Imperial College
London

25 years of ILP

Stephen Muggleton
Department of Computing
Imperial College, London
June, 2010

ILP biography

Infancy (1991-1994) - prodigious beginnings

- Formal basis - eg inverse resolution, saturation
- Initial implementations - eg FOIL, Golem, Clint, Linus
- Early applications - eg nite element mesh, proteins

Childhood (1995-2001) - logical development

- Rigorous theoretical foundations
- Advanced implementations, eg Progol, Aleph, Tilde
- Signi cant applications - eg mutagenesis

Teenage years (2002-2010) - indecision

- Probabilistic logic representations, eg SLPs, DTLPs
- PLL implementations - eg Prism, ProbLog
- Larger signi cant applications - eg robot scientist

Now and Next

Young adulthood (2011-2015) - action and dynamism

- Metalogical and functional extensions
- Learning actions and strategies
- Hard applications - eg systems biology

Middle adulthood (2015-2020) - socialisation

- Integration of learning, perception and action
- Learning social skills
- Hard applications - eg synthetic ecology

Human vs Statistical Learning

UK EPSRC Priority 2016-2021 - Human-like Computing

Characteristic	Human	Statistical
Examples per concept	Few (≈ 1) [Tenenbaum, 2011]	Many ($\geq 10K$)
Concepts	Many ($\geq 10K$) [Brown et al, 2008]	Few (≈ 1)
Background knowledge	Large [Brown, 2000]	Small
Structure	Modular, re-useable [Omrod et al, 2004]	Monolithic

Trends

Endogenous

- Probabilistic ILP (aging)
- Declarative learning/mining
- Higher-order, meta-intepretative, predicate invention

Exogenous

- Lifted modeling (graphical models and **beyond**)
- Semantic web, description logics & co (decline?)
- Deep learning

Deep Learning Inspirations

Tricks for predicate invention

- unlabeled data, autoencoders, denoisers, ...
- mainly in the presence of big data

Revival of NN/Logic techniques in the deep context

- KBANN, C-ILP, ...

Lifted neural networks

- Logical prescriptions for unfolding ground networks
- Co-evolution of weights

Trends at ILP 2011-2015

- Strong subfields
  1. Propositionalization
  2. Graph Mining
  3. Feature construction
  4. Hybrid architectures: SVM, Neural Nets
- Growing number of applications
  1. Robotics
  2. Actions
  3. Reinforcement learning
  4. Bioinformatics
  5. Cognition
- Ongoing work on logical foundations
- Growing work on description logics
- Probabilistic Logic Learning/Statistical Relational Learning

The future of ILP

- Stronger connection to the Semantic Web
- Exploiting Linked Open Data
- Probabilistic Logic Learning/Statistical Relational Learning
  - tractable languages
  - tractable inference (lifted)
  - scaling
- Declarative learning: Constraints, ASP
- Applications to Big Data/Scaling: clusters, GPUs

Digits and Letters

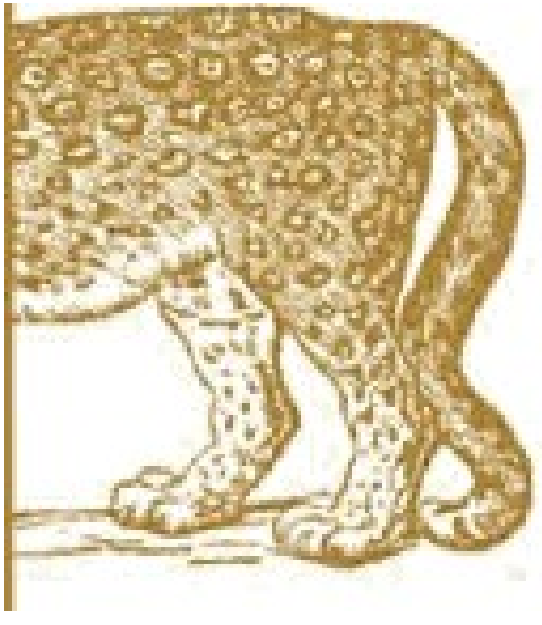

Gerson Zaverucha
Vítor Santos Costa

ILP'13

Real **hard** progress in bridging numeric and symbolic approaches:

- ProbLog is now used in real applications
- NELL has 50 million beliefs
- CP-Lint applies sophisticated search works
- Relational Dependency Networks + Boosting is the relational algorithm to beat

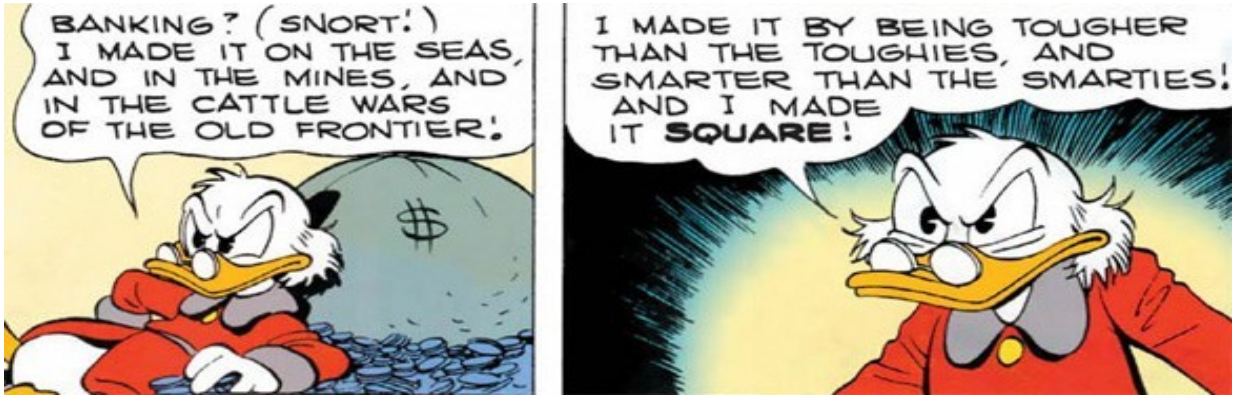

Progress in Logic:

- Abduction of Rules (Meta-Rules)
  - Probabilities
  - Predicate Invention
  - Recursion
- Learning with Negation

Looking Ahead

- Big Data?
- Propositionalization
- Parallelism
- Applications

ILP 2014: Nancy, France

Jesse Davis and Jan Ramon: PC chairs

jesse.davis@cs.kuleuven.be
KU Leuven

Trends and Outlook

- Looking back
  - Applications play a central role: Biology, medicine, robotics, natural language, vision, etc.
  - Different learning settings: Combinations of logic and probability, graphs, time and dynamics, etc.
- Going forward
  - Learn from very few examples and lots of knowledge
  - Automate ILP setup: Discover background knowledge in one domain and reuse it to solve other problems
  - Learn from both discrete and continuous data
  - Continue to focus on theory and applications
  - Make systems and data publicly available

ILP 2015

PANEL FROM VIEWPOINT OF ILP 2015

Katsumi Inoue, NII

Hayato Ohwada, Tokyo University of Science

Akihiro Yamamoto, Kyoto University

Trends

- Declarative learning: Some declarativeness everywhere in the forms of constraints, graphs, actions, kernels, ontologies, etc.
- Modeling rather than simple learning: Biology, robotics, data, natural language, cognition, vision
- New applications: logic, proof, strategy, etc.
- Diversity of approaches and topics: Continue to focus on theory and applications
- Learning with meta-theories: Meta-interpretive learning, meta-level abduction, etc.
- Learning from time-series data
- Learning from both discrete and continuous data

Outlooks

- Learning from few examples and commonsense
- Learning from state transitions with lots of fluents
- Integration of inference with high-level logical representation and recognition with low-level numerical data
- Deep relational learning—How they look like?
- General intelligence and learning

New horizon

Taisuke Sato / AIST, SONAR

- Recent technologies (word2vector, relation extraction, FreeBase,⋯) provide millions of propositions learned from big data that ILP can ex ploit:

  big data propositions ILP knowledge

- Propositions such as friend(X,Y), buys(X,Y,Z) can be non-traditional:
  - X,Y,Z are random vectors of (latent) features representing entities
  - Relations are matrixes and their truth value is computed like:
    friend(e1,e2) = σ((friendM(e1) • e2)) ∈ [0,1]
  - Inferences are made logically or possibly by matrix (tensor) operations