Probability for Worlds with Things in Them

Stuart Russell
UC Berkeley

Lectures

1. (me) Motivation, technical background, overview, applications
2. (Brian) Relational probability models for worlds with known sets of objects
3. (Brian) BLOG models for worlds with unknown objects

2

Outline

- Logical knowledge representation
  - Syntax, semantics, inference in general
  - Expressiveness
  - Atomic, propositional, and first-order representations
- First-order logic
- Herbrand vs full first-order semantics
- Probabilistic knowledge representations
  - Probabilistic logics
  - First-order probability models
- Open-world applications
- Open problems

3

Knowledge representation

- Knowledge base = set of sentences in a formal language
- Declarative systems: perception/learning create sentences in KB; decisions *follow* from KB
- Can be understood in terms of what they *know*, regardless of implementation

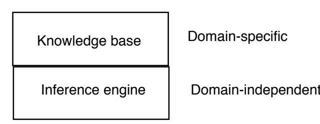

4

Logic in general

- Logics are formal languages for representing information such that conclusions can be drawn
- Syntax defines the sentences in the language
- Semantics define the meaning of sentences; i.e., define truth of a sentence in a world
- E.g., the language of arithmetic
  - x+2 > y is a sentence;
  - x2+y > is not a sentence;
  - x+2 > y is true in a world where x=7,y=1;
  - x+2 > y is false in a world where x=0,y=6

5

Entailment

- Entailment means that one thing follows from another: $\alpha \models \beta$ iff β is true in all worlds where α is true
- E.g., x+y = 4 entails 4 = x+y
- Entailment is a relationship between sentences (i.e., syntax) that is based on semantics
- Note: brains process syntax (of some sort)

6

Entailment and inference

- Possible worlds $\omega \in \Omega$
- α is true in ω: $\omega \models \alpha$
- To determine if $\alpha \models \beta$
  1) Every ω satisfying α also satisfies β
  2) No ω satisfies $\alpha \wedge \neg \beta$
  3) Every ω satisfies $\alpha \Rightarrow \beta$
- Problem: often infinitely many $\omega \in \Omega$; some logics undecidable

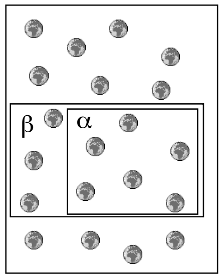

7

Kinds of possible worlds

- Atomic: each possible world is an atom or token with no internal structure. E.g., Heads or Tails
- Propositional: each possible world defined by values assigned to variables. E.g., propositional/Boolean logic
- First-order: each possible world defined by objects and relations
- [Probabilistic logics: each “possible world” is a complete probability model, wrt which probabilistic assertions are true or false]

8

Expressiveness matters

- Expressive language => concise models
  => fast learning, sometimes fast reasoning

E.g., rules of chess:
  1 page in first-order logic,
  ~100000 pages in propositional logic,
  ~1000000000000000000000000000000000000000 pages as atomic-state model
  [Note: chess is a teeny problem]

9

Expressiveness contd.

- Expressive language => general-purpose inference and learning methods

10

Expressiveness contd.

- Expressive language => general-purpose inference and learning methods

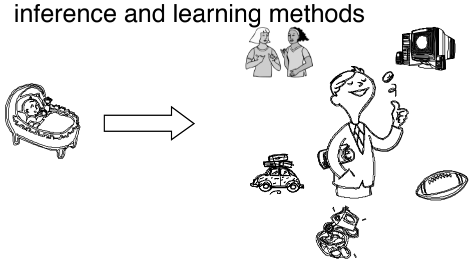

11

Expressiveness contd.

- Expressive language => general-purpose inference and learning methods

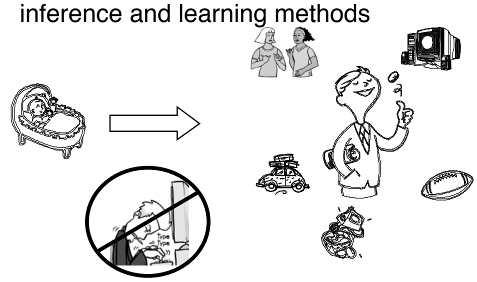

12

Brief history of expressiveness

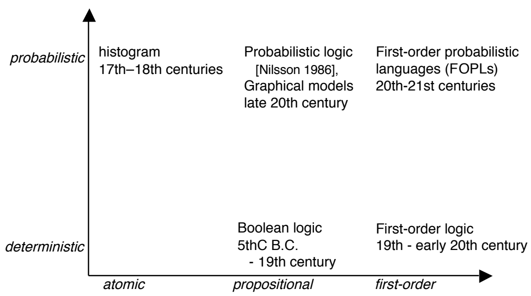

13

Examples

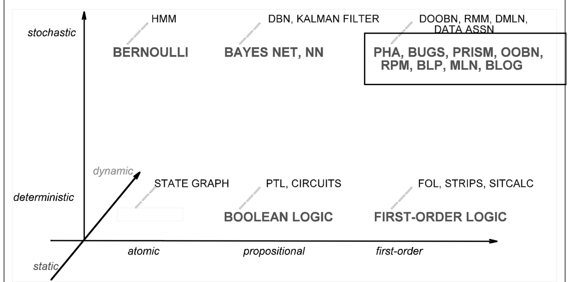

14

First-order logic

- The basic mathematics for worlds with things in them
- Expressive enough to say almost anything (definite) of interest
- There exists a sound and complete inference procedure
  - I.e., the procedure will answer any question whose answer follows from what is known

15

FOL syntax: basic elements

- Constants: Brian, 2, AIMA2e, UCB,...
- Predicates: AuthorOf, >,...
- Functions: PublicationYear,√,...
- Variables: x,y,a,b,...
- Connectives: $\wedge$ $\vee$ $\neg$ $\Rightarrow$ $\Leftrightarrow$
- Equality: =
- Quantifiers: $\forall$ $\exists$

16

Terms

- A term refers (according to a given possible world) to an object in that world
- Term =
  - function($term_1$,...,$term_n$) or
  - constant symbol or
  - variable
- E.g., PublicationYear(AIMA2e)
- Arbitrary nesting $\Rightarrow$ infinitely many terms

17

Atomic sentences

- Atomic sentence =
  - predicate($term_1$,...,$term_n$) or
  - term1=term2
- E.g.,
  - AuthorOf(Norvig,AIMA2e)
  - NthAuthor(AIMA2e,2) = Norvig
- Can be combined using connectives, e.g.,
  (Peter=Norvig) $\Rightarrow$(NthAuthor(AIMA2e,2) = Peter)

18

Semantics: Truth in a world

- Each possible world contains ≥1 objects (domain elements), and maps…
  - Constant symbols → objects
  - Predicate symbols → relations (sets of tuples of objects satisfying the predicate)
  - Function symbols → functional relations
- An atomic sentence predicate($term_1$,...,$term_n$) is true iff the objects referred to by $term_1$,...,$term_n$ are in the relation referred to by predicate

19

Example

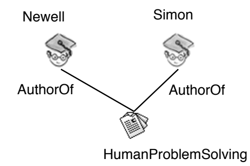

AuthorOf(Newell,HumanProblemSolving) is true in this world

20

Universal quantification

- ∀ <variable> <sentence>
- ∀p ∀c Venue(p)=c ∧Conference(c)⇒ ConferencePaper(p)
- ∀x P is true in ω iff P is true with x being each possible object in ω
- Roughly speaking, equivalent to the conjunction of instantiations of P

(Venue(AIMA1e)=AAAI80 ∧Conference(AAAI80) ⇒ ConferencePaper(AIMA1e))
∧(Venue(AIMA1e)=AAAI82 ∧Conference(AAAI82) ⇒ ConferencePaper(AIMA1e)) ∧...

21

Existential quantification

- ∃ <variable> <sentence>
- ∃p AuthorOf(Gates,p)
- ∃x P is true in ω iff P is true with x being some possible object in ω
- Roughly speaking, equivalent to the disjunction of instantiations of P

(AuthorOf(Gates,AIMA1e) ∨
AuthorOf(Gates,AIMA2e) ∨
AuthorOf(Gates,WayAhead)...)

22

Grounded Inference

Loop for d = 1 to infinity
- Instantiate variables with all possible ground terms of nesting depth up to d
- Each atomic sentence -> proposition symbol
- Apply propositional solver

- Will always find a proof if one exists
- Fails if no proof exists
- Unavoidable: FOL is semidecidable

23

Lifted inference

∀p ∀c Venue(p)=c ∧Conference(c)⇒ ConferencePaper(p)
Venue(BLOGPaper1)=IJCAI05
Conference(IJCAI05)
ConferencePaper(BLOGPaper1)

Works by unifying facts and premise using substitution {p/BLOGPaper1,c/IJCAI05}

24

Expressiveness and complexity in logic

[Poole, Mackworth & Goebel, 1998]

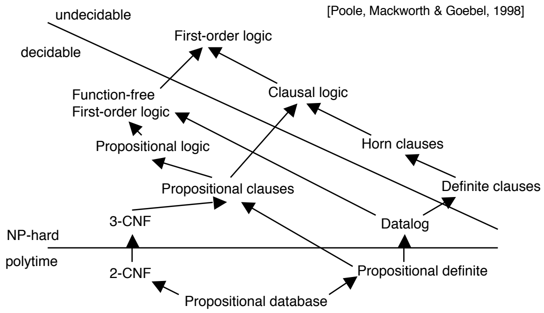

25

Herbrand vs full first-order

- Databases and logic programs use a more restricted "Herbrand" semantics:
  - Unique names assumption
    - Every term refers to a distinct object
  - Domain closure assumption
    - There are no other objects
  - Closed world assumption
    - Every atomic sentence not asserted true is false

26

Herbrand vs full first-order

Given

Father(Bill,William) and Father(Bill,Junior)

How many children does Bill have?

27

Herbrand vs full first-order

Given

Father(Bill,William) and Father(Bill,Junior)

How many children does Bill have?

Herbrand semantics:

2

First-order logical semantics:

Between 1 and ∞

28

Possible worlds

Restricted first-order (Herbrand)

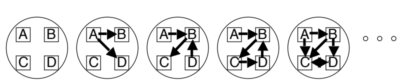

Full first-order (open-world)

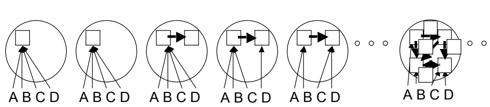

29

First-order probabilistic logics

- Gaifman [1964]:
  - distributions over first-order possible worlds
- Halpern [1990]:
  - syntax for constraints on such distributions

A probabilistic assertion such as

$\forall a_1 a_2 p\ P(AuthorOf(a_1,p)) = P(AuthorOf(a_2,p))$

is true or false wrt a distribution μ over worlds

Logical entailment works exactly as before.

Problem: almost impossible to avoid making inconsistent assertions or leave the model underconstrained

30

First-order probability models

- Poole [1993] and several others:
  - KB defines distribution exactly
  - assumes unique names and domain closure like Prolog, databases (Herbrand semantics)
- Milch et al [2005] and others (Laskey, Mjolsness):
  - distributions over full first-order possible worlds
    - generative models for events and existence
  - complete inference for all well-defined models

31

Possible Worlds

(not showing attribute values)

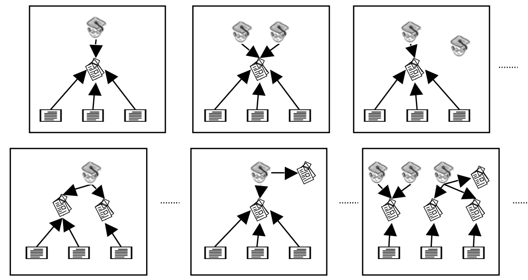

How can we define a distribution over such outcomes?

32

Generative Process (Bayes nets)

- Construct worlds using one kind of step, proceeding in topological order:
  - Set the value of a variable conditioned on its parent values

33

Generative Process (Herbrand)

- Construct worlds using one kind of step, proceeding in topological order:
  - Set the value of a function or relation on a tuple of arguments, conditioned on parent values

34

Generative Process (First-Order)

[Milch *et al.*, IJCAI 2005]

- Construct worlds using two kinds of steps, proceeding in topological order:
  - Add some objects to the world, conditioned on what objects and relations exist so far
  - Set the value of a function or relation on a tuple of arguments, conditioned on parent values
    - Includes setting the referent of a constant symbol (0-ary function)

35

Application: IRS investigation

- Given tax returns from many tax entities and transactions
  - "To" and "from" Taxpayer IDs and addresses often scrambled
  - Several thousand fictitious entities
- Work out who owns whom, which entities are which, which ones are real, etc.

36

Application: 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? How many papers are there? Who wrote them and which papers did they cite?

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

37

CiteSeer02: Russell w/4 Norvig

38

CiteSeer02: 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.

39

- 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.

40

- 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.
- 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.

41

- 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.
- 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.

42

- 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.

43

- 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.

44

- 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.

45

- 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.

46

- 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.
- 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.

47

- 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

48

- 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.

49

- 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).

50

Example: classical data association

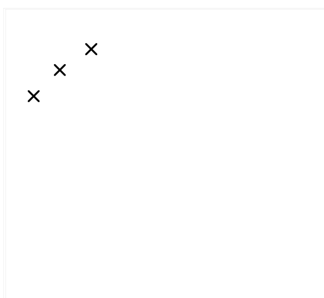

51

Example: classical data association

52

Example: classical data association

53

Example: classical data association

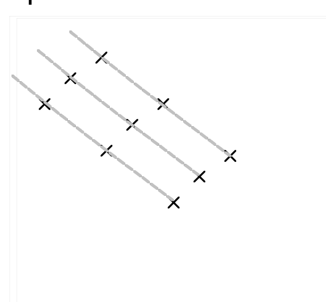

54

Example: classical data association

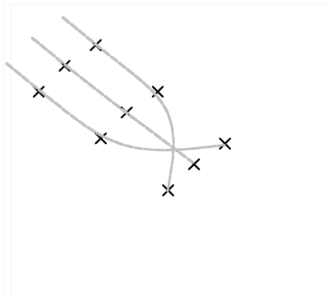

55

Example: classical data association

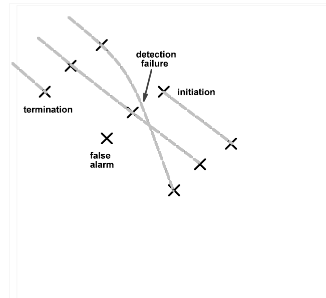

56

Example: modern data association

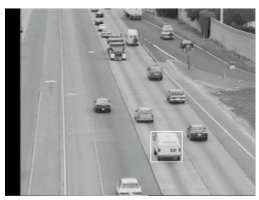

Are these two vehicles the same??

57

Modern data association

Same car?

Need to take into account competing matches!

58

Example: natural language

- What objects are referred to in the following natural language utterance?

59

Example: vision

- What objects appear in this image sequence?

60

Summary

- Expressive probability languages are important for many real-world problems and for real AI
- We must be able to specify probability distributions over possible worlds with unknown objects and relations
- BLOG does this

61

Open Problems

- Inference
  - Applying “lifted” inference to BLOG (like Prolog)
  - Approximation algorithms for problems with huge numbers of objects
  - Effective filtering algorithm for DBLOG
- Structure learning
  - Learning new dependencies within BLOG
  - Hypothesizing new random functions, new types
- Applications to vision, NLP, WWW => KB

62

Merging Logical and Probabilistic AI

- Logical languages are special cases of probabilistic!
- Does probabilistic inference operate in the deterministic limit like logical inference?
  - One hopes so! The brain does not have a big switch!!!
  - MCMC asymptotically identical to greedy SAT algorithms
  - Several algorithms exhibit “default logic” phenomenology
- Modern AI should build on, not discard, classical AI

63