

Dystopia:

- Super-intelligence: humans are mere mosquitoes, kill them!
- Killer robots: technical failure or pure evil.
- Changes in job market.
- Loss of control, privacy, understanding.

Errors



- **Input:** gVor elf Jahren sitzt Sufjan Stevens auf der Buhne im Kolner Prime Club (heute: Luxor). **Translation:** About eleven years ago
- Sufjan Stevens sits on the stage in Cologne Prime Club (today: Luxor).

Hallucination: EPEFA is the first year of the year.

Interpreting/Explaining ML

Often, machine learning resorts to very complex models.



DARPA's Explainable AI



Conclusion

- ► AI is going through a hot summer...
 - scruffy, pragmatic (not "ideal"),
 - data-driven (statistical/neural machine learning).
- ► Hot debate on how to mix data-driven and knowledge-based AI.
 - Neuro-symbolic AI?
 - Knowledge-enhanced Machine Learning?

- Throughout its history, AI has been interested in formalisms that include probabilistic indeterminacy and imprecision.
 - Probabilistic logic and probabilistic argumentation, credal networks, causal reasoning, imprecise MDPs.