The Future of AI & Law

L. Karl Branting
BAE Systems, Inc.
Columbia, MD, USA

ICAIL 2005, Bologna, Italy

Overview of Talk

- Substance
  - The goals of our discipline and community
  - What we have already achieved
  - How we should focus our efforts to achieve our remaining goals
- Method
  - Identify lessons from larger AI community
  - Apply those lessons to our community

Overview of Talk

- Scope
  - The last ≈ 18 years, since ICAIL 1987
- Themes
  - Rigorous task analysis
  - Decomposition into independent sub-problems
  - Replicated empirical evaluation

Goals of AI *Discipline*

- A computational theory of the mind
- Automated agents with human-like social, learning, and problem-solving characteristics
- Tools that solve problems using explicit knowledge that is
  - expert
  - common-sense, or
  - automatically acquired

Goals of AI *Community*

- Shared research objectives, evaluation criteria, and resources
- Literature stream embodying key research results

Successes of AI Discipline

- Combinatorial search, e.g.,
  - 500 city traveling salesman
  - 10^6 queens problem
  - Hubble space telescope scheduling

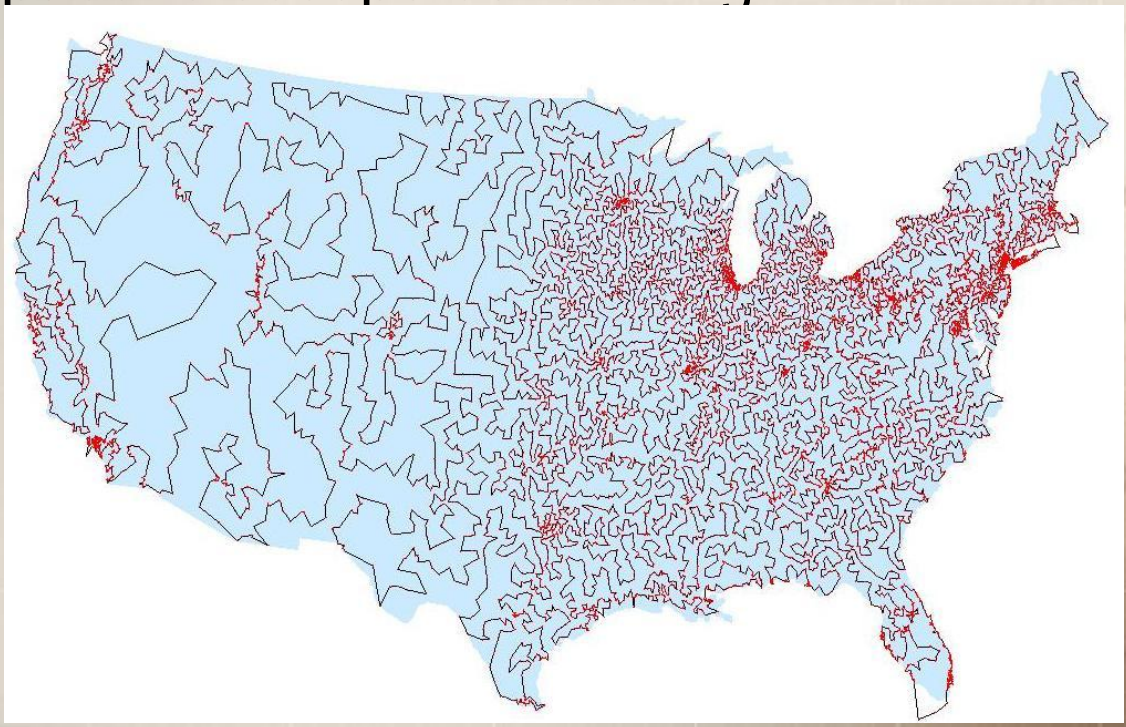

Successes of AI Discipline

Combinatorial search – Success came from:

1. Analysis of heuristics in a neural network scheduler for the Hubble Space Telescope
2. Investigation of phase transitions in the space of satisfiability problems
3. Well-defined, shared tasks

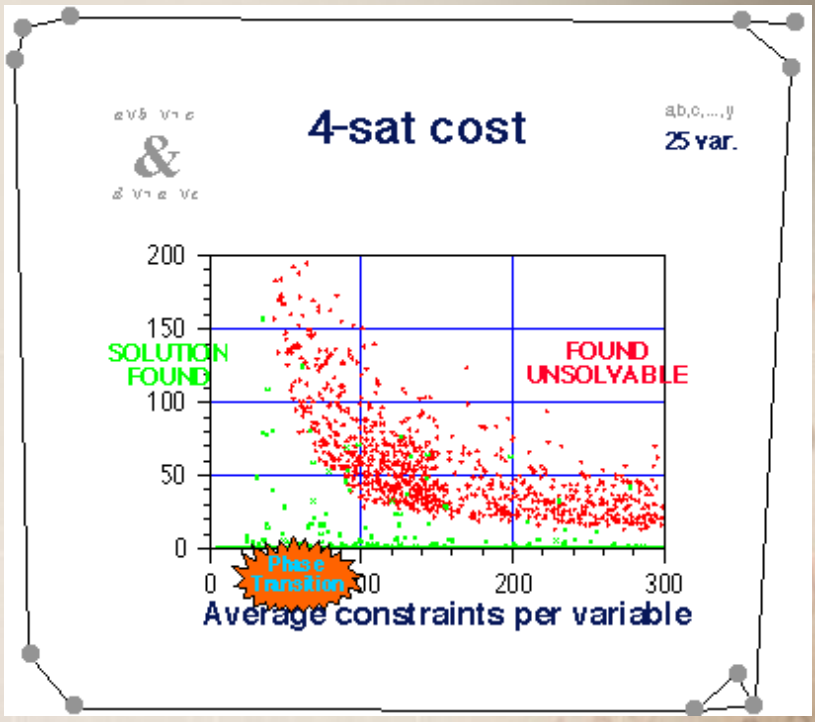

Successes of AI Discipline

Human ability reached or exceeded

- Chess (Deep Junior and X3dFritz)
- Backgammon (TD-gammon)
- Checkers
- Othello
- Many others

Successes of AI Discipline

Theorem proving

- The Robbins Conjecture, open since the 1930s, proved by EQP.
- Nqthm validated entire computer system, including circuit design, operating system, and compiler.

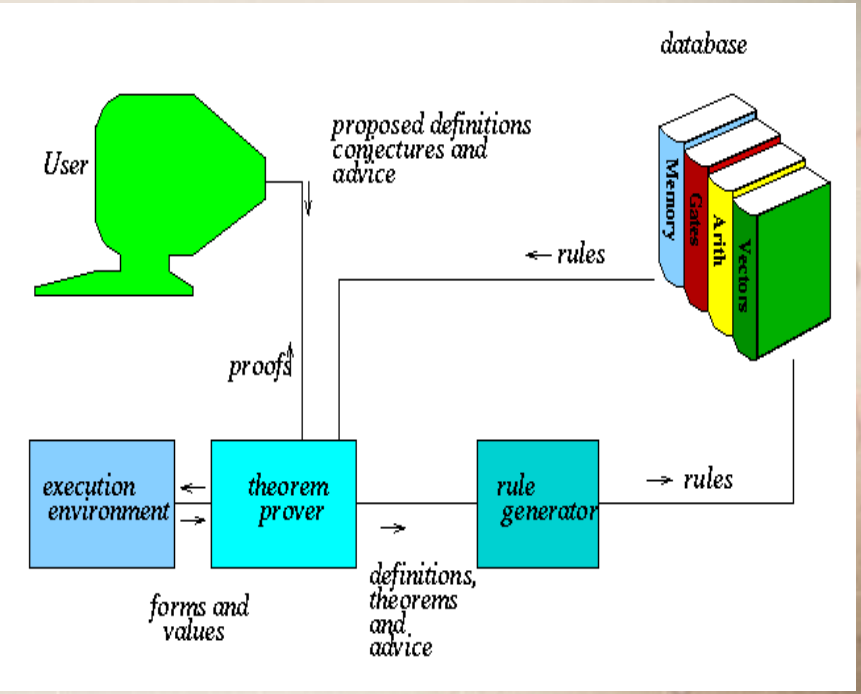

Successes of AI Discipline

NASA's autonomous robots

- Deep Space One
  - Autonomous navigation system, AutoNav, handled unanticipated system failure
- Technology
  - Automated planning and scheduling
  - Machine learning
  - Knowledge representation
  - Automated reasoning

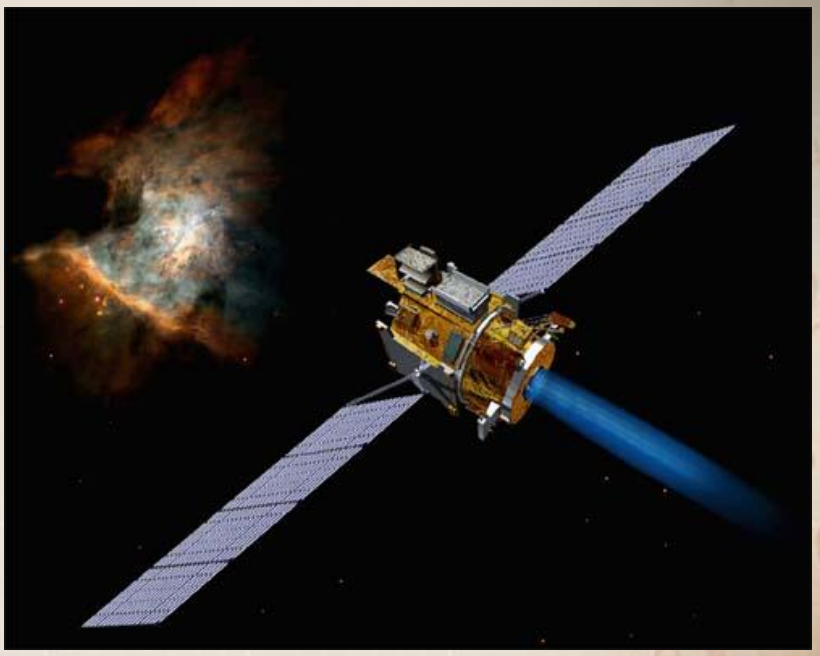

Successes of AI Discipline

Continuous, speaker independent, speech understanding

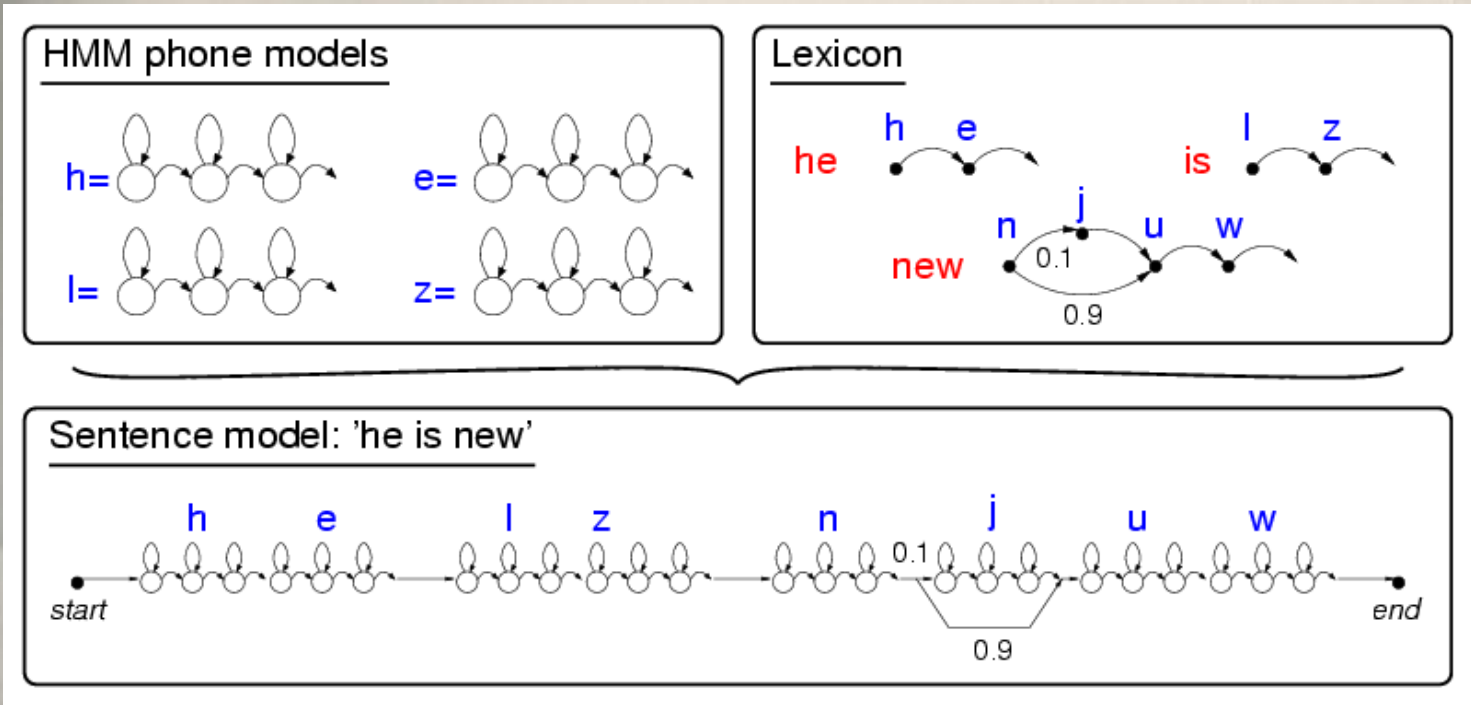

Successes of AI Discipline

Open-domain question answering

Replicable, high accuracy performance in factoid question answering, e.g.

- "How many symphonies did Shostakovich compose?"
- "What are light bulb filaments made of?"

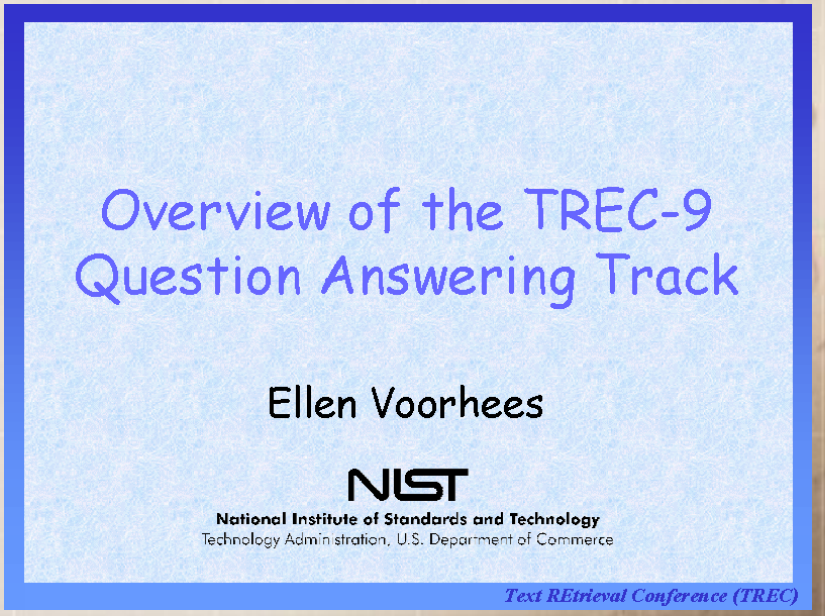

Successes of AI Discipline

Open-domain question answering

What is responsible for the success of this work?

North Carolina

North Carolina

“First In Flight”?

“First In Flight”

Wright Brother's Competitors

Alberto Santos-Dumont

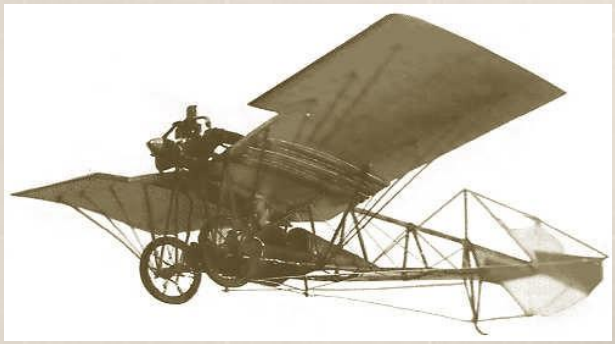

ICAIL 2005, Bologna, Italy

Wright Brother's Competitors

Richard Pearse

Richard William Pearse
Courtesy Geoff Rodliffe

Wright Brother's Competitors

Wright Brother's Competitors

James C. Mars

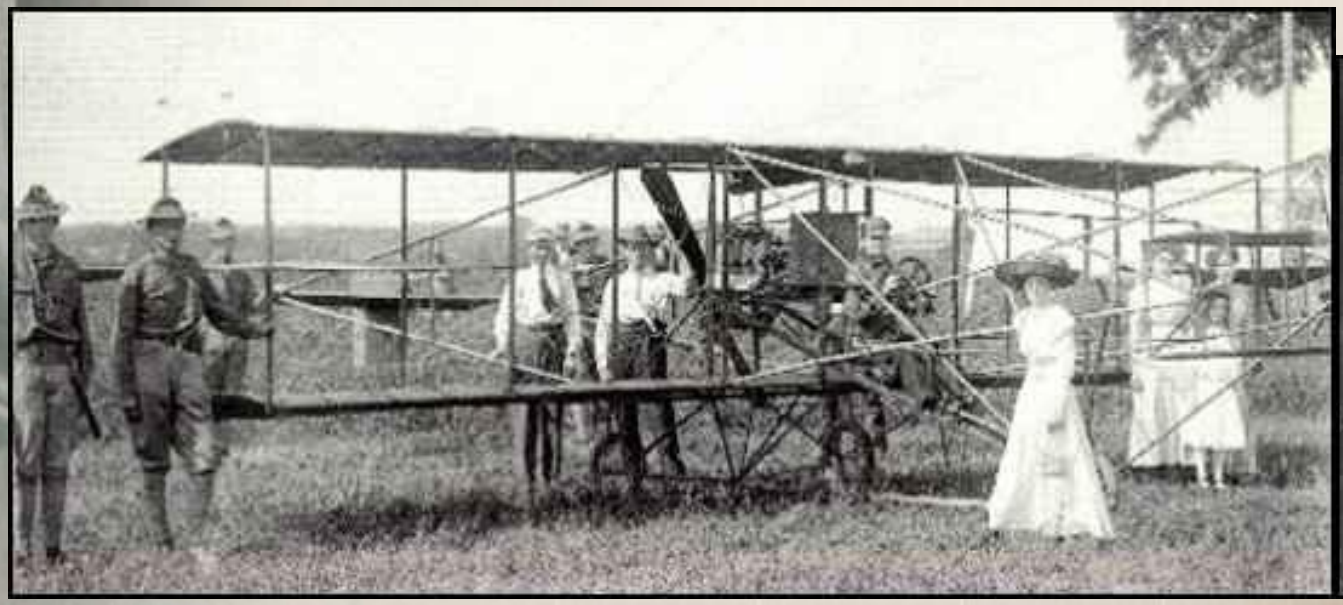

Wright Brother's Competitors

Professor Samuel Langley

Smithsonian Institute Secretary

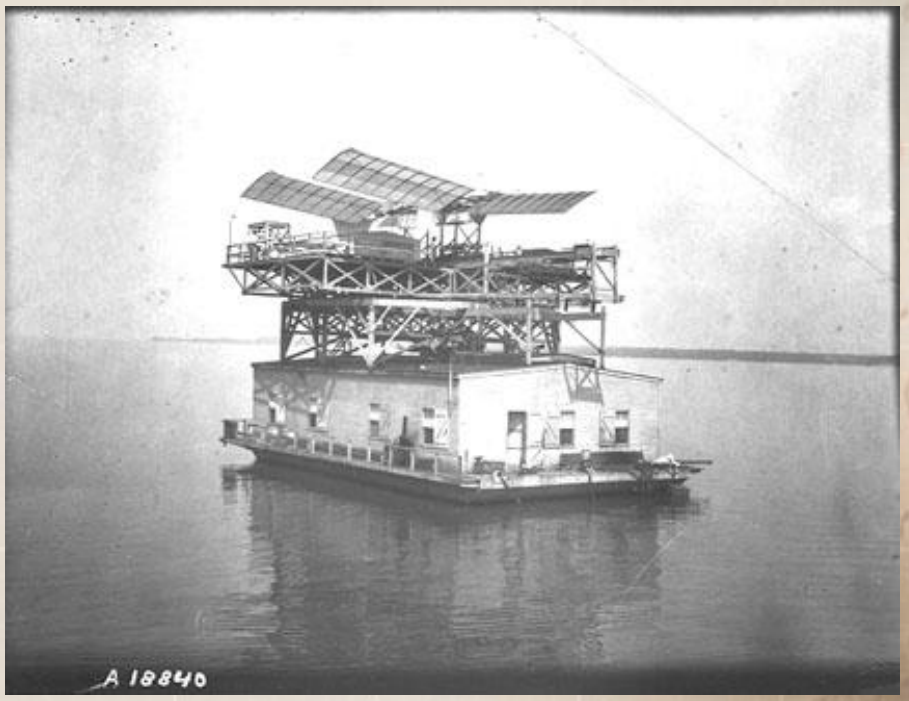

Key To Wright Brother's Success

Were resources the key?

- Langley – supported by US Department of Defense

- Orville and Wilbur Wright – supported by an Ohio bicycle shop

Key To Wright Brother's Success

Decoupling, and independently solving, key problems:

- Lift

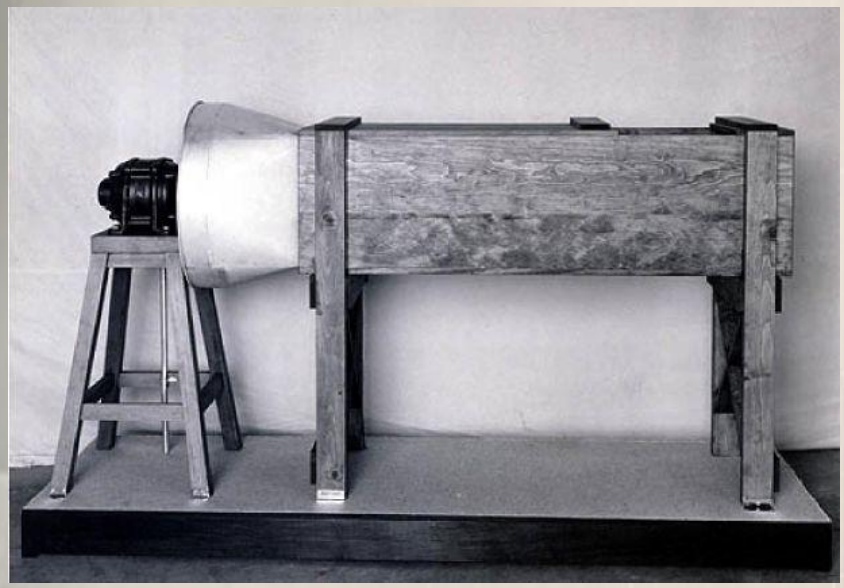

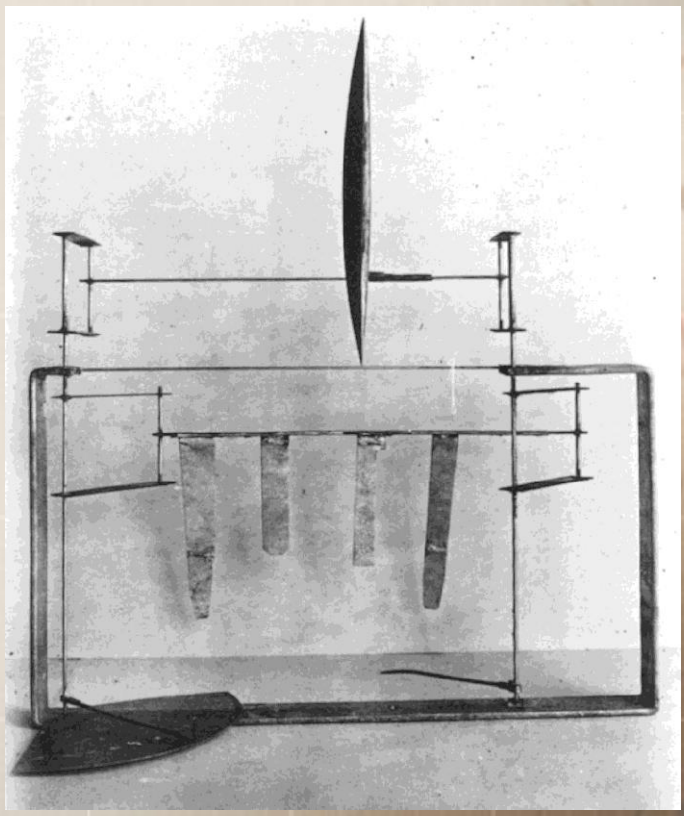

Key To Wright Brother's Success

Decoupling, and independently solving, key problems:

- Lift
- Flight control

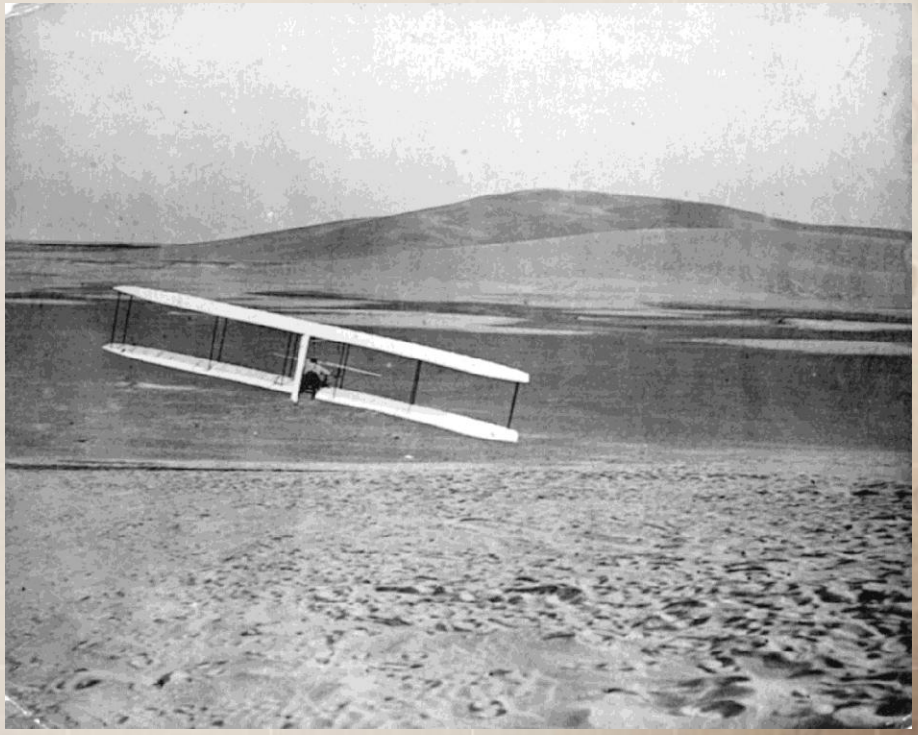

Key To Wright Brother's Success

Decoupling, and independently solving, key problems:

- Lift
- Flight control
- Propulsion

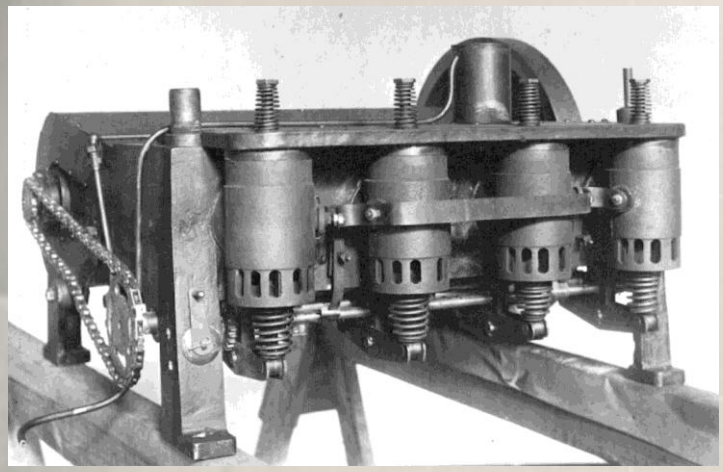

Key To Wright Brother's Success

Less successful competitors tried to solve all problems at once

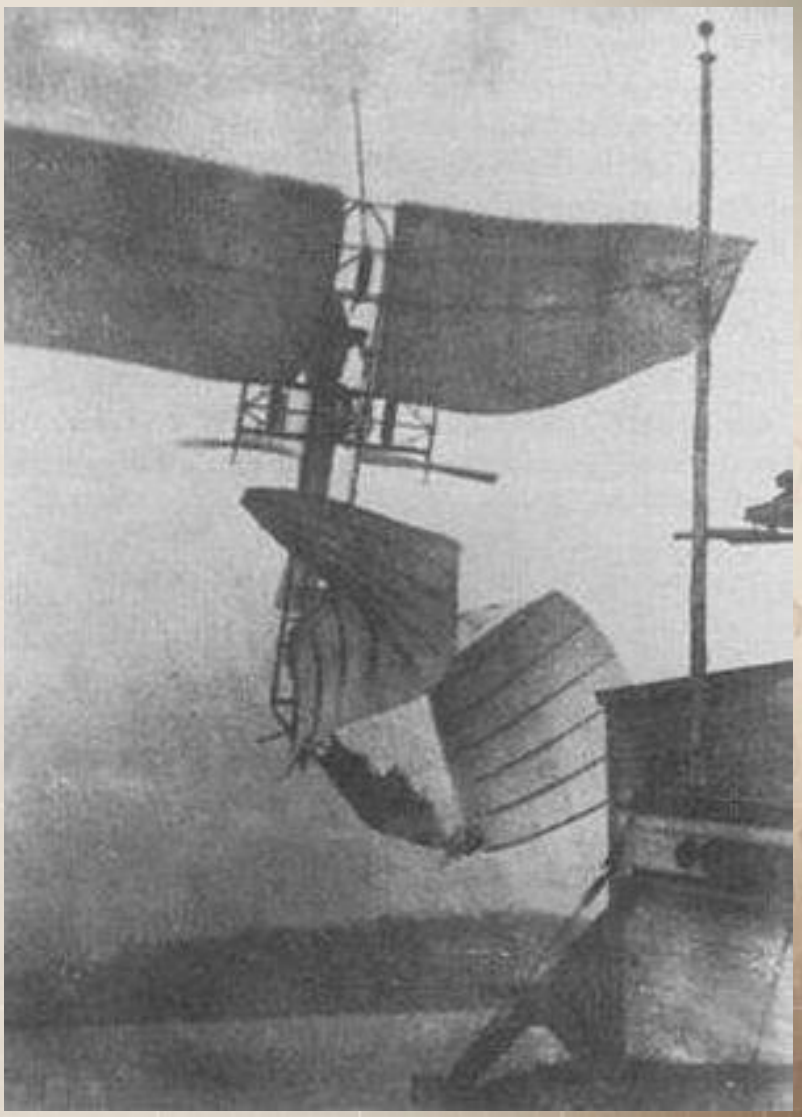

Key To Progress in Open-Domain Question Answering

Decoupling, and independently solving, key problems:

- Information retrieval
- Text segmentation
- POS tagging
- Parsing (shallow, deep)
- Word-sense disambiguation

Key To Progress in Open-Domain Question Answering

Decoupling, and independently solving, key problems:

- Named-entity recognition
- Information Extraction
- Theorem proving

Key To Progress in Open-Domain Question Answering

Each sub-problem:

- Well-defined task
- Separate theoretical and empirical evaluation criteria
- Separate literature

Key To Progress in Open-Domain Question Answering

- Task analysis
- Decomposition into well-defined sub-problems
- Theoretical analysis and reproducible empirical evaluation of proposed solutions
- Rigorous evaluation of entire system

Not so successful AI enterprises

- Circumscription (e.g. Yale shooting problem)
- Explanation-based generalization

Neither clearly tied to a specific task

Successes of AI *Community*

- Growing balkanization, but within each sub-discipline well-defined:
  - Research objectives
  - Evaluation criteria
  - Resources
- One can, and must, consult AI literature to solve AI problems

Testimonial: RealDialog, Inc.

- Customers include
  - Circuit City
  - Ford
  - WaterPik
- Dialogue architecture based on NIST TREC literature

Testimonial: RealDialog, Inc.

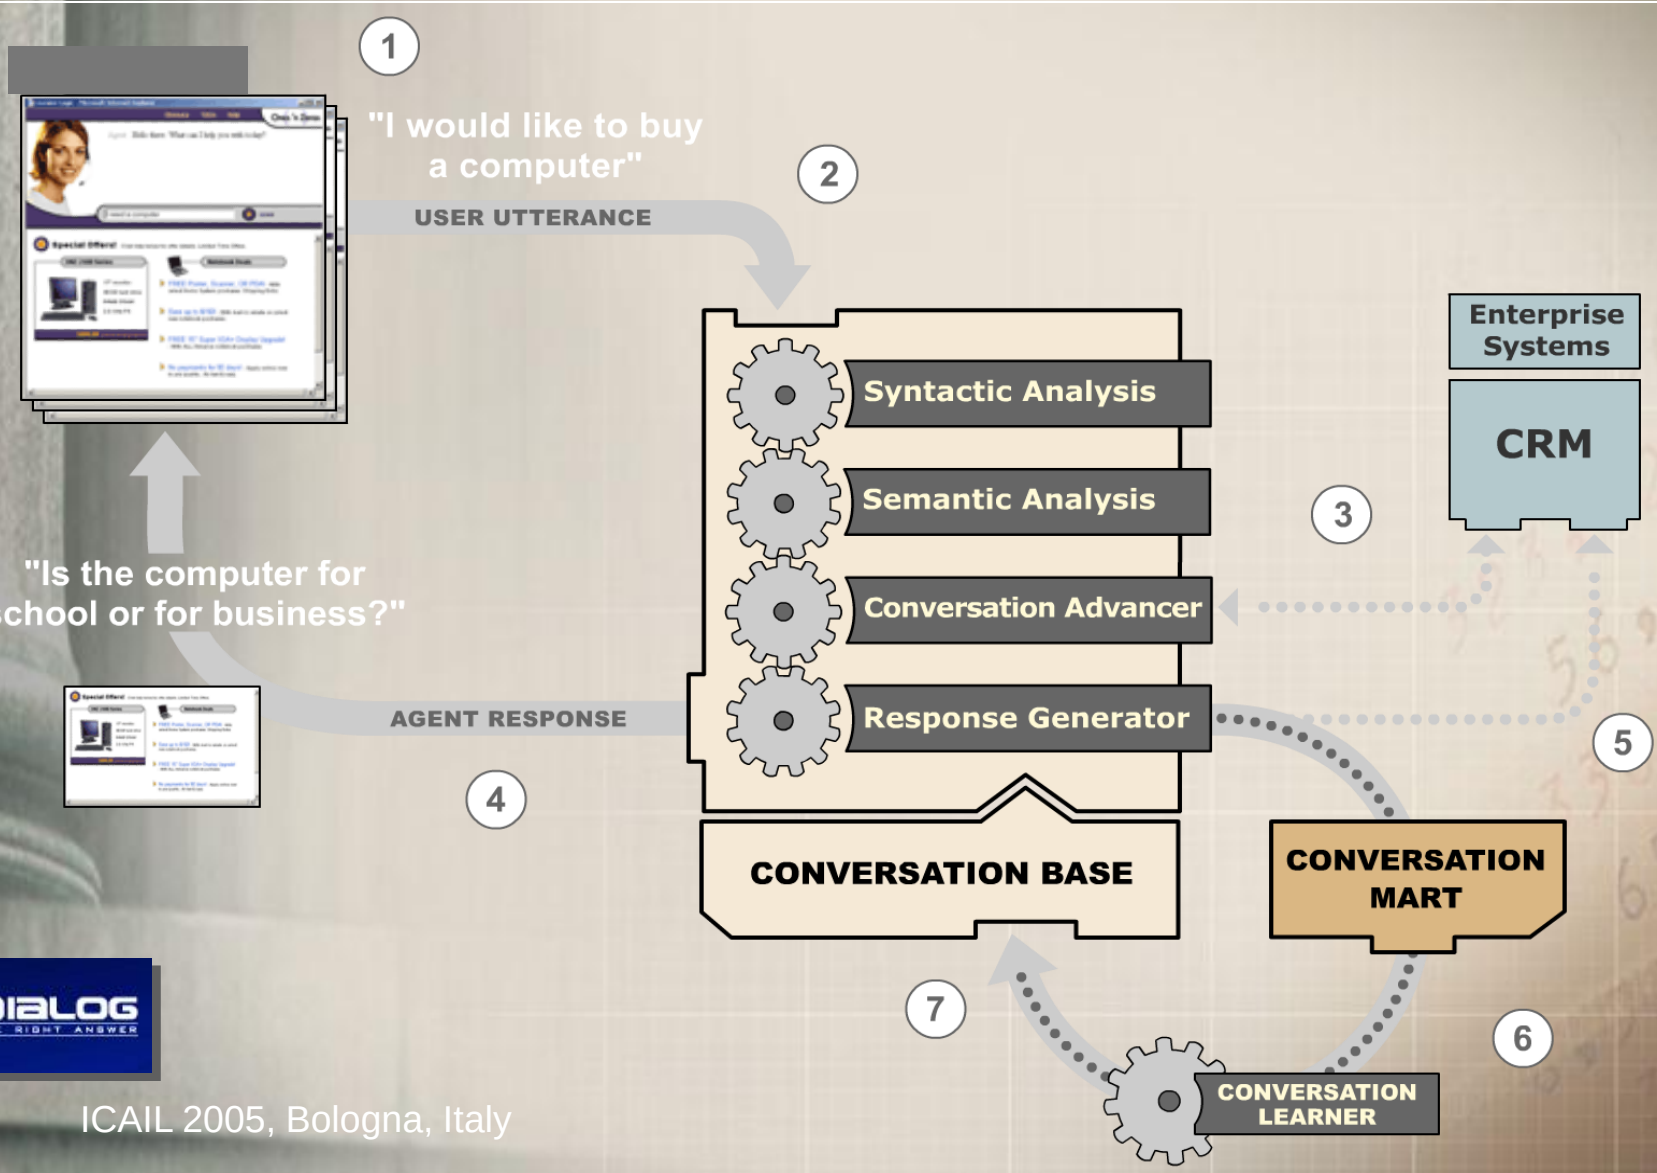

Goals of AI & Law Research

- A computational theory of legal reasoning
- Automated agents with lawyer-like interactive and problem-solving characteristics
- Practical computational tools for participants in legal system
  - Increased citizen access, understanding, and participation
  - Decreased costs
  - Increased compliance

How far have we come since ICAIL 1987?

Potential beneficiaries of AI and law technology

- Citizens
- Attorneys
- Judges
- Juries
- Clerical staff
- Legislators and regulatory rule-makers
- Scholars (e.g., legal philosophers, law professors)

How far have we come since ICAIL1987?

Citizens

- Market for routine legal advice has proliferated
- Milestone: Texas Judge Barefoot Sanders' ruling that Quicken Family Lawyer guilty of unauthorized practice of law

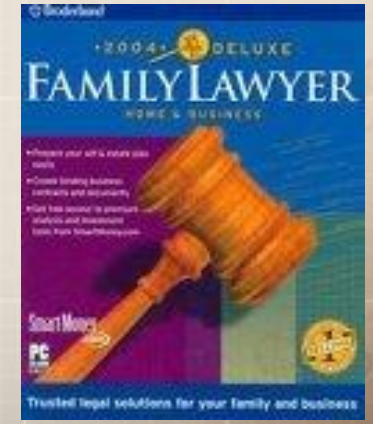

How far have we come since ICAIL1987?

Citizens

- E-government has flourished
- Many pro se litigant systems in US state courts
- Generally limited to simple rule-based systems

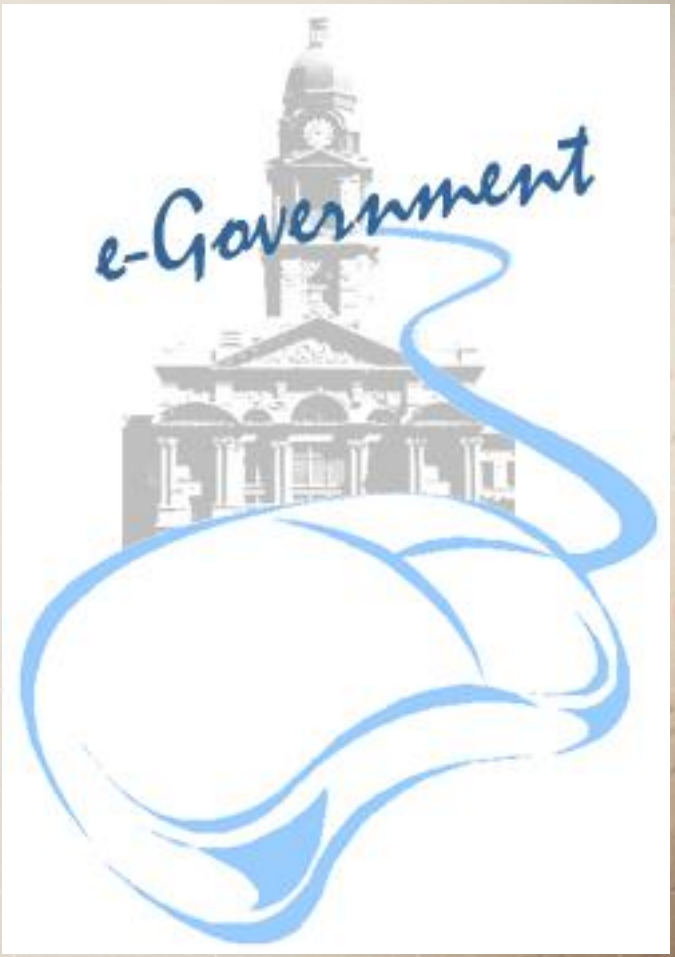

How far have we come since ICAIL1987?

Attorneys

- Reese Morrison's (ICAIL 89) barriers to acceptance of rule-based legal systems
- Some have abated:
  - Reluctance of lawyers to type
  - High expense of computers
  - Proliferation of incompatible operating systems and hardware

How far have we come since ICAIL1987?

Attorneys

- Some have diminished:
  - Incompatibility with law firm's revenue model

  - Alternative revenue models have been developed

How far have we come since ICAIL1987?

Attorneys

- Some remain:
  - High knowledge-engineering costs
  - Ignorance about legal AI systems
- Many proprietary legal expert systems used by insurance companies and large law firms.
- A significant proportion of attorneys use document-drafting software.

How far have we come since ICAIL1987?

Judges

- US Judges are late adopters with little interest in technology
- Pilot intelligent decision drafting projects have gotten nowhere
- Suspicion and resistance to sentencing systems

How far have we come since ICAIL1987?

Juries

- Research in the US indicates that juries seldom understand jury instructions
- An opportunity for tutorial or collaborative tools

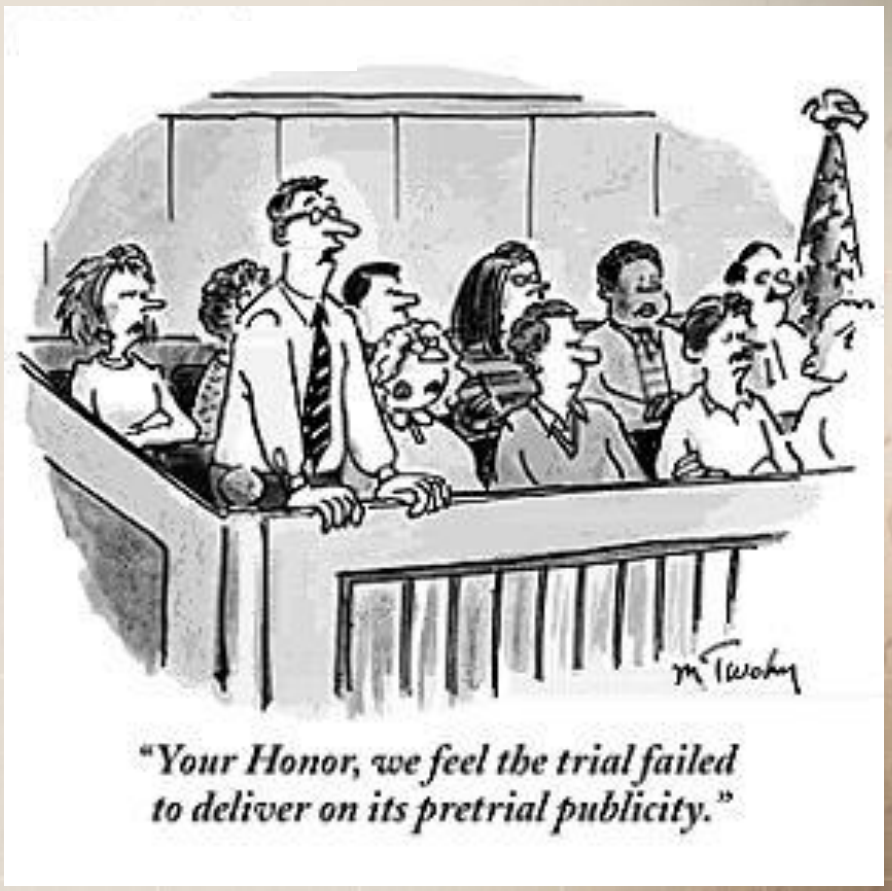

How far have we come since ICAIL1987?

Legislators and regulatory rule-makers

- Active research
  - Van Engers
  - Winkels
  - Arnold-Moore
  - Tiscornia
  - Many others …
- Limited adoption

How far have we come since ICAIL1987?

Scholars

- In US, little perceptible influence
- AI & law is not recognized as providing fruitful, novel insights into jurisprudence

Successes of AI & Law Discipline

Successes in goal 3, developing practical computational tools, at least for:

- Citizens
- Attorneys

Successes of AI & Law Research *Community*

- ICAIL is a successful forum
- Less success at standards, repositories, shared evaluation criteria
- Still less success at embodying key research results
  - Commercial development largely *independent* of AI & law literature
  - Contrast speech understanding, data mining, planning, question answering, or robotics

How has AI & law research changed since ICAIL 87?

- Number of applied papers has shrunk
- Number of argumentation papers has grown
- Task analysis papers have disappeared

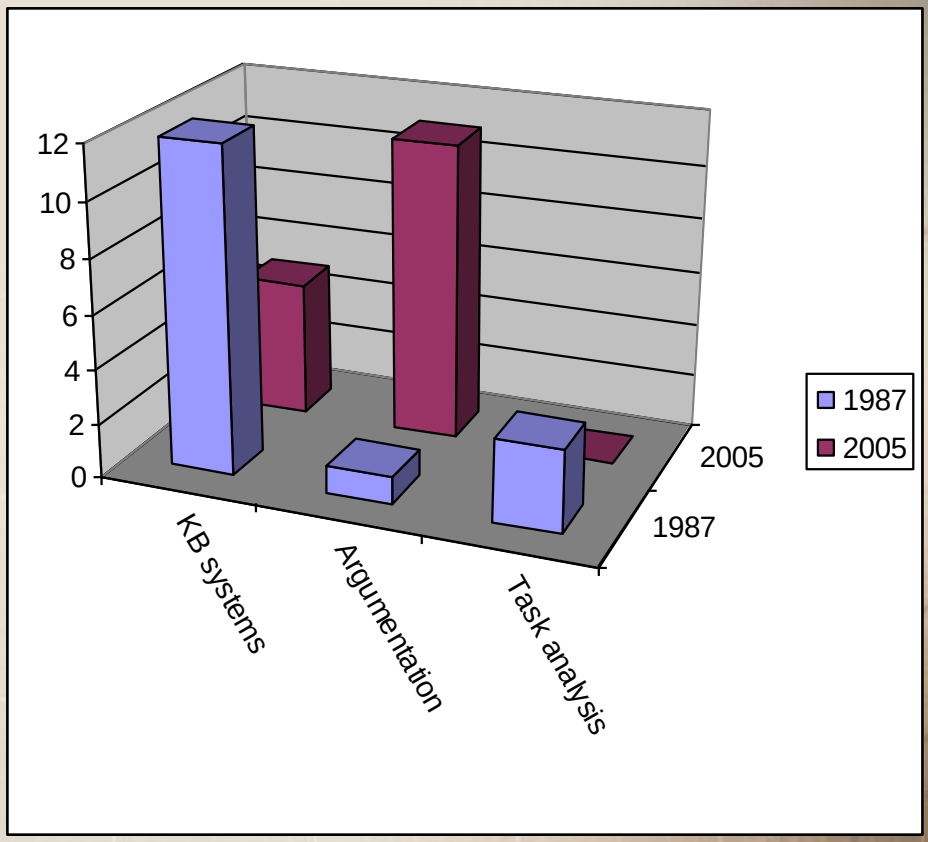

Task Analysis

Task analysis is necessary to

- Specify the I/O behavior of a successful program
- Connect formal models to actual legal discourse
- Permit comparative evaluation of alternative approaches

Task Analysis

- Formal models are not sufficient per se to specify information-processing tasks
- Example:
  - These queries have same normative model, but different task.

Query 1

“At the place where I work, our bookkeeper didn't give me my paycheck last month. Instead, she signed my name on it, cashed it, and left town. I don't know where she went. I asked my boss to give me a new check for my salary, but he said that he had already paid me once and that he didn't have to pay me again. He says that if anyone owes me the money, it is the bookkeeper. Is he right that he doesn't owe me my wages anymore?”

Query 2

"Under Article 3 of the Uniform Commercial Code, is a payer's obligation to a payee discharged by a negotiable instrument if the negotiable instrument is paid to a third party over a forged endorsement?"

Task Analysis

- Query 1 subsumes, and is much harder than, Query 2.
- Laypersons pose queries like Query 1; legal experts pose queries like Query 2;

Language and Narrative

- No implemented system can handle queries like Query 2, much less those like Query 1, much less an interactive dialogue.
- This is odd because language is as central to legal reasoning as vision is to robotics or probability is to Bayes nets.

Language and Narrative

- Legal argument is about characterizing facts, not the meaning or effect of legal norms
- Lawyers and judges believe that the meaning and effect of legal rules and cases is almost always clear
- Once the facts are fixed, the outcome of most cases is highly predictable
- Opposing lawyers therefore dispute about facts

Language and Narrative

Negotiable instruments example

- Attorney for boss would tell a story about an employee who is trying to avoid responsibility for her own carelessness
- Attorney for employee would tell a story about a boss who is trying to avoid paying an employee what she is owed
- Argument would consist of clashing narratives
- Attack and support relations among legal predicates are too obvious to be overtly discussed.

Recommendations for AI & Law Discipline

- Study actual legal discourse
- Return to task analysis (e.g., O'Neil 1987)
- Go to a computational linguistics conference – big advances since 1987
- Exploit developments in computational narrative theory

Recommendations for AI & Law Discipline

- Emulate the Wright Brothers,
  - Decompose overall task into well-defined sub-problems
  - Rigorously evaluate alternative approaches to each sub-problem
- Eschew the “not-invented-here” syndrome

Recommendations for AI & Law Community

- Formally recognize projects that have achieved independence (like IAAI).
- Develop techniques usable by commercial developers.
- Develop corpora and data repositories
- Let disinterested domain experts judge models

Conclusion

- We have come far since 1987
- Economics will make AI & law increasingly importance
- The greater AI community has important methodological lessons
- Lessons from history of aviation
  - Don’t limit yourself to lift if you also need guidance and propulsion
  - Don’t try to solve the entire problem all at once
  - Don’t mistake models for the real thing