Stochastic Search And Phase Transitions: AI Meets Physics

Bart Selman

Cornell University

Computational Challenges In AI

Many core AI tasks have been shown to be computationally intractable.

We have results in:

- **reasoning**
- **planning**
- **learning**

A Few Examples

Reasoning

- many forms of deduction
- abduction / diagnosis (e.g. de Kleer 1989)
- default reasoning (e.g. Kautz and Selman 1989)
- Bayesian inference (e.g. Dagum and Luby 1993)

Planning

- domain-dependent and independent (STRIPS)
  (e.g. Chapman 1987; Gupta and Nau 1991; Bylander 1994)

Learning

- neural net “loading” problem (e.g. Blum and Rivest 1989)

Complexity Results, Cont.

- An abundance of **negative** complexity results.

  Sets AI apart from other areas in CS.

- Results often apply to **very restricted** formalisms, and also to finding **approximate** solutions.

What Is The Impact Of These Results?

- Results are based on a **worst-case** analysis and there continues to be a debate on their **practical** relevance.
- On the one hand, there are successful systems that do not appear to be hampered by the negative complexity results.

Examples: Bayesian net applications,
Neural nets,
CLASSIC KR system (Brachman et al. 1989)

- On the other hand, in other domains, negative complexity properties are a clear obstacle in **scaling-up** the systems.

Examples: ATMS diagnosis: 25+ components
planning systems: 20+ objects and operators
(Real domains: 1,000+ elements.)

- Contradictory experiences lead to the question:

When and where do computationally hard instances show up?

Recent Developments

A --- A better understanding of the nature of computationally hard problems.

B --- New stochastic methods for solving such problems.

BS - 08/95 7

Overview

PART A. Computationally Hard Instances

worst-case vs. average-case

critically-constrained problems

phase transitions

PART B. Stochastic Methods

heuristic repair, GSAT, and simulated annealing

comparison with systematic methods

asymmetry consistency / inconsistency

Summary

BS - 08/95 8

PART A. Computationally Hard Instances

- I'll use the **propositional satisfiability problem (SAT)** to illustrate ideas and concepts throughout this talk.
- SAT: prototypical hard combinatorial search and reasoning problem.

Several of these concepts have also been studied in the context of **Constraint Satisfaction Problems**. In particular, see the work by Cheeseman and colleagues (1991).

BS - 08/95 9

Satisfiability

- SAT: Given a formula in propositional calculus, is there an assignment to its variables making it true?
- We consider clausal form, e.g.:

$$(a \vee \neg b \vee c) \wedge (\neg b \vee d) \wedge (b \vee c \vee e) \wedge \ \ldots$$

- Problem is NP-Complete. (Cook 1971)
- Shows surprising "power" of SAT for encoding computational problems.

BS - 08/95 10

Some Example Applications Of SAT

- constraint satisfaction
  - scheduling and planning
  - temporal reasoning (Allen 1983)
  - VLSI design and testing (Larrabee 1992)
- direct connection to deductive reasoning

  $\Sigma \models \alpha$ iff $\Sigma \cup \{\neg \alpha\}$ is not satisfiable
- part of other AI reasoning tasks
  - diagnosis / abduction
  - default reasoning
- learning

BS - 08/95 11

How well can SAT be solved in practice?

BS - 08/95 12

Average-Case Analysis

- Goldberg (1979) reported very good performance of Davis-Putnam (DP) procedure on random instances.

  But distribution favored easy instances. (Franco and Paull 1983)

- Problem: Many randomly generated SAT problems are **surprisingly easy**.
- Goldberg used variable-clause-length model:

  For each clause, pick each literal with probability p.

Variable Clause Size Model

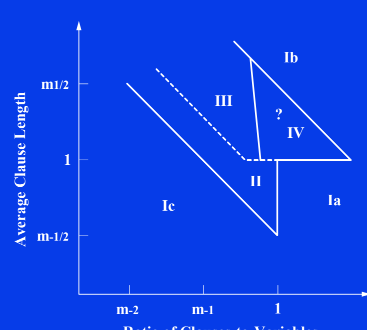

Polynominal average time in regions:

Ia Ð Purdom 1987 - backtracking
Ib Ð Iwama 1989 - counting alg.
Ic Ð Brown and Purdom 1985 - pure literal rule
II Ð Franco 1991
III Ð Franco 1994

Open: region IV

But the problem is NP-complete ...

where are the hard instances?

BS - 08/95 15

Aside: Small Hard Instances Do Exist!

- Job-Shop Scheduling: 10 jobs on 10 machines.
- Proposed by Fischer and Tompson in **1963**.
- Solved by Carlier and Pinson in **1990**!
- Open: 15 jobs on 15 machines.

BS - 08/95 16

The Instance

Job#	Machine Order / Duration									
1	0 29	1 78	2 09	3 36	4 49	5 11	6 62	7 56	8 44	9 21
2	0 43	2 90	4 75	9 11	3 69	1 28	6 46	5 46	7 72	8 30
3	1 91	0 85	3 39	2 74	8 90	5 10	7 12	6 89	9 45	4 33
4	1 81	2 95	0 71	4 99	6 09	8 52	7 85	3 98	9 22	5 43
5	2 14	0 06	1 22	5 61	3 26	4 69	8 21	7 49	9 72	6 73
6	2 84	1 02	5 52	3 95	8 48	9 72	0 47	6 65	4 06	7 25
7	1 46	0 37	3 61	2 13	6 32	5 21	9 32	8 89	7 30	4 55
8	2 31	0 86	1 46	5 74	4 32	6 88	8 19	9 48	7 36	3 79
9	0 76	1 69	3 76	5 51	2 85	9 11	6 40	7 89	4 26	8 74
10	1 85	0 13	2 61	6 07	8 64	9 76	5 47	3 52	4 90	7 45

Generating Hard Random Formulas

- Key: Use **fixed-clause-length** model.

  (Mitchell, Selman, and Levesque 1992)
- Critical parameter: ratio of the number of clauses to the number of variables.
- Hardest 3SAT problems **at ratio = 4.3**

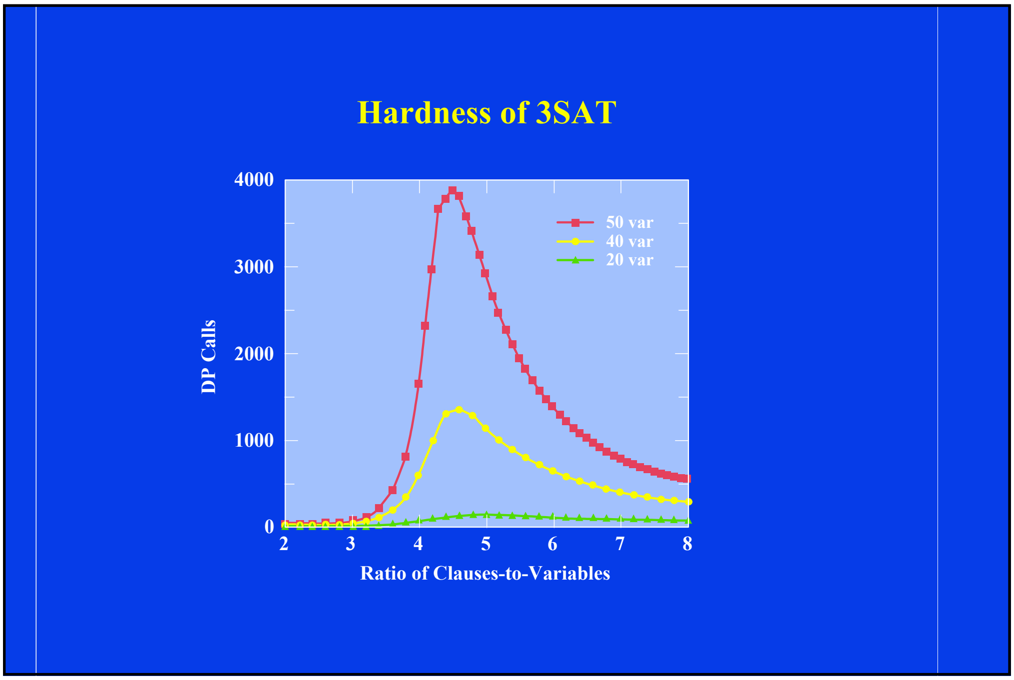

Intuition

- At low ratios:
  - few clauses (constraints)
  - many assignments
  - easily found
- At high ratios:
  - many clauses
  - inconsistencies easily detected

BS - 08/95 20

The 4.3 Point

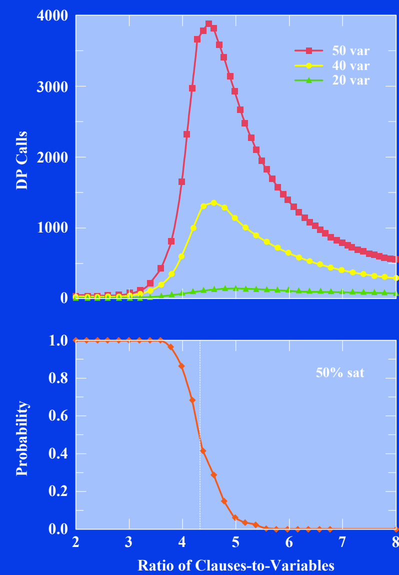

Mitchell, Selman, and Levesque 1991

Theoretical Status Of Threshold

- Very challenging problem **...**
- Current status:

  3SAT threshold lies between **3.003** and **4.8**

  (Motwani et al. 1994; Broder and Suen 1993; Broder et al. 1992; Dubois 1990)

A Closer Look At The 3SAT Phase Transition

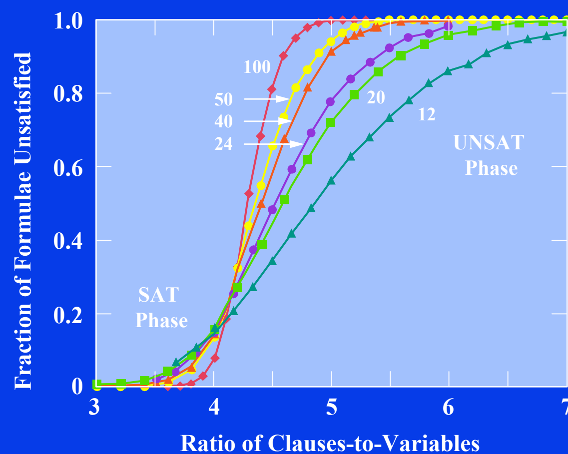

Transition sharpens up for higher values of N

The Physics Of Thresholds

- Threshold phenomena studied in physics as part of **phase-transition** phenomena.

  E.g., transition from ice to water.

- Analyzed by **rescaling** ("stretching") the X-axis using a function dependent on N (number of variables) to **slow down** transition for large N.

  (Kirkpatrick and Selman, *Science* 1994)

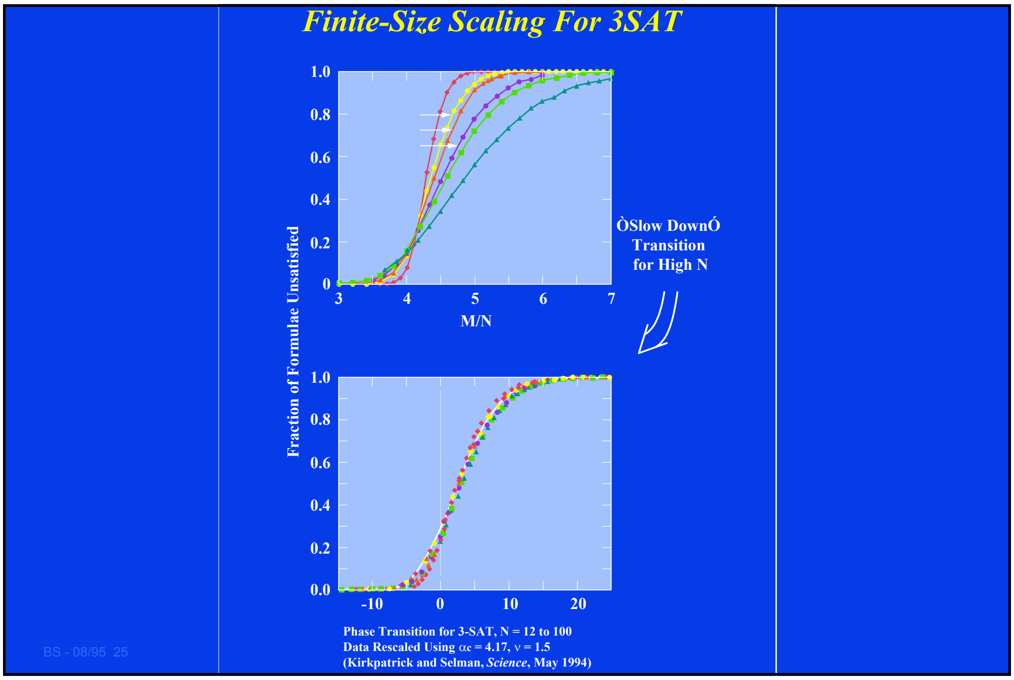
Finite-Size Scaling For 3SAT
Fraction of Formulae Unsatisfied
1.0
0.8
0.6
0.4
0.2
0
3
4
5
6
7
M/N
ÒSlow DownÓ
Transition
for High N
1.0
0.8
0.6
0.4
0.2
0.0
-10
0
10
20
Phase Transition for 3-SAT, N = 12 to 100
Data Rescaled Using αc = 4.17, ν = 1.5
(Kirkpatrick and Selman, Science, May 1994)
BS - 08/95 25

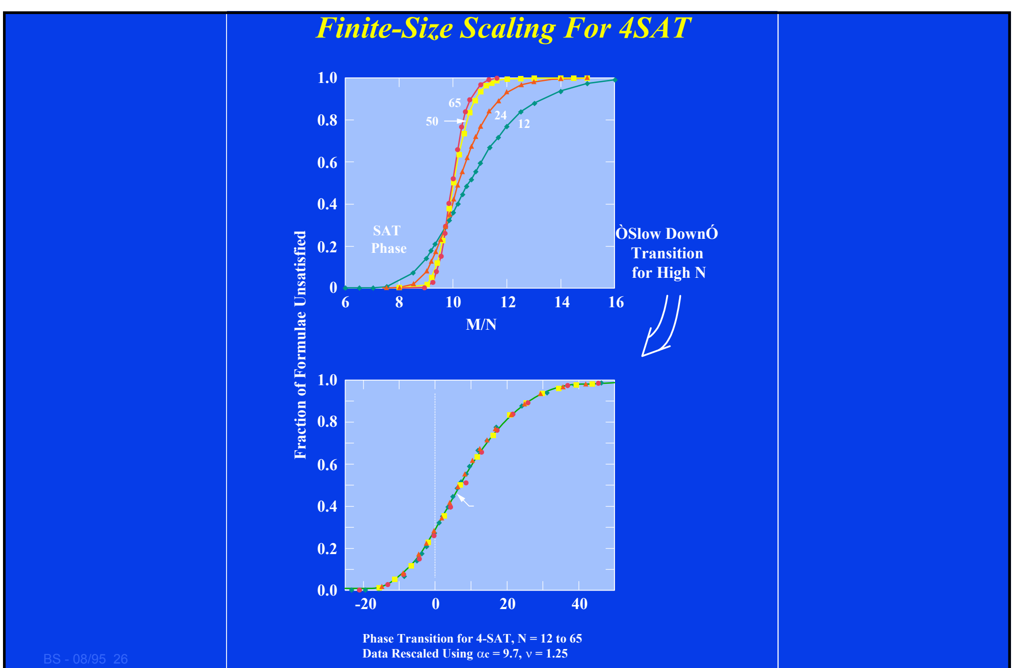
Finite-Size Scaling For 4SAT
Fraction of Formulae Unsatisfied
1.0
0.8
0.6
0.4
0.2
0
6
8
10
12
14
16
M/N
65
50
24
12
SAT
Phase
ÒSlow DownÓ
Transition
for High N
1.0
0.8
0.6
0.4
0.2
0.0
-20
0
20
40
Phase Transition for 4-SAT, N = 12 to 65
Data Rescaled Using αc = 9.7, ν = 1.25
BS - 08/95 26

Summary Phase Transition Effect

- Coincides with hardest instances.
- Behavior at threshold can be analyzed with tools from statistical physics:
  - Threshold has universal form with predictable corrections for N (number of vars).
  - Inverse transformation gives 50% point for testing.

    (Also, rescaling cost function; Selman and Kirkpatrick 1995, Gent and Walsh 1995)

- Similar phenomenon for graph coloring.
  - random graphs
  - 3-coloring; threshold around 4.6 (connectivity)

    (Cheeseman et al. 1991)
- **Critically-constrained** --- Practical relevance
  - Airline fleet scheduling (Nemhauser 1994)
  - VLSI design (Agrawal 1991)
  - Traveling Salesperson Problem (Gent and Walsh 1995)

See also Hogg, Huberman, and Williams 1996; Crawford and Auton 1993; Frost and Dechter 1994; Larrabee and Tsuji 1993; Schrag and Crawford 1996; Smith and Grant 1994; Smith and Dyer 1996; and more!

PART B. Fast Stochastic Methods

- After having identified hard instances, can we find better algorithms for solving them?
- Answer: Yes (at least for half of them...)

BS - 08/95 29

Standard Procedures For SAT

- **Systematic search** for a satisfying assignment.
- Interesting situation:
  - Davis-Putnam (DP) procedure, proposed in **1960**, is still the fastest complete method!
  - Backtrack-style procedure with unit propagation.

    SAT Competition 1992; DIMACS Challenge 1993 / 1994

BS - 08/95 30

- DP provides **very** challenging benchmark for comparisons with other systematic (complete) procedures.

Not just on random formulas!

- Many other methods have been tried, e.g.,

1) Backtracking with sophisticated heuristics

(Purdom 1984; Zabih and McAllester 1988; Andre and Dubois 1993; Bhom 1992; Crawford and Auton 1993; Freeman 1993, etc.)

2) Translations to integer programming

(Jeroslow 1986; Hooker 1988; Karmarkar et al. 1992; Gu 1993)

3) Exploiting hidden structure

(Stamm 1992; Larrabee 1991; Gallo and Urbani 1989; Boros et al. 1993)

4) Limited resolution at the backtrack nodes

(Billionet and Sutter 1992; van Gelder and Tsuji 1993)

- **And others!**

Open Question: Why don't they beat DP?

- Let's try something completely different **...**

Randomized Greedy Local Search: GSAT

Begin with a random truth assignment.

Flip the value assigned to the variable that yields greatest number of satisfied clauses.

Repeat until a model is found, or have performed specified maximum number of flips.

If model is still not found, repeat entire process, starting from different random assignment.

(Selman, Levesque, and Mitchell 1992)

How Well Does It Work?

- First intuition: Will get **stuck in local minimum**, with a few unsatisfied clauses.
- No use for **almost** satisfying assignments.
  E.g., a plan with a “magic” step is useless.
  Contrast with optimization problems.
- Surprise: It often finds **global** minimum!
  I.e., finds satisfying assignments.
- Inspired by local search for CSP initially used on N-Queens: **Heuristic Repair Method.** (Minton et al. 1991)

GSAT outperforms Davis-Putnam on, e.g.:

- Hard random formulas
  - DP: up to 400 vars; GSAT: 2000+ var formulas.
- Boolean encodings of graph coloring problems.
  - GSAT competitive with direct encodings.
- Encodings of Boolean circuit synthesis and diagnosis problems.

GSATÕs Search Space

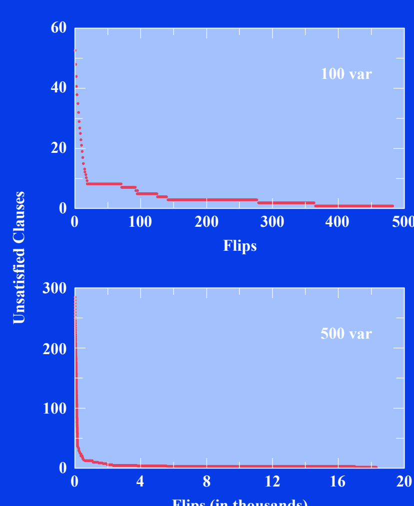

Improvements Of Basic Local Search

Issue: How to move more quickly to successively lower plateaus?

Idea: Introduce **uphill** moves ("noise") to escape from long plateaus (or true local minima).

Noise strategies:

a) Simulated Annealing

(Kirkpatrick et al. 1982)

b) Biased Random Walk

(Selman, Kautz , and Cohen 1993)

BS - 08/95 37

Simulated Annealing

- Noise model based on statistical mechanics.
- Pick a random variable

δ = change in number of unsatisfied clauses

If $\delta < 0$ make flip ("downward")

else flip with probability $e^{-\delta/T}$ ("upward").

Slowly decrease T from high temperature to near zero.

BS - 08/95 38

Random Walk

- Random walk SAT algorithm:

  1) Pick random truth assignment.

  2) Repeat until all clauses are satisfied:
  Flip random variable from unsatisfied clause.

- Solves 2SAT in O(n^2) flips. (Papadimitriou 1992)
- Does **not** work for hard k-SAT (k >= 3).

BS - 08/95 39

Biased Random Walk

1) With probability p, **“walk”**, i.e.,
flip variable in some unsatisfied clause.

2) With probability 1-p, **“greedy move”**, i.e.,
flip variable that yields greatest number of satisfied clauses.

BS - 08/95 40

Experimental Results: Hard Random 3SAT

	GSAT				Sim. Ann.	
	basic		walk			
vars	time	eff.	time	eff.	time	eff.
100	.4	.12	.2	1.0	.6	.88
200	22	.01	4	.97	21	.86
400	122	.02	7	.95	75	.93
600	1471	.01	35	1.0	427	.3
800	*	*	286	.95	*	*
1000	*	*	1095	.85	*	*
2000	*	*	3255	.95	*	*

Biased Walk better than Sim. Ann. better than Basic GSAT better than DP.

Other Applications Of GSAT

- VLSI circuit diagnosis

  SAT formulation by Larrabee (1992)
  approx. 10,000 var 5,000 clause problems

- Planning and scheduling

  approx. 20,000 var 100,000 clause problems
  (Crawford and Baker 1994)

- Finite algebra

  search for algebraic structures

  GSAT+walk outperforms systematic method on large instances. Currently exploring remaining open problems. (Fujita et al. 1993)

For other work on stochastic, incomplete methods, see e.g.:

Adorf and Johnston 1990; Beringer et al. 1994; Davenport et al. 1994 (GENET); Kask and Dechter 1995; Ginsberg and McAllester 1994; Gu 1992; Hampson and Kibler 1993; Konolige 1994; Langley 1992; Minton et al. 1991; Morris 1993; Pinkas and Dechter 1993; Resende and Feo 1993; Spears 1995, and others!

- **GSAT-style procedures are now a promising alternative to systematic methods.**
- **Drawback: cannot show unsatisfiability.**

Showing UNSAT / Inconsistencies

Given the success of stochastic search methods on satisfiable instances, a natural question is:

Can we do something similar for unsatisfiable instances?

To show a set of clauses S **unsatisfiable**, we need to demonstrate ("**prove**") that none of the 2^N truth assignments satisfies S.

This "truth-table" method is very time consuming.

Compare this with having to check a single satisfying assignment to verify the satisfiability of a formula.

Can we do better? --- Surprisingly difficult!

Length Of Proofs

- Best know improvement on truth tables: **resolution**
  - Resolve clauses until empty clause is reached.
  - Widely used in automated theorem proving.
- DP is a form of resolution.

BS - 08/95 47

Limitations Of Resolution

- **Method can't "count"!** Pigeon-hole formulas:

  Can't place N+1 objects in N holes.

  Shortest resolution proof is exponentially long.

  (Cook / Karp 1972; Haken 1985)

- **Random unsat formulas: exponential size proofs.**

  Explains why we can't push DP over 400 vars:

  400 vars requires search tree of about 10 million nodes

  1000 vars unsat requires 10^15 nodes!

  (Chvatal and Szemeredi 1988; Crawford 1995)

BS - 08/95 48

Stochastic Search For Proofs

- **GSAT:** start with random truth assignment (size linear in N), and try to "fix" it.
- **Proposal for UNSAT:** start with random proof structure, and try to fix it.
- *Completely unfeasible if the structure that we're fixing has trillions of nodes (exponential in N).*
- We need short proofs! (O(N) or something**...**)
  (Using abstractions / symmetrries?)

BS - 08/95 49

Recap Of Results

A) Computationally hard problem instances

- Hardest ones are critically-constrained.
- Under- and over-constrained ones can be surprisingly easy.
- Critically-constrained instances at phase-transition boundaries.

  Properties of transition can be analyzed with tools from statistical physics.

BS - 08/95 50

B) Stochastic Search Methods

- **GSAT:** Randomized local search for SAT testing.
  Viable alternative to systematic, complete methods.
- **Progress:**
  - 1991: 10 vars, 500 clause theories.
  - 1995: 2,000 to 20,000 vars, up to 500,000 clauses
- **Approaches size of practical applications.**
  E.g. in scheduling, planning, diagnosis, circuit design, and constraint-logic programming.

See proceedings for many additional pointers.

Impact And Future Directions

Fast Incomplete Methods

- Shift in Reasoning and Search from **Systematic** / **Complete** methods to **Stochastic** / **Incomplete** methods.
- Key issue: Better scaling properties.
- Analogy in OR: Shift from finding optimal to finding approximate solns.
- Also, little progress on heuristic guidance of complete methods. DP still rules…

Impact, Cont.

Message for KR&R

- Asymmetry between our ability to show **satisfiability vs. unsatisfiability**, argues for **model-finding** (show sat) over **theorem proving** (show unsat).
- Examples:
  - Vivid repr. (Levesque 1985)
  - Planning (Kautz and Selman 1992)
  - Abduction / diagnosis / deduction
    - Model-based repr. versus formula-based repr. (Kautz, Kearns, and Selman 1994; Khardon and Roth 1994)
    - Case-based reasoning (Kolodner 1991)

BS - 08/95 53

Some Challenges

- Fast incomplete strategies for UNSAT (deduction)?

  Need for short proofs. Human proofs O(N)? Need automatic discovery of abstractions, symmetries, useful lemmas...

- Need for more **model-based** reformulations:

  Where solutions are **compact structures** --- allowing for randomized local search strategies.

- Can we syntactically characterize the class of instances solved by incomplete, stochastic methods?
  Running algorithm may be the best and only characterization!

BS - 08/95 54

Possible Limits Of Syntactic Characterization

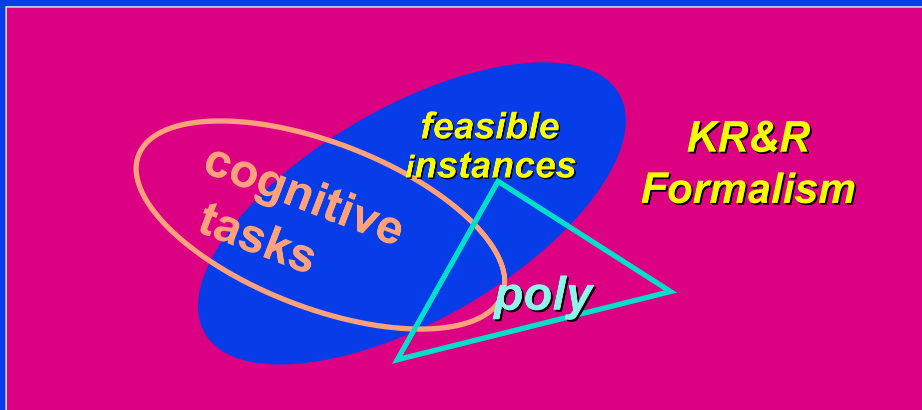

Would suggest fundamental role for incomplete methods.