Uncertainty in an Unknown World

Stuart Russell
Computer Science Division, UC Berkeley

Joint work with Hanna Pasula, Brian Milch, and Bhaskara Marthi

Outline

Objects, existence, identity, uncertainty

Examples

First-order probabilistic languages

BLOG

- example BLOG models
- semantics (briefly)
- inference (briefly)

Random comments on prior knowledge

Existence, Identity, Uncertainty

The world has individual objects in it

Raw data provide propositional observations (e.g., pixels, text strings) not necessarily direct observations of the objects

Postulating *existence* of objects is the first step

Ascertaining *identity* of objects is important for:

- ♢ Accumulation of knowledge (e.g., tracking systems, databases)
- ♢ Certain relations (ownership, marriage, ...)
- ♢ Navigation (places), counting, etc.

Uncertainty is inevitable

Urns and balls

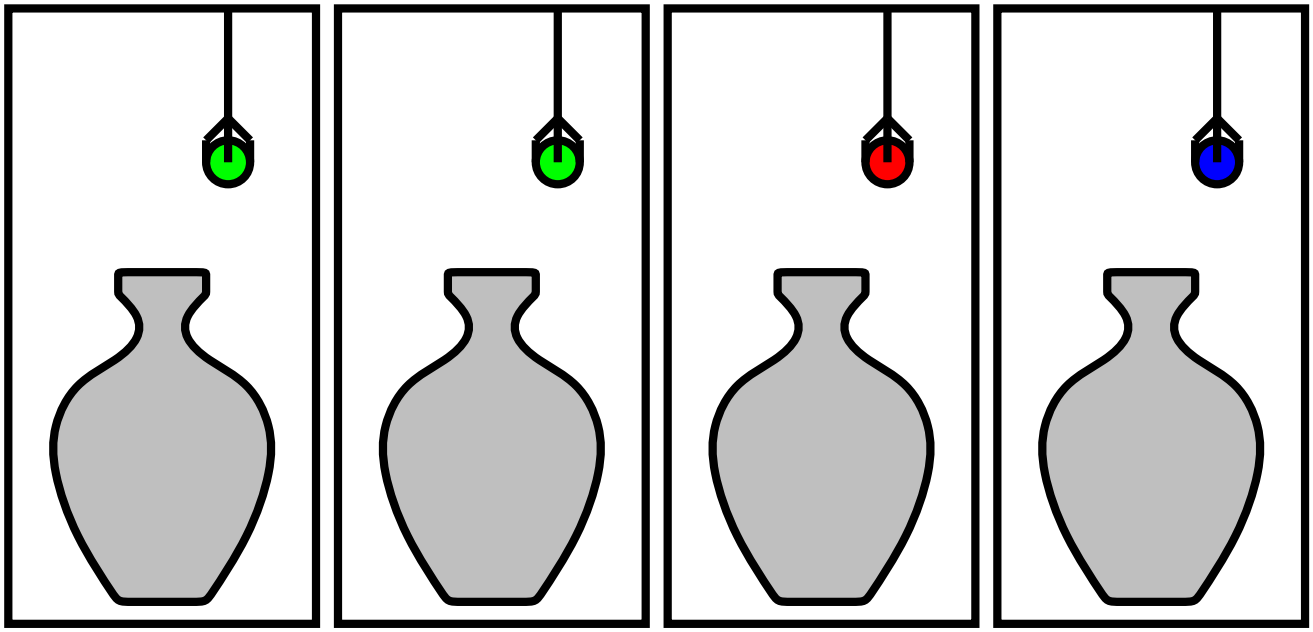

Repeat:

shake the urn, pick out one ball, examine colour, replace

How many balls are there in the urn?

Urns and balls

N balls, prior distribution $P(N)$

True colours $C_1, \ldots C_N$, identical priors $P(C_i)$

k observations, observed colours $\mathbf{O} = O_1, \ldots, O_k$

Assignment $\omega = \{\mathbf{B}_1, \ldots, \mathbf{B}_m\}$ where each $\mathbf{B}_i$ is the set of observations purportedly generated by a single real ball:

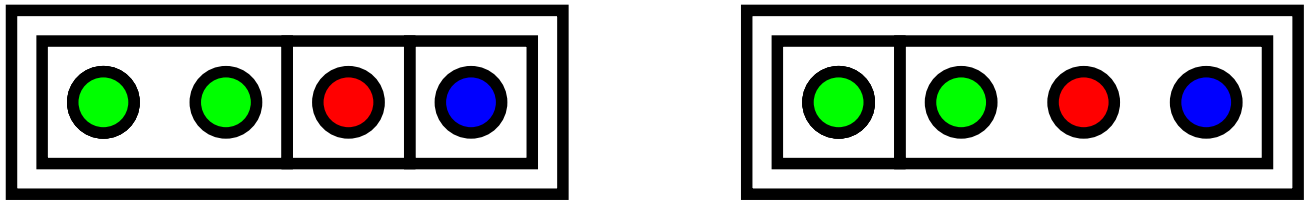

Sensor model $P(O_j \mid C_{\omega_j})$

Counting the balls

$$
\begin{aligned}
P(N \mid \mathbf{O}) &= \alpha P(N) P(\mathbf{O} \mid N) \\
&= \alpha P(N) \sum_{\omega} P(\omega \mid N) P(\mathbf{O} \mid N, \omega) \\
&= \alpha P(N) \sum_{\omega} P(\omega \mid N) P(\mathbf{O} \mid \omega) \\
&= \alpha P(N) \sum_{\omega} \frac{N!}{(N-m)!} \left(\frac{1}{N}\right)^{k} \prod_{i=1}^{m} \sum_{C_i} \prod_{O_j \in \mathbf{B}_i} P(O_j \mid C_i) P(C_i)
\end{aligned}
$$

Counting the balls contd.

Two qualitatively distinct cases:

- ♢ No identical balls
  $\Rightarrow$ converge to true N as $k \to \infty$
- ♢ Identical balls possible
  $\Rightarrow$ all multiples of minimal N possible as $k \to \infty$

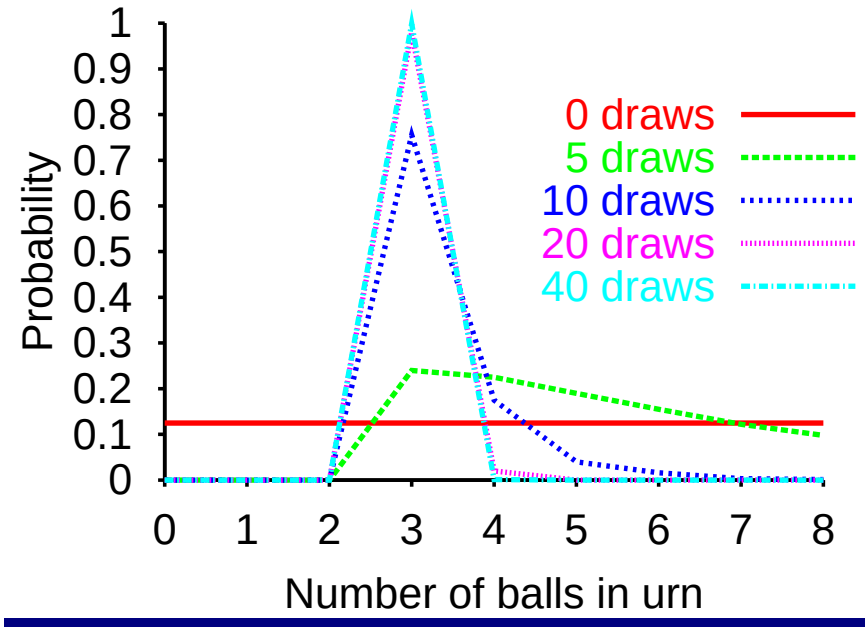

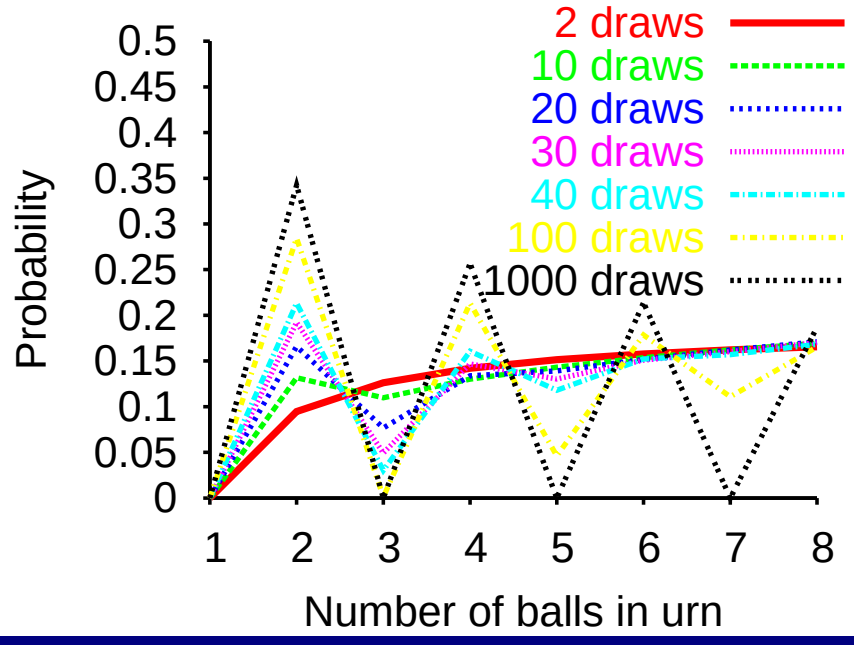

Citation matching

[Lashkari et al 94] Collaborative Interface Agents, Yezdi Lashkari, Max Metral, and Pattie Maes, Proceedings of the Twelfth National Conference on Articial Intelligence, MIT Press, Cambridge, MA, 1994.

Metral M. Lashkari, Y. and P. Maes. Collaborative interface agents. In Conference of the American Association for Artificial Intelligence, Seattle, WA, August 1994.

Are these descriptions of the same object?

This problem is ubiquitous with real data sources,
hence the *record linkage* industry

Large databases: CiteSeer

Unavailability of exact key matches can make large databases useless

`Russell w/4 Norvig`

Large databases: CiteSeer

Unavailability of exact key matches can make large databases useless

`Russell w/4 Norvig`

Russell S, Norvig P (1995) Artificial Intelligence: A Modern Approach, Prentice Hall Series in Artificial Intelligence. Englewood Cliffs, New Jersey

Stuart Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, Prentice Hall, 1995.

Russell S.; Norvig, P. Articial Intelligence - A Modern Approach. Prentice-Hall International Editions, 1995.

Russell S.J., Norvig P., (1995) Artificial Intelligence, A Modern Approach. Prentice Hall.

S. Russell and P. Norvig. Articial Intelligence, a Modern Approach. Prentice Hall, New Jersey, NJ, 1995.

Stuart Russell and Peter Norvig. Artificial intelligence: A modern approach. Prentice-Hall Series on Artificial Intelligence. Prentice-Hall, Englewood Cliffs, New Jersey, 1995.

S. Russell and P Norvig. Artifical Intelligence: a Modern Approach. Prentice Hall, 1995. Book Details from Amazon or Barnes & Noble

Stuart Russell and Peter Norvig. Articial Intelligence: A Modern Approach. Prentice Hall, 1995.

S. J. Russell and P. Norvig. Artificial Intelligence, a modern approach. Prentice Hall, Upper Saddle River, New Jersey 07458, 1995.

Stuart Russell and Peter Norvig. Artificial Intelligence. A modern approach. Prentice-Hall, 1995.

S. J. Russell and P. Norvig. Articial Intelligence: A Modern Approach. Prentice Hall. 1995.

S. Russell and P. Norvig, Artificial Intelligence A Modern Approach Prentice Hall 1995.

S. Russell and P. Norvig. **Introduction to Artificial Intelligence**. Prentice Hall, 1995.

Stuart Russell and Peter Norvig. Artficial Intelligence: A Modern Approach. Prentice-Hall, Saddle River, NJ, 1995.

Stuart Russell and Peter Norvig. Articial Intelligence a modern approach. Prentice Hall series in articial intelligence. Prentice Hall, Upper Saddle River, New Jersey, 1995.

Chapter 18 Artificial Intelligence: A Modern Approach by Stuart Russell and Peter Norvig, Prentice-Hall, 2000.

S. Russell and P. Norvig. A modern, agent-oriented approach to introductory artificial intelligence. ACM SIGART Bulletin, 6(2):24–26, 1995.

Dynamics of computational ecosystems. Physical Review A 40:404–421. Russell, S., and Norvig, P. 1995. Artificial Intelligence: A Modern Approach. Prentice Hall.

S. Russell, P. Norvig: Artificial Intelligence – A Modern Approach, Prentice Hall, 1995.

Russell, S. & Norvig, P. (1995) Artificial Intelligence: A Modern Appraoch (Englewood Cliffs, NJ: Prentice-Hall). Book Details from Amazon or Barnes & Noble

Stuart Russell and Peter Norvig. AI: A Modern Approach. Prentice Hall, NJ, 1995.

S. Russell, P. Norvig. Artificial Intelligence: A Modem Approach. Prentice- Hall, Inc., 1995.

391-414. Russell SJ, Norvig P (

Russell and Peter Norvig, "Artificial Intelligence - A Modern Approach (AIMA)", pp. 33

Stuart Russell and Peter Norvig: Artificial Intelligence: A Modern Approach, Prentice-Hall, 1994.

N.J. Nilsson. Stuart Russell and Peter Norvig artificial intelligence: a modern approach (book review). Artificial Intelligence, 82:369–380, 1996.

Russell, S. & Norvig, P., An Introduction to Artificial Intelligence: A Modern Approach, Prentice Hall International, 1996.

S. Russell, P. Norvig. Artician Intelligence. A modern approach. Prentice Hall, 1995.

Stuart Russell and Peter Norvig. Artificial Intelligence: A Modern Approach. Prentice Hall, 1995. Contributing writers: John F. Canny, Jitendra M. Malik, Douglas D. Edwards. ISBN 0-13-103805-2.

Stuart Russell and Peter Norvig. Artificial Intelligence: A Mordern Approach. Prentice Hall, Englewood Cliffs, New Jersey 07632, 1995.

In Proceedings of the Third Annual Conference on Evolutionary Programming (pp. 131–139). River Edge, NJ: World Scientific. Russell,

S.J., & Norvig, P. 1995. Artificial Intelligence, A Modern Approach. Englewood Cliffs, NJ: Prentice Hall.

John Wiley. Russell, S., & Norvig, P. (1995). Artificial Intelligence: A Modern Approach. Prentice-Hall, Inc.

Stuart Russell and Peter Norvig: Artifcial Intelligence A Modern Approach, Englewood Clioes, NJ: Prentice Hall, 1995.

In Scherer, K.R. & Ekman, P. Approaches to Emotion, 13–38. Hillsdale, NJ: Lawrence Erlbaum. Russell, S.J. and Norvig, P. 1995. Artificial Intelligent: A Modern Approach. Englewood Cliffs, NJ: Prentice Hall.

Rosales E, Forthcoming Masters dissertation, Department of Computer Science, University of Essex, Colchester UK Russell S and Norvig P (1995) Artificial Intelligence: A Modern Approach. Prentice Hall: Englewood Cliffs, New Jersey.

S. Russell and P. Norvig (1995) Artificial Intelligence; A Modern Approach, Prentice Hall, New Jersey.

S. Russell and P. Norvig. Articial Intelligence. A Modern Approach. Prentice-Hall, 1995. ISBN 0-13-360124-2.

Stuart J. Russell and Peter Norvig. Articial Intelligence: A Modern Approach, chapter 17. Number 0-13-103805-2 in Series in Articial Intelligence. Prentice Hall, 1995.

Stuart J. Russell and Peter Norvig. Articial Intelligence A Modern Approach. Prentice Hall, Englewood Cli s, New Jersey, USA, 1995. 32

Morgan Kaufmann Publishers. Russell, S., and Norvig, P. 1995. Artificial Intelligence: A Modern Approach. Prentice Hall.

Stuart J. Russell and Peter Norvig. Articial Intelligence: AModern Approach,chapter 17. Number 0-13-103805-2 in Series in Articial Intelligence. Prentice Hall, 1995.

W. Shavlik and T. G. Dietterich, eds., Morgan Kaufmann, San Mateo, CA. Russell, S. and Norvig, P. (1995). Artificial Intelligence - A Morden Approach. Englewood Cliffs, NJ: Prentice-Hall.

KeyGraph: Automatic indexing by co-occurrence graph based on building construction metaphor. In Advanced Digital Library Conference. to appear. Russell, S., and Norvig, P. 1995. Artificial Intelligence –A Modern Approach–. Prentice-Hall.

Formal derivation of rule-based programs. IEEE Transactions on Software Engineering 19(3):277–296. Russell, S., and Norvig, P. 1995. Artificial Intelligence: A Modern Approach. Prentice Hall.

Russell, Stuart and Peter Norvig, Artificial Intelligence, A Modern Approach, New Jersey, Prentice Hall, 1995.

S. Russell, P. Norvig: Articial Intelligence: A modern approach; Prentice Hall (1995).

Rechenberg, I. (89). Artificial evolution and artificial intelligence. In Forsyth, R. (Ed.), Machine Learning, pp. 83–103 London. Chapman. Russell, S., & Norvig, P. (1995). Artificial Intelligence: A Modern Approach. Prentice Hall.

Russell, S and Norvig, P. 1995. Articial Intelligence: A Modern Approach Prentice-Hall, Englewood Cli s, New Jersey, 1995.

Russell, S., & Norvig, P. (1995) . Artificial intelligence: A modern monitoring methods for information retrieval systems: From search approach. Prentice-Hall series on artificial intelligence. Upper Saddle product to search process. Journal of the American Society for Information Science, 47, 568 – 583. River, NJ: Prentice-Hall.

Stuart J. Russell and Peter Norvig. Artificial Intelligence: A Modern Approach, chapter 17. Number 0-13-103805-2 in Series in Artificial Intelligence. Prentice Hall, 1995.

S. Russell and P. Norvig. Articial Intelligence A Modern Approach. Prentice Hall, Englewood Cli s, 1995.

Russell, Stuart and Norvig, Peter: Artificial Intelligence: A Modern Approach, Prentice Hall, Englewood Cliffs NJ, 1995

S. Russell and P. Norvig. ????????? ????????????? ? ?????? ????????. Prentice Hall, Englewood Cli s, NJ, 1995.

S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach - The Intelligent Agent Book, Prentice Hall, NY, 1995.

S. Russell and P. Norvig. Artificial Intelligence-aModern Approach. Prentice Hall International, Englewood Cliffs, NJ,USA,1995.

S.J.Russell, P.Norvig: Arti cial intelligence. A modern approach", Prentice-Hall International, 1995.

In Proceedings of the Third Annual Conference on Evolutionary Programming (pp. 131–139). River Edge, NJ: World Scientific. Russell, S.J., & Norvig, P. 1995. Artificial Intelligence, A Modern Approach. Englewood Cliffs, NJ: Prentice Hall.

In Working Notes of the IJCAI-95 Workshop on Entertainment and AI/ALife, 19–24. Russell, S., and Norvig, P. 1995. Artificial Intelligence: A Modern Approach. Prentice Hall.

Stuart J. Russell and Peter Norvig. Artiilcial Intelligence: A Modern Approach. Prentice Hall, Englewood Cliffs, N J, 1995.

Academic Press. 359–380. Russell, S., and Norvig, P. 1994. Artificial Intelligence: A Modern Approach. Prentice Hall.

Russell, S., and Norvig, P. A modern, agentoriented approach to AI instruction. In Proceedings of the AAAI Fall Symposium on Innovative Instruction for Introductory AI (New Orleans, Nov. 1994), AAAI Press.

Stuart J. Russell, Peter Norvig, Artifical Intelligence: A Modern Appraoch, Prentice-Hall, Englewood Cliffs, New Jersey. 1994.

Cambridge, MA: MIT Press. Russell, S. J., and Norvig, P. (1994). Artificial Intelligence: A Modern Approach. Englewood Cliffs, NJ: Prentice-Hall.

Morgan Kauffman. Russell, S., and Norvig, P. 1994. Artificial Intelligence: A Modern Approach. Prentice Hall.

Fast Plan Generation Through Heuristic Search Russell, S., & Norvig, P. (1995). Artificial Intelligence: A Modern Approach. Prentice-Hall, Englewood Cliffs, NJ.

Hoffmann & Nebel Russell, S., & Norvig, P. (1995). Artificial Intelligence: A Modern Approach. Prentice-Hall, Englewood Cliffs, NJ.

Available on the Web as: http://cdps.umcs.maine.edu/Papers/Theses/
Russell, S. & Norvig, P. (1995). AI - A modern approach. Prentice Hall, 1995.

Stuart Russel and Peter Norvig. Artificial Intelligence: A Modern Approach, chapter 12.1 - 12.3, pages 367–380. Prentice Hall, 1995.

Stuart Russel and Peter Norvig. Artificial Intelligence, A Modern Approach. PrenticeHall, 1996. 2

Stuart Russel, Peter Norvig, Articial Intelligence: A Modern Approach, Prentice Hall, New Jersey, US, 1995

Russel, S., and Norvig, P. Articial Intelligence. A Modern Approach. Prentice Hall Series in Artificial Intelligence. 1995.

S. Russel and P. Norvig. Artificial Intelligence, A Modern Approach, Prentice Hall: 1995. Book Details from Amazon or Barnes & Noble

S. J. Russel and P. Norvig. Articial Intelligence A Modern Approach, chapter 14, pages 426-435. Prentice Hall Series in Articial Intelligence. Prentice Hall International, Inc., London, UK, rst edition, 1995. Exercise 14.3.

Russel, S. and P. Norvig. Articial intelligence: A modern approach, Prentice Hall, 1995. Book Details from Amazon or Barnes & Noble

S. Russel and P. Norvig Artificial Intelligence: A Modern Approach, MIT Press 1995.

Russel, S. and Norvig, P., "Artificial Intelligence: A Modern Approch," p. 111-114, Prentice-Hall.

J. Russel and P. Norvig. Artificial Intelligence, A Modern Approach. Prentice Hall, Upper Saddle River, NJ, 1995. 71

Stuart Russel and Peter Norvig. A Modern, Agent-Oriented Approach to Introductory Artificial Intelligence. 1995.

Stuart J. Russel and Peter Norvig. Artificial Intelligence—A Modern Approach, chapter 14, pages 426–435. Prentice Hall Series in Artificial Intelligence. Prentice Hall Internationall, Inc., London, UK, first edition, 1995. Excersice 14.3.

Russel S. and Norvig P. (1995). Articial Intelligence. A Modern Approach. Prentice Hall Series in Artificial Intelligence.

S. Russel, P. Norvig Articial Intelligence - A Modern Approach Prentice Hall, 1995

Russel, S., P. Norvig. Artificial Intelligence: A Modern Approach Prentice Hall 1995.

Artificial Intelligence, S Russel & P Norvig, Prentice Hall, 1995 21

Russel, S.J, Norvig P: Artificial Intelligence. A Modern Approach, Prentice Hall Inc. 1995

Russel, S., Norvig, P. (1995) Artificial Intellience - A modern approach. (Englewood Cliffs: Prentice Hall International).

Data association: classical

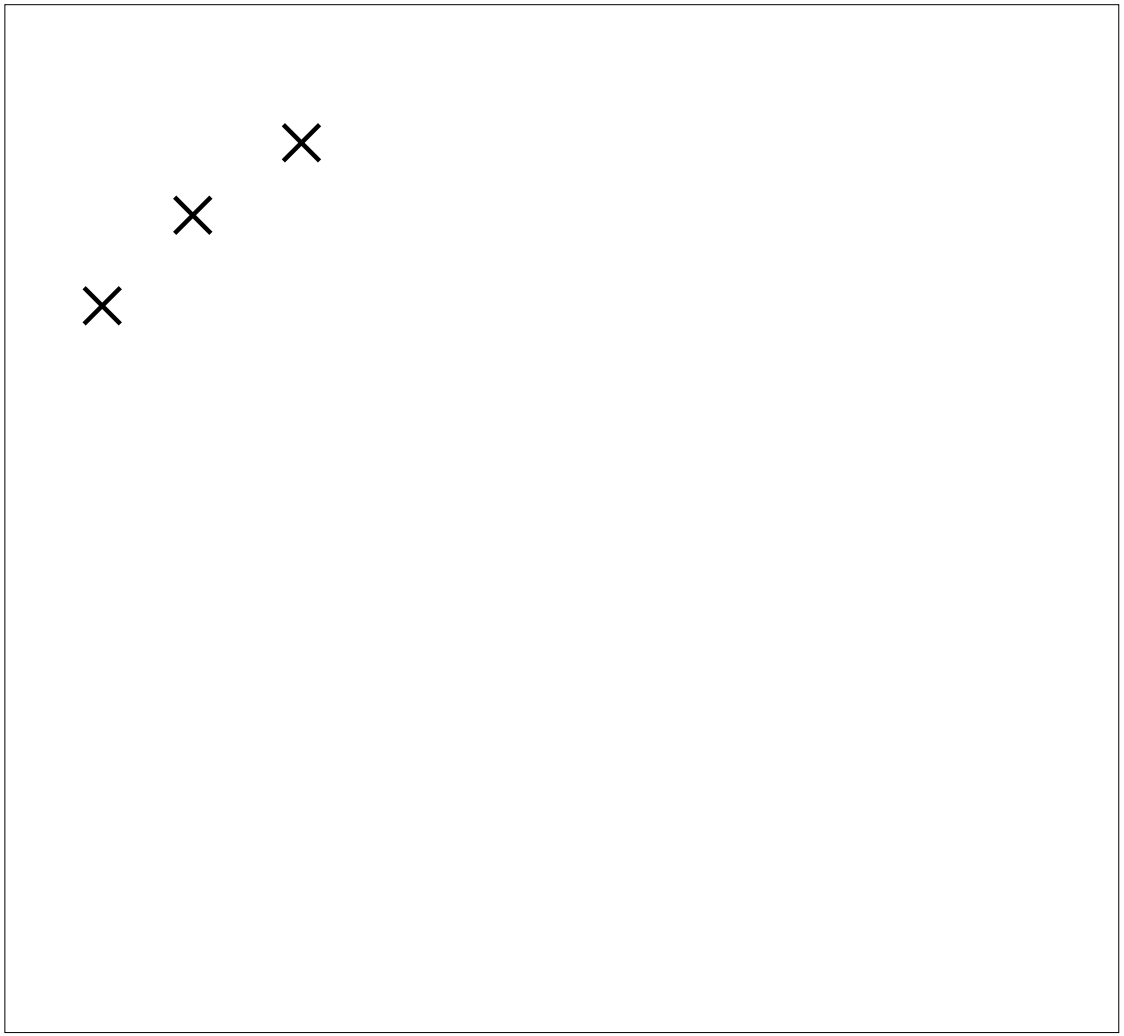

Data association: classical

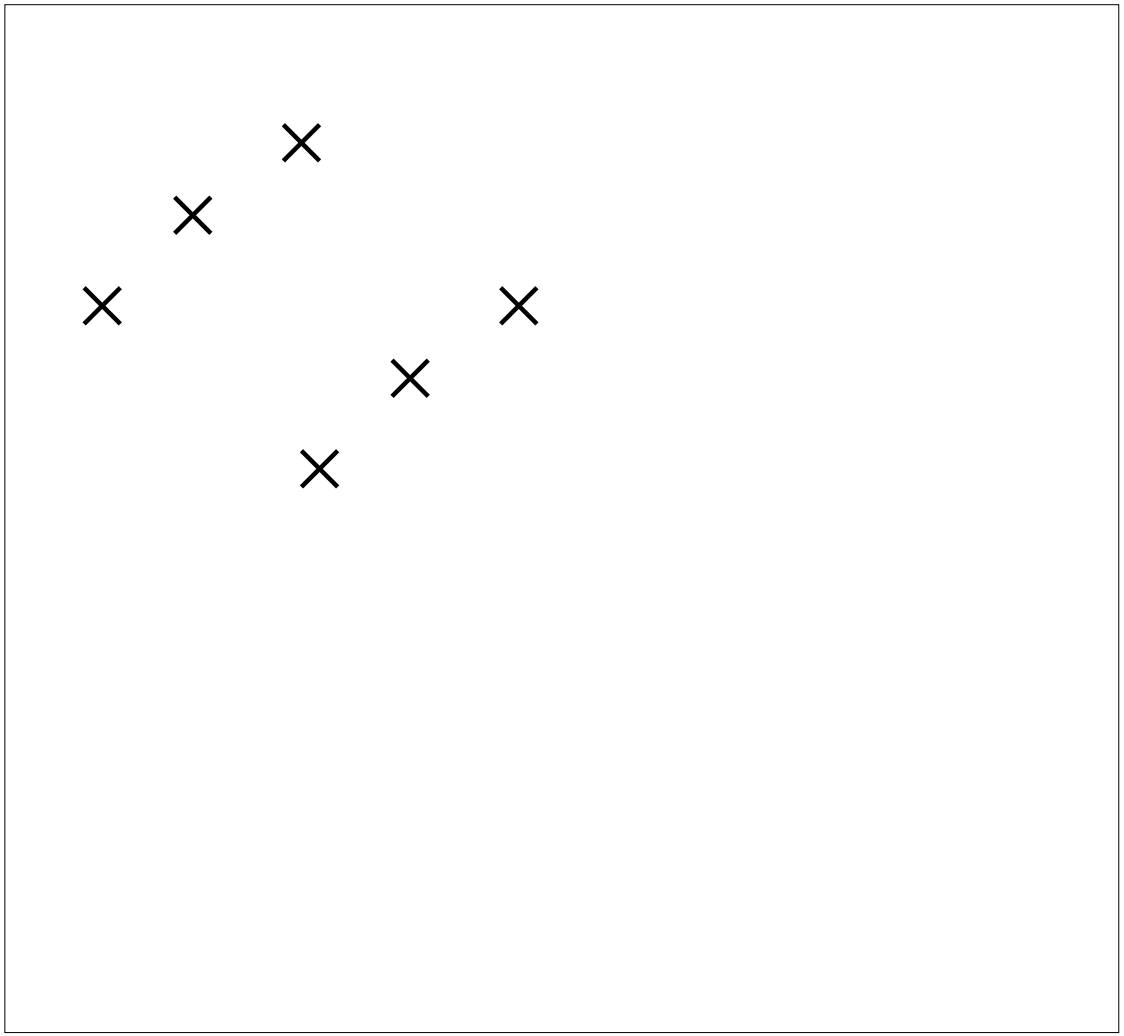

Data association: classical

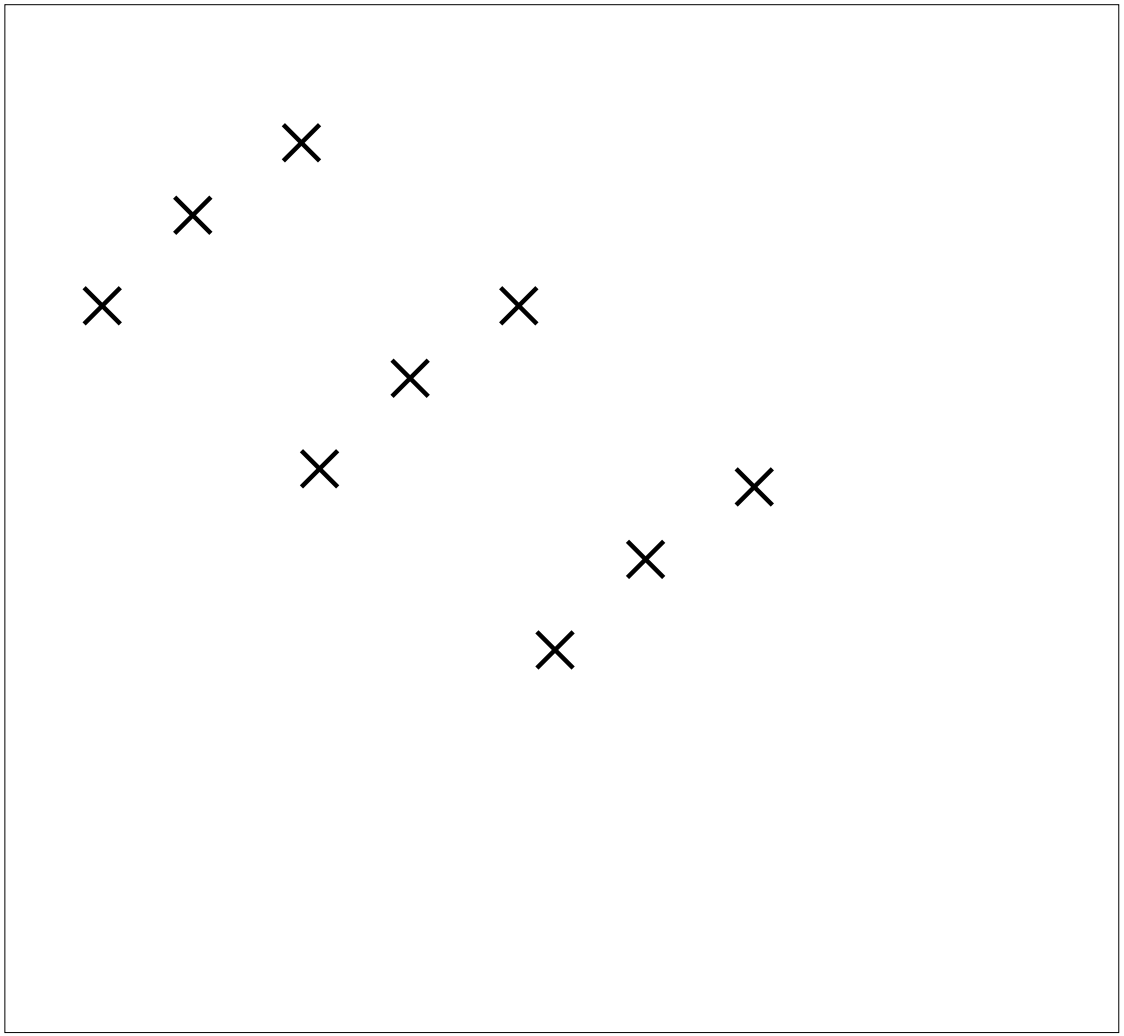

Data association: classical

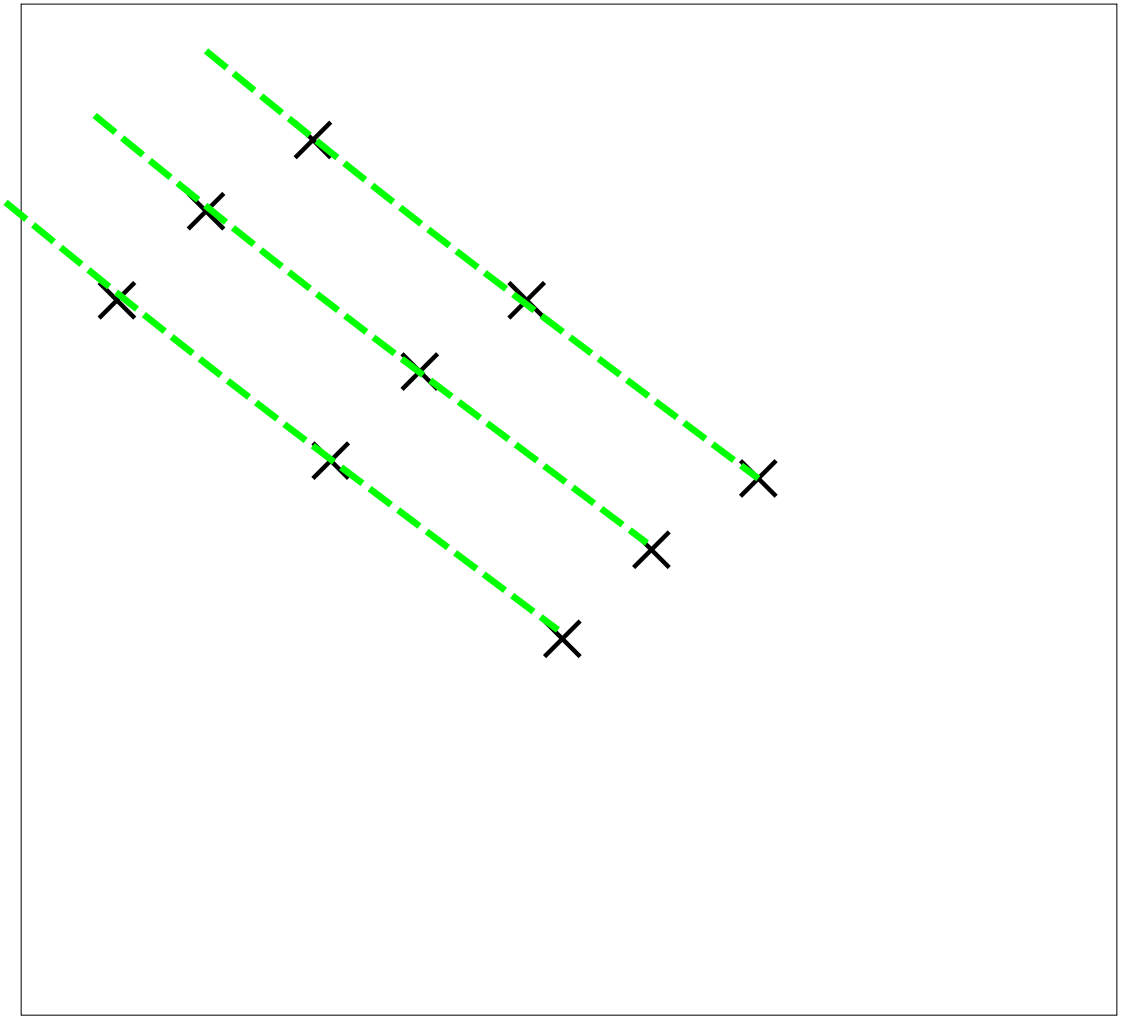

Data association: classical

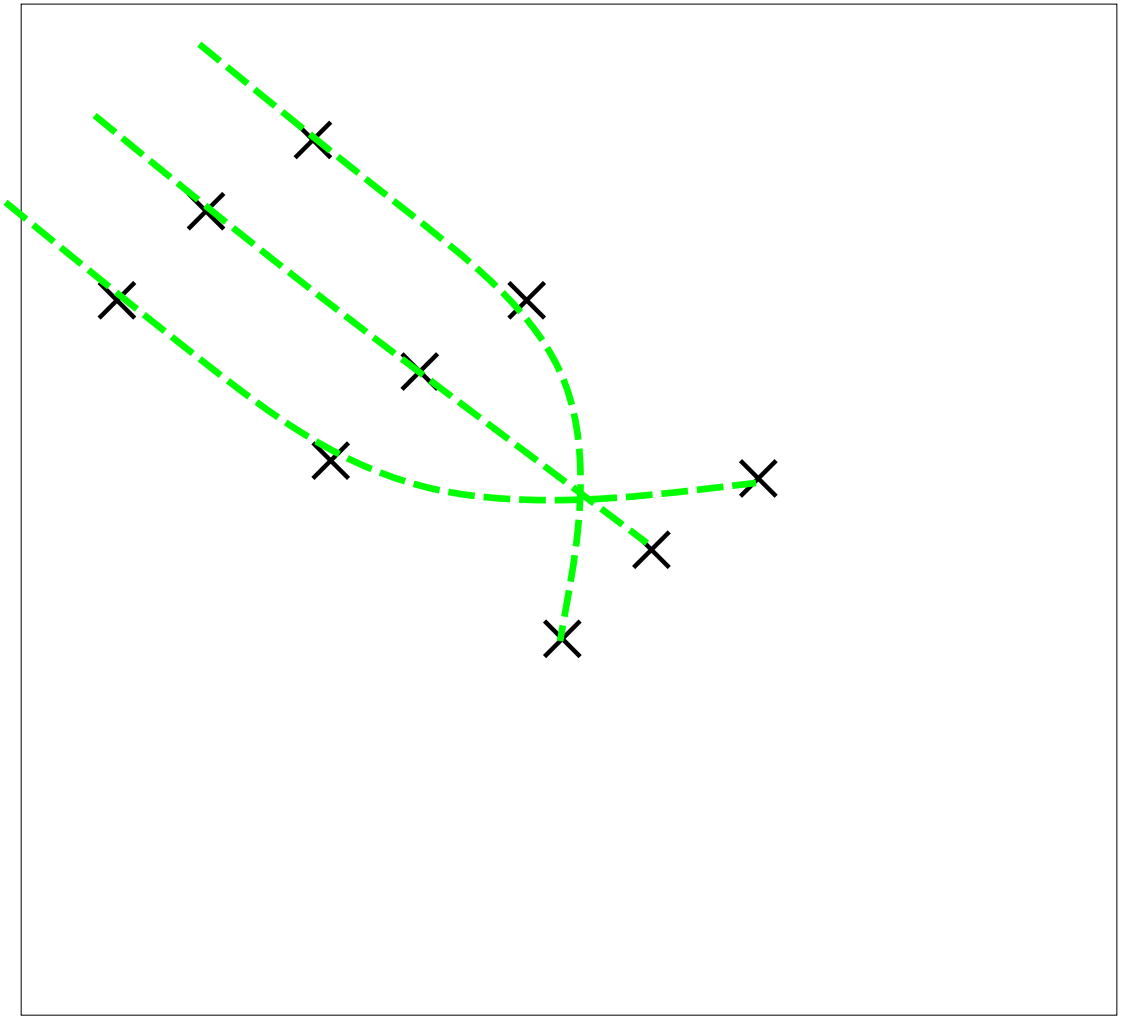

Data association: classical

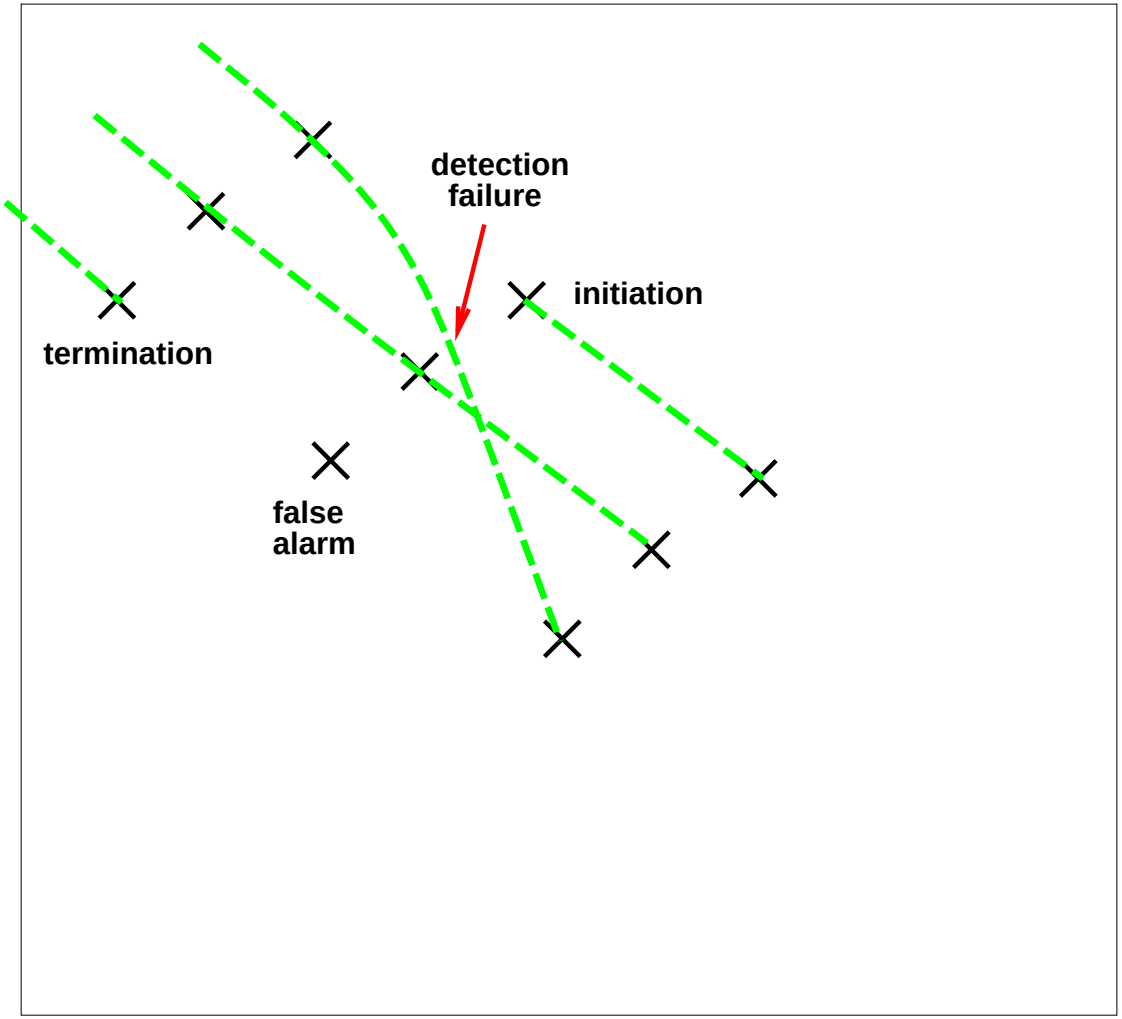

First-order probabilistic languages

(Gaifman, 1964; Halpern, 1990; lots of recent papers)

Idea: KB defines probability measure μ over all possible logical worlds

Idea [PR 01]: FOPL inference by sampling possible logical worlds

- ♢ Not easy to define consistent and complete theories
- ♢ Most FOPLs assume unique names and domain closure:
  possible world = *Herbrand model on fixed object set*
- ♢ Real world violates unique names and domain closure
- ♢ BLOG [MMR 04] for open worlds:
  possible world = *object sets + relation structure + symbol mapping*

BLOG models

Guaranteed objects are known to exist:

$\{\text{Color } c\} : \quad = \{\text{Red}, \text{Blue}\}$

Also includes built-in types (integers, strings, reals, etc.)

Number statements specify distributions for objects that might exist:

$\#\{\text{Ball } b\} : \quad \sim \mathit{Uniform}(\{1, \ldots, 8\})$

Can also depend on existence of other objects

Dependency statements specify conditional distributions for predicates and functions

$\text{TrueColor}(b) : \quad \sim \mathit{Uniform}(\{\text{Color } c\})$
$\text{BallDrawn}(d) : \quad \sim \mathit{Uniform}(\{\text{Ball } b\})$
$\text{ObservedColor}(d) : \quad \sim \mathit{ColorSensorModel}(\text{TrueColor}(\text{BallDrawn}(d)))$

BLOG model for citations

#{Researcher r}:
    $\sim$ *NumResearchersDistrib*()
Name(r):
    $\sim$ *NamePrior*()
#{Publication p}:
    $\sim$ *NumPubsDistrib*({Researcher r})
NumAuthors(p):
    $\sim$ *NumAuthorsDistrib*()
Author(p, i):
    `if` $i <$ NumAuthors(p)
        `then` $\sim$ *SampleUnused*({Reseacher r},
                {Author(p, j) : $j < i$})

BLOG model for citations, contd.

Title(p):
    $\sim$ *TitlePrior*()
PubCited(c):
    $\sim$ *Uniform*({Publication p})
NameAsCited(c, i):
    `if` $i <$ NumAuthors(PubCited(c))
        `then` $\sim$ *NameObs*(Name(Author(PubCited(c), i)))
TitleAsCited(c):
    $\sim$ *TitleObs*(Title(PubCited(c)))
CitString(c):
    $\sim$ *CitDistrib*({NameAsCited(c, i) : $i <$ NumAuthors(PubCited(c))},
                TitleAsCited(c))

BLOG model for aircraft tracking

#{AirBase}:
    $\sim$ *NumBasesDistrib*()
Location(b):
    $\sim$ *UniformLocation*()
#{Aircraft : HomeBase $\mapsto b$}:
    $\sim$ *NumAircraftDistrib*()
TakesOff(a, t):
    `if` $t > 0 \wedge \neg$InFlight($a, t-1$)
        `then` $\sim$ *TakeoffBernoulli*()
Lands(a, t):
    `if` $t > 0 \wedge$ InFlight($a, t-1$)
        `then` $\sim$ *LandingDistrib*(State($a, t-1$),
                Location(Dest($a, t-1$)))

BLOG model for aircraft contd.

CurBase(a, t):

    if $t = 0$ then = HomeBase(a)
    elseif TakesOff(a, t) then = null
    elseif Lands(a, t) then = Dest($a, t-1$)
    else = CurBase($a, t-1$)

InFlight(a, t):

    = (CurBase(a, t) = null)

State(a, t):

    if TakesOff(a, t)
        then $\sim$ *InitState*(Location(CurBase($a, t-1$)))
    elseif InFlight(a, t)
        then $\sim$ *StateTransition*(State($a, t-1$),
                                      Location(Dest(a, t)))

BLOG model for aircraft contd.

Dest(a, t):
    `if` TakesOff(a, t)
        `then` $\sim$ *UniformChoice*({AirBase a})
    `elseif` InFlight(a, t)
        `then` = Dest($a, t-1$)
#{RadarBlip : BlipSource $\mapsto a$, BlipTime $\mapsto t$)}:
    `if` InFlight(a, t) `then` $\sim$ *NumDetectionsDistrib*()
#{RadarBlip : BlipSource $\mapsto$ null, BlipTime $\mapsto t$}:
    $\sim$ *NumFalseAlarmsDistrib*()
ApparentPos(r):
    `if` BlipSource(r) = null
        `then` $\sim$ *FalseDetectionDistrib*()
    `else` $\sim$ *ObsDistrib*(State(BlipSource(r), BlipTime(r)))

Semantics

Possible worlds defined by

- possibly infinite sets of objects for each type,
  each tagged by generation record (if any)
  e.g., $Publication_{12}$ from $Researcher_3$
- a relation over the objects for each random predicate symbol
- a function over the objects for each random function symbol
- an object for each random constant symbol

Theorem: every well-formed BLOG KB defines a unique measure over possible worlds

Well-formed BLOG models

Dependency structure specified by BLOG model must

- have no infinite receding chains
- have finite parent contexts (but may have infinite parent sets)
- be *conditionally acyclic*

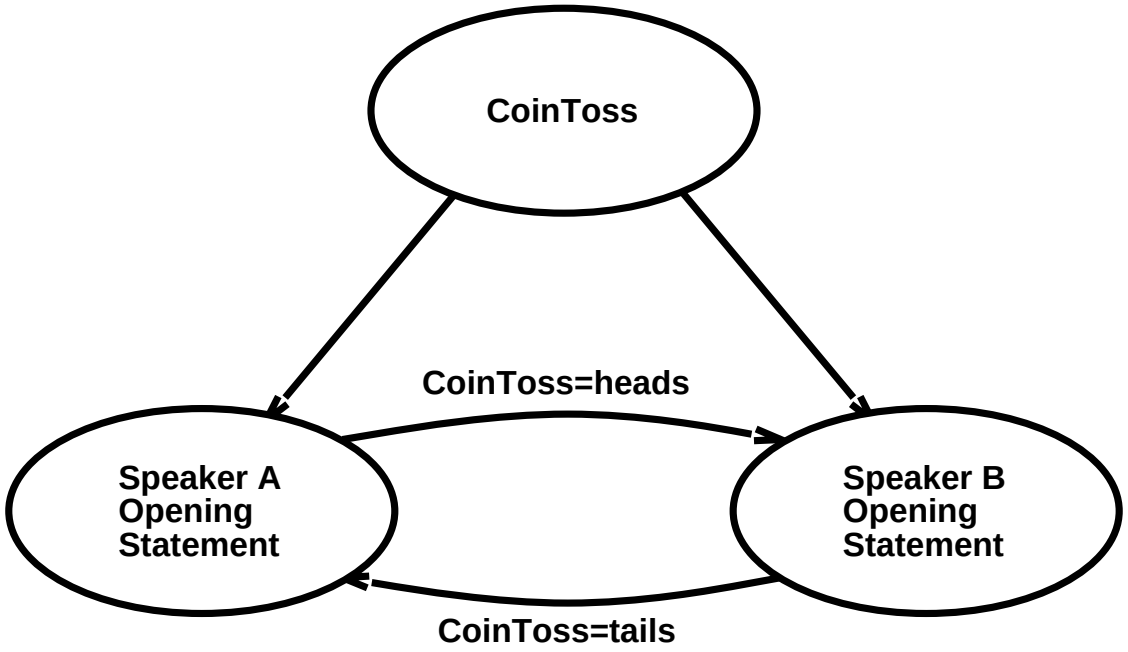

Inference

Completely generic importance sampling and MCMC algorithms based on query/evidence ancestor search

Theorem: Algorithms are consistent and have finite time per sample, for any well-formed BLOG model

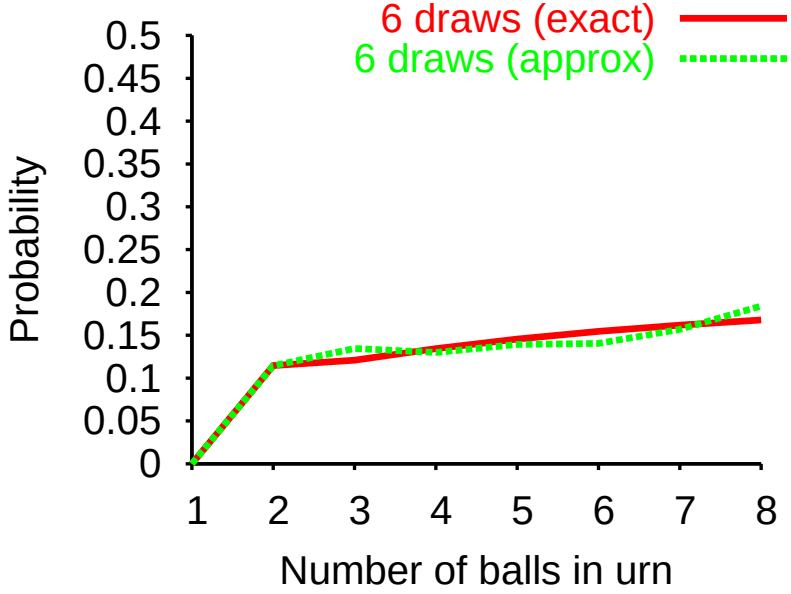

Application of MCMC to FOPLs

Markov chain states are possible worlds

Proposals are modifications to worlds $\omega \longrightarrow \omega'$

e.g., change relations between objects, merge objects, etc.

Transition ratio $\pi(\omega')/\pi(\omega)$ computable from KB,
ideally by incremental modification

Computing $f(\omega) =$ evaluating arbitrary query sentence on a world,
ideally by incremental modification

MCMC on possible worlds

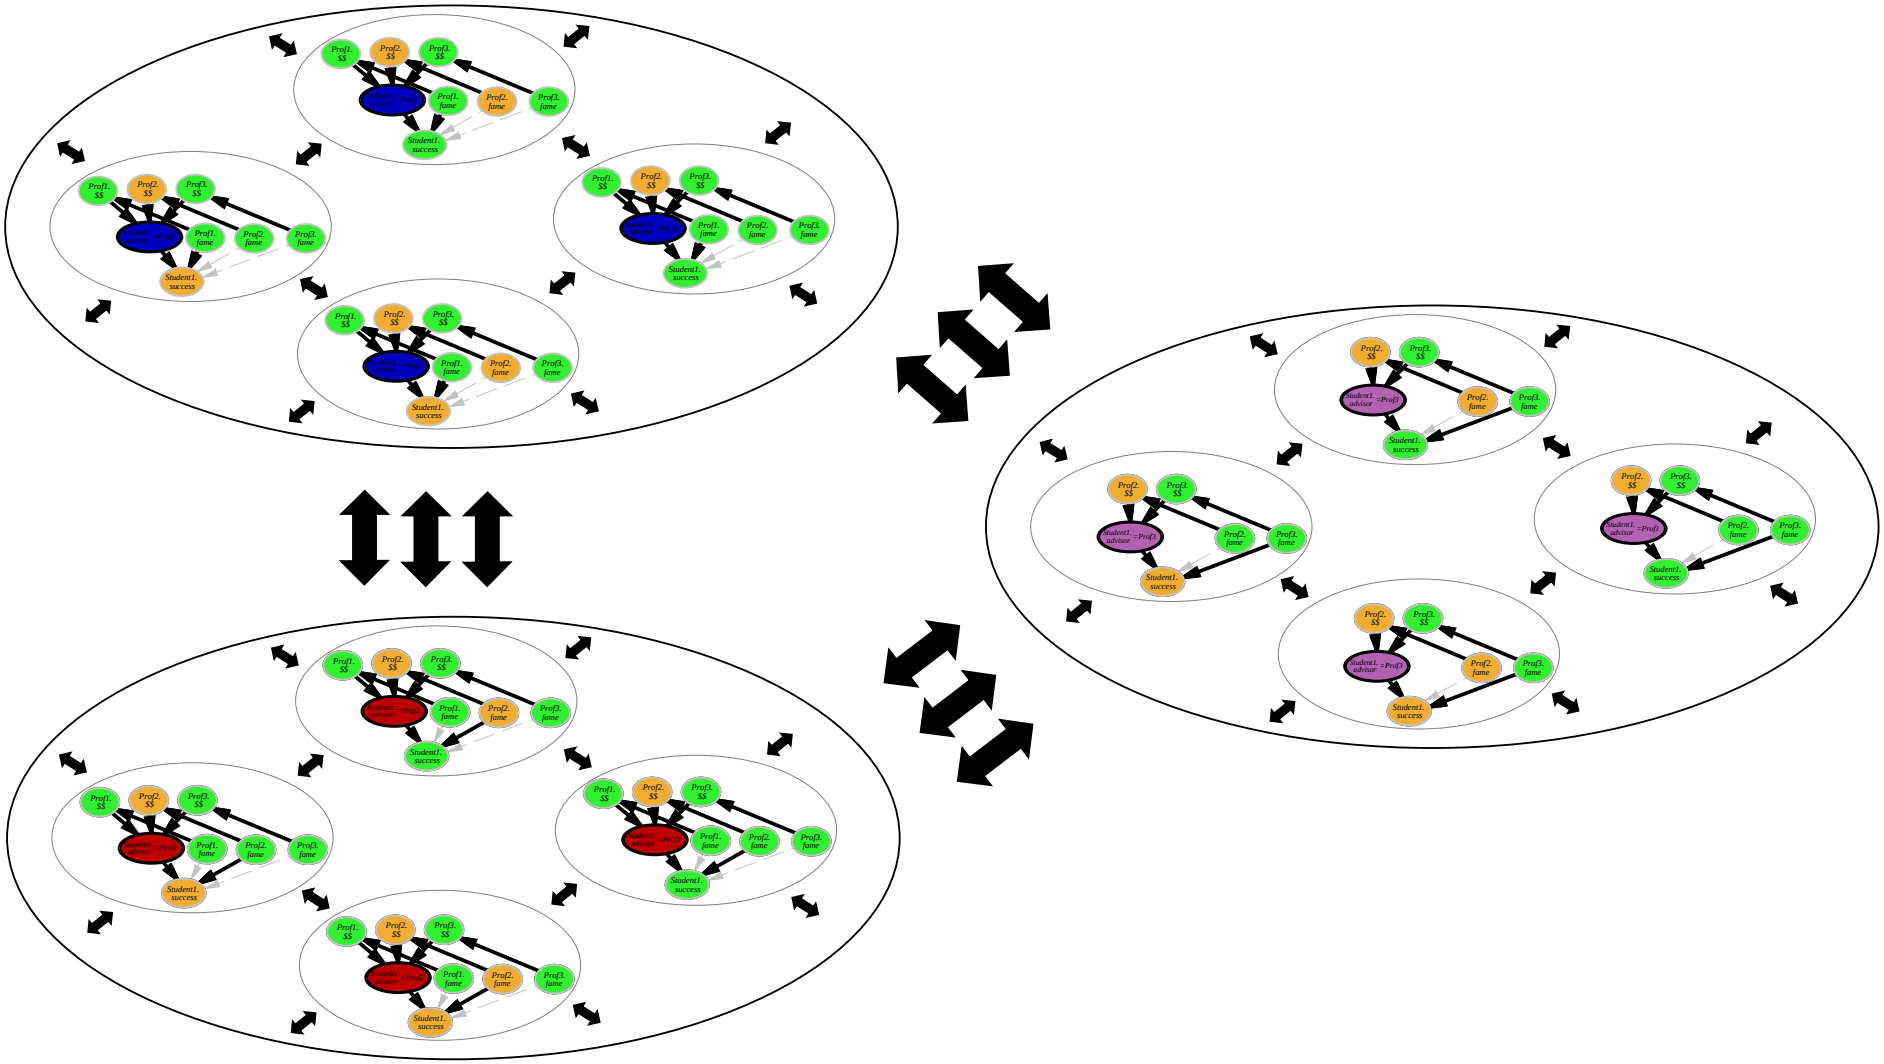

Citation matching

[PMMRS 02]: MCMC state is a complete publication database

Transitions (these all happen automatically via generic inference):

- change cited publication
- change authorship
- change true title, true author names
- change citation parsing decisions
- add/remove papers, researchers

Citation results

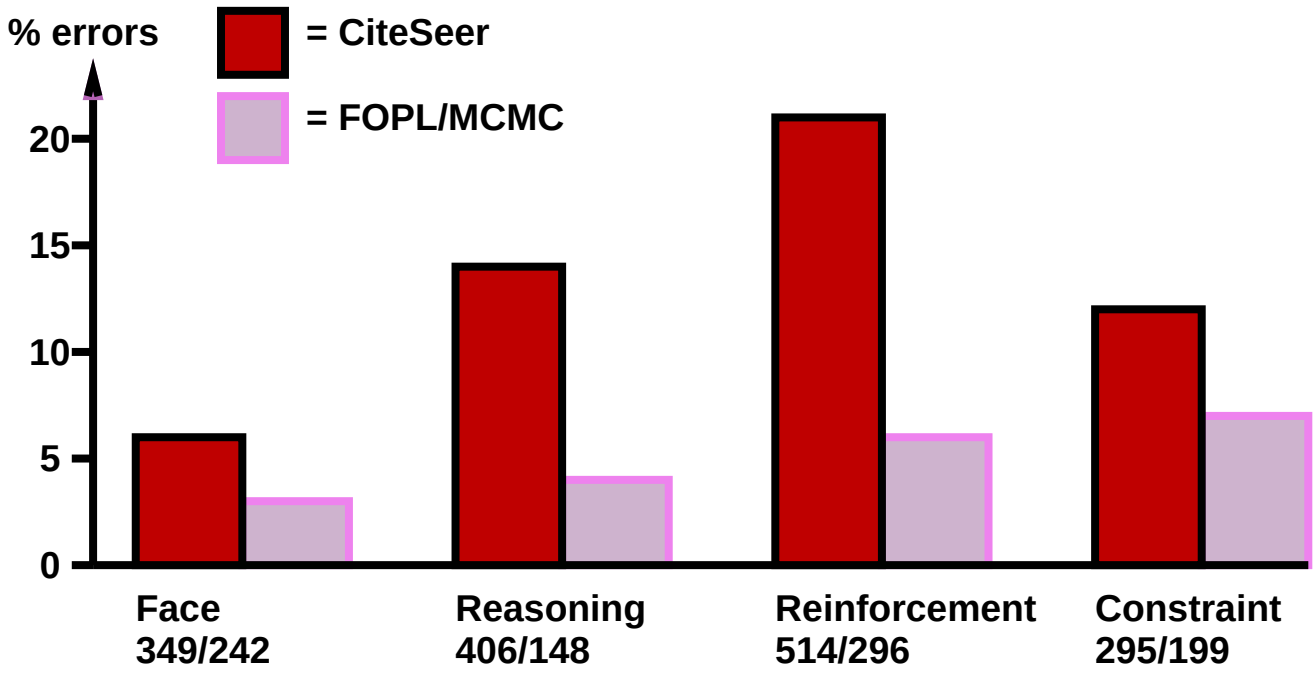

Integrated Parsing and Extraction

Citation parsing is nontrivial: HMM gets 31% correct

Wauchope, K. Eucalyptus: Integrating Natural Language Input with a Graphical User Interface. NRL Report NRL/FR/5510-94-9711 (1994).

Integrated Parsing and Extraction

Citation parsing is nontrivial: HMM gets 31% correct

Wauchope, K. Eucalyptus: Integrating Natural Language Input with a Graphical User Interface. NRL Report NRL/FR/5510-94-9711 (1994).

Adducing other evidence helps:

Kenneth Wauchope (1994). Eucalyptus: Integrating natural language input with a graphical user interface. Technical Report, Naval Research Laboratory, Washington, DC, 39pp.

64% correct parsing by combining two citations

Knowledge and Parsing

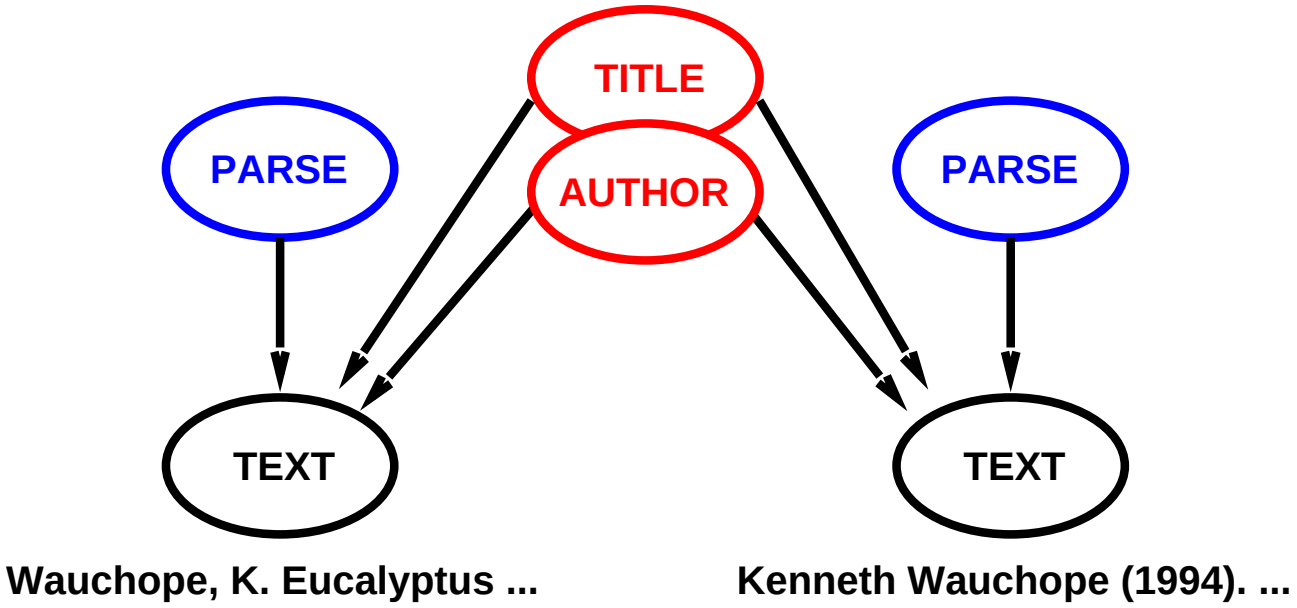

Knowledge directly affects MCMC parsing decisions

Transfer of information depends on paper identity assumption
i.e., combining top-down and bottom-up processing

Summary

Probability models for full FOL model structures—unknown worlds

Inference currently slow except for application-specific proposals

Needs a generic adaptive proposal mechanism

Parametric learning easy via stochastic EM; structural learning harder