ECE/CS 584: Verification of Embedded and Cyberphysical Systems

Lecture 4: Solving Satisifiability Problems

Prof. Huan Zhang

Office: CSL 262

huan@huan-zhang.com

Some slides adapted from Prof. Sayan Mitra's ECE584
Some of the slides for this lecture are adapted from slides by Clark Barrett

Review: Boolean satisfiability problem

Given a *well-formed boolean formula* α, determine whether there exists a satisfying solution

We will assume α to be in *conjunctive normal form (CNF)*

literals: variable or its negation, e.g., x_3, $\neg x_3$

clause: disjunction (or) of literals, e.g., $(x_1 \vee x_2 \vee \neg x_3)$

CNF formula: conjunction (and) of clauses,

e.g., $(x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_2 \vee x_1)$

A variable may appear *positively* or *negatively* in a clause

Review: Boolean satisfiability problem

Restatement: $\exists \boldsymbol{x} \in val(X) : \boldsymbol{x} \vDash \alpha ?$

If the answer is “No” then α is said to be *unsatisfiable*

SAT problem example:

$$\alpha := (\neg x_1 \vee x_2) \wedge (\neg x_3 \vee x_4) \\ \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \\ \wedge (x_5 \vee x_7) \wedge (x_1 \vee x_6 \vee \neg x_7)$$

Review: SAT is NP-complete

SAT was the first problem shown to be NP-complete [Cook 71]

1. Essentially we don't know better (in terms of asymptotic complexity) than naïve enumeration
2. A solver for SAT can be used to solve any other problem in the NP class with only polytime slowdown. i.e., makes a lot of sense to build SAT solvers
3. SAT/SMT solving is the cornerstone of *many* verification procedures

Stephen Cook, The complexity of theorem-proving procedures. In Proceedings of the third annual ACM symposium on theory of computing. STOC '71.

A simple greedy algorithm for SAT (GSAT)

Input: Set of clauses *C* over *X*, parameters *max-flips, max-tires*

Output: A satisfying assignment for C, or ∅ if none found

for i = 1 to *max-tries*

v := random truth assignment in *val(X)*

for j = 1 to *max-flips*

if $v \vDash C$ then return v

$p \coloneqq$ variable in C such that flipping its value gives the largest increase in the number of clauses of C that are satisfied by v

$v \coloneqq v$ with the assignment to p flipped

e.g., $x_1x_2x_3x_4x_5$ = 00100 -> 00110

return ∅

GSAT is a stochastic local search (SLS) algorithm

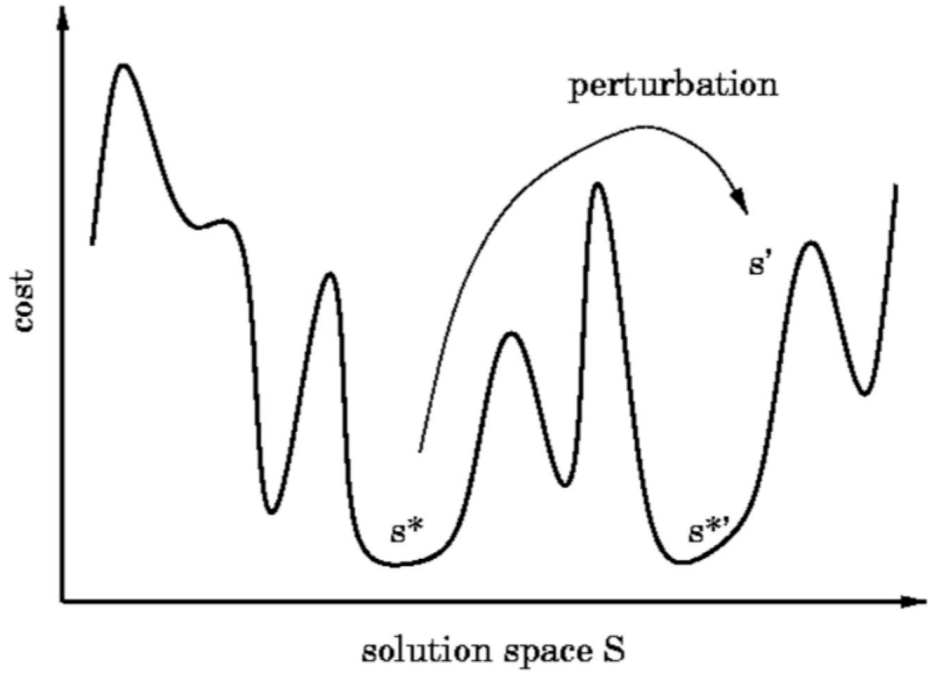

Limitation of this approach?

Local search algorithms are usually **incomplete**: they cannot show unsatisfiability!

Image Source: Alan Mackworth

SAT Solving with backtracking

Sysmetically enumerating all possibilities!

$$\alpha: = (\neg x_1 \vee x_2) \wedge (\neg x_3 \vee x_4) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7) \wedge (x_1 \vee x_6 \vee \neg x_7)$$

First, assume x_1 is True, and substitute

Because (True ∨ A) = True, (False ∨ A) = A, we can simplify α and it becomes:

$$x_2 \wedge (\neg x_3 \vee x_4) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee x_6 \vee \neg x_7)$$

But we still don't know if it is satisfiable!

SAT Solving with backtracking

After assuming x_1 is True and we get:

$$x_2 \wedge (\neg x_3 \vee x_4) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Search tree

SAT Solving with backtracking

Then, let's substitute x_2 = True in

$$x_2 \wedge (\neg x_3 \vee x_4) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

α is still unresolved:

$$(\neg x_3 \vee x_4) \wedge (\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

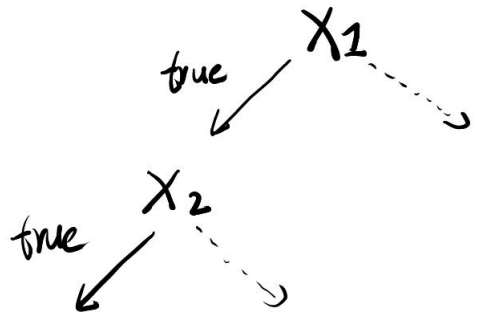

SAT Solving with backtracking

Keep setting and substituting variables

$$(\neg x_3 \vee x_4) \wedge (\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Set x_3 = True $\quad x_4 \wedge (\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$

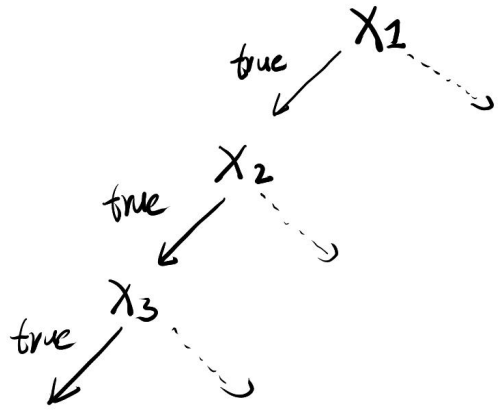

SAT Solving with backtracking

Keep setting variables

$$x_4 \wedge (\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Set x_4 = True $\quad (\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$

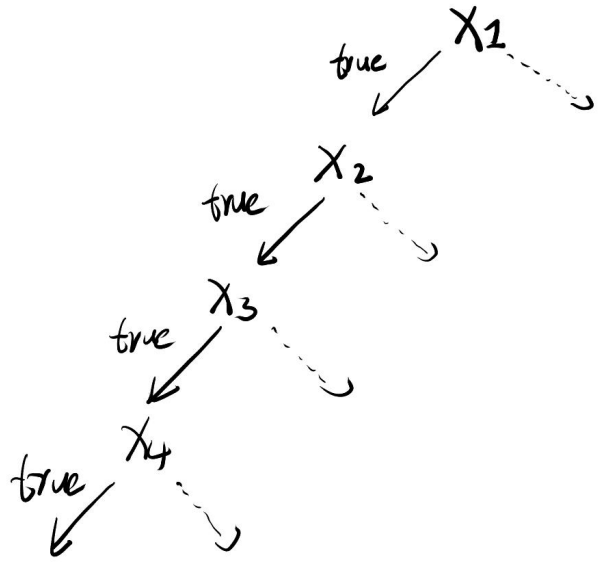

SAT Solving with backtracking

Keep setting variables

Set x_5 = True $\quad (\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$

Set x_6 = True $\quad \neg x_6 \wedge x_6$

Conflict, α evaluates to False!

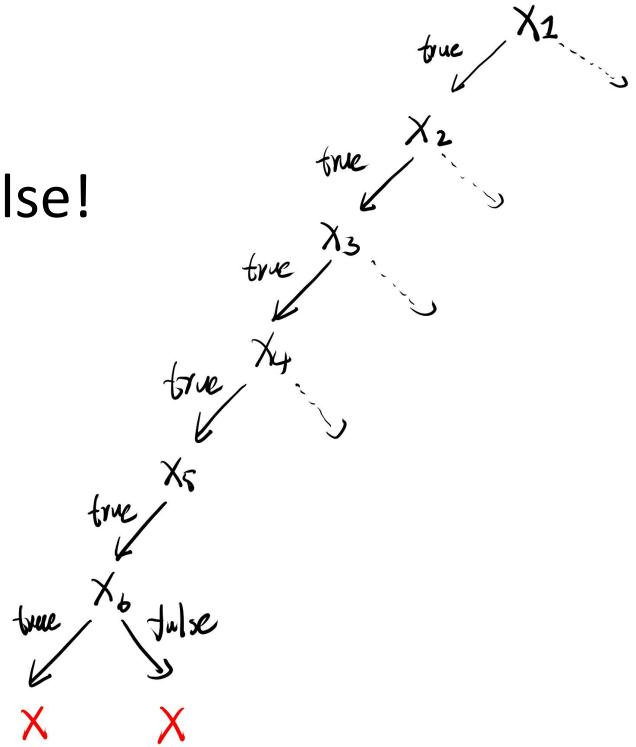

SAT Solving with backtracking

$$\neg x_6 \wedge x_6$$

Set x_6 = True does not work. Backtrack and try a different x_6.

Setting x_6 = False also does not work. Backtrack one-level up and try x_5

$$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

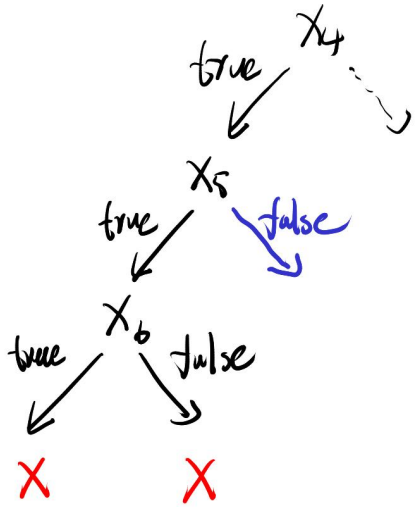

SAT Solving with backtracking

$$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Now set x_5 = False $\quad \neg x_6 \vee \neg x_7$

Now set x_6 = True $\quad \neg x_7$

Now set x_7 = True, does not work.

Set x_7= False, α is now true!

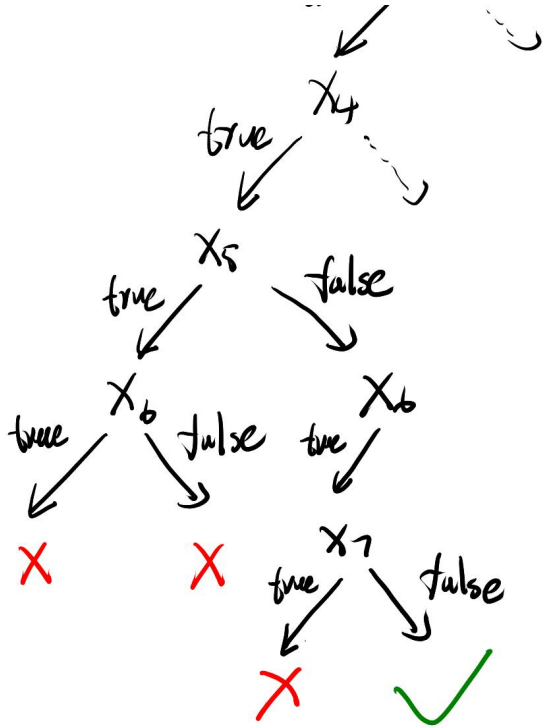

SAT Solving with naive backtracking

```
function BackTracking(α)
    if α is true then return true;
    if α is false then return false;
    // α is unresolved, need to decide on a literal
    l ← choose-literal(α);
    return (BackTracking(substitute l in α with true) or
            BackTracking(substitute l in α with false));
```

Search tree can be large!

Each variable is tested with two cases (true and false).
Complexity exponential to the number of variables.

To prove unsatifiability, the entire tree must be visited.

We need to reduce the number of variables that requires **decision** (try both true and false cases).

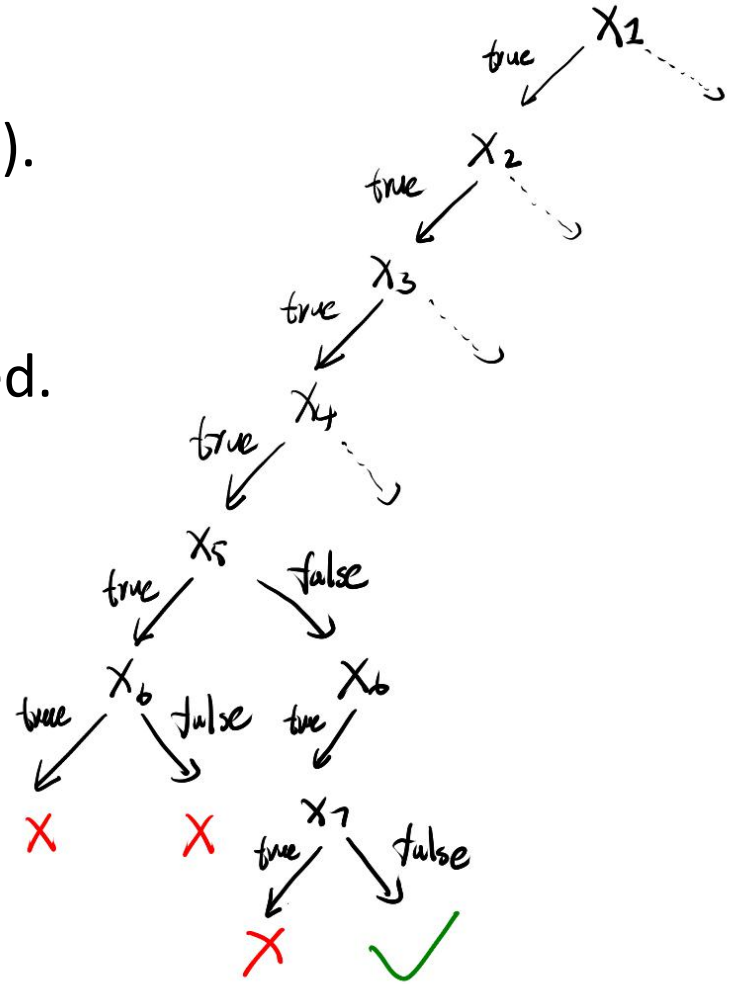

Davis Putnam Logemann Loveland Algorithm (DPLL) 1962

Backtracking with a few transformation rules to improve efficiency (reduce decision variables and search tree depth)

Transform the given formula α by applying a **sequence of satisfiability preserving rules**

If final result has **an empty clause** then *unsatisfiable*

if final result has **no clauses** then the formula is *satisfiable*

Transform 1: Unit propagation

A clause has a single literal

$$\alpha \equiv \ldots \wedge \ldots \wedge p \wedge \ldots \wedge \ldots$$

What choice do we really have?

$$\alpha \equiv \ldots \wedge (x_1 \vee \neg p \vee x_2) \wedge p \wedge \ldots \wedge (\neg x_3 \vee \neg p \vee x_1) \ldots$$

Transform 1: Unit propagation

A clause has a single literal

$$\alpha \equiv \ldots \wedge \ldots \wedge p \wedge \ldots \wedge \ldots$$

All clauses mentioning $\neg p$ have this literal deleted

All clauses mentioning p are deleted

$$\alpha \equiv \ldots \wedge (x_1 \vee \neg p \vee x_2) \wedge p \wedge \ldots \wedge (\neg x_3 \vee \neg p \vee x_1) \ldots$$
$$\alpha' \equiv \ldots \wedge (x_1 \vee x_2) \wedge \ldots \wedge (\neg x_3 \vee x_1) \ldots$$

α and α' are **equisatisfiable**

Transform 1: Unit propagation

How about

$$\alpha \equiv \ldots \wedge \ldots \wedge p \wedge \ldots \wedge (\neg p) \wedge \ldots$$

By deleting $\neg p$, we have an "**empty clause**" which means α is unsatisifiable

$$\alpha \equiv \ldots \wedge \ldots \wedge \ldots \wedge () \wedge \ldots$$

Transform 2: Pure literal

A literal appears only positively (or negatively) in α

$$\alpha \equiv \ldots \wedge (x_1 \vee \neg p \vee x_2) \wedge (x_4 \vee \neg p) \wedge \ldots \wedge (\neg x_3 \vee \neg p \vee x_1) \ldots$$

p does not appear anywhere; only $\neg p$ is used

Transform 2: Pure literal

A literal appears only positively (or negatively) in α

$$\alpha \equiv \ldots \wedge (x_1 \vee \neg p \vee x_2) \wedge (x_4 \vee \neg p) \wedge \ldots \wedge (\neg x_3 \vee \neg p \vee x_1) \ldots$$

p does not appear anywhere

Makes sense to set $p = 0$ and remove all occurrences of $\neg p$

Transform 2: Pure literal

A literal appears only positively (or negatively) in α

$$\alpha \equiv \ldots \wedge (x_1 \vee \neg p \vee x_2) \wedge (x_4 \vee \neg p) \wedge \ldots \wedge (\neg x_3 \vee \neg p \vee x_1) \wedge (\neg x_3 \vee x_4) \ldots$$

p does not appear anywhere

Makes sense to set $p = 0$ and remove all occurrences of $\neg p$

$$\alpha' \equiv \ldots \wedge \ldots \wedge \ldots \wedge (\neg x_3 \vee x_4) \ldots [p = 0]$$

α and α' are **equisatisfiable**

DPLL Algorithm: backtracking with unit-propagation and pure-literal assignment

```
function DPLL(α)
    α ← unit-propagate(α);
    α ← pure-literal-assign(α);
    // stopping conditions:
    if α is empty then return true;
    if α contains an empty clause then return false;
    l ← choose-literal(α);
    return DPLL(α ∧ {l}) or DPLL(α ∧ {¬l});
```

repeat until no further transformations possible

DPLL Algorithm example

$\alpha: = (\neg x_1 \vee x_2) \wedge (\neg x_3 \vee x_4)$
$\wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6)$
$\wedge (x_5 \vee \neg x_6 \vee \neg x_7) \wedge (x_1 \vee x_6 \vee x_7)$

```
function DPLL(α)
    unit-propagate
    pure-literal-assign
    check-stopping-conditions
    l ← choose-literal(α);
    return (DPLL(α ∧ {l}) or
            DPLL(α ∧ {¬l}));
```

Possible to apply unit propagation?

Possible to apply pure literal assignment?

We can essentially remove x_3 and x_4 from the search tree!

DPLL Algorithm example

$$\alpha: = (\neg x_1 \vee x_2) \wedge (\neg x_3 \vee x_4) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7) \wedge (x_1 \vee x_6 \vee x_7)$$

We decide to choose x_1 and search the two cases

```
function DPLL(α)
    unit-propagate
    pure-literal-assign
    check-stopping-conditions
    l ← choose-literal(α);
    return (DPLL(α ∧ {l}) or
            DPLL(α ∧ {¬l}));
```

DPLL Algorithm example

$(\neg x_1 \vee x_2) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$
$\wedge (x_1 \vee x_6 \vee x_7) \wedge x_1$

Possible to apply unit propagation? Always Yes!

DPLL Algorithm example

$(\neg x_1 \vee x_2) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$
$\wedge (x_1 \vee x_6 \vee x_7) \wedge x_1$

Possible to apply unit propagation? Always Yes!

$$x_2 \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Possible to apply unit propagation again?

DPLL Algorithm example

$(\neg x_1 \vee x_2) \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$
$\wedge (x_1 \vee x_6 \vee x_7) \wedge x_1$

Possible to apply unit propagation? Always Yes!

$x_2 \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$

Possible to apply unit propagation again?

$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$

Possible to apply unit propagation again?

DPLL Algorithm example

$$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Possible to apply pure-literal assignment?

DPLL Algorithm example

$$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Possible to apply pure-literal assignment?

$$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6)$$

Possible to apply pure-literal assignment again?

DPLL Algorithm example

$$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \wedge (x_5 \vee \neg x_6 \vee \neg x_7)$$

Possible to apply pure-literal assignment?

$$(\neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6)$$

Possible to apply pure-literal assignment again?

No clause left, we are done (return true).

With unit-propagation and pure-literal assignment, search process is much facster!

DPLL Algorithm

```
function DPLL(α)
    α ← unit-propagate(α);
    α ← pure-literal-assign(α);
    // stopping conditions:
    if α is empty then return true;
    if α contains an empty clause then return false;
    l ← choose-literal(α); // We decided to choose x1
    return DPLL(α ∧ {l}) or DPLL(α ∧ {¬l});
```

First condition returns true. No need to execute the second DPLL call.

DPLL Algorithm

```
function DPLL(α)
    α ← unit-propagate(α);
    α ← pure-literal-assign(α);
    // stopping conditions:
    if α is empty then return true;
    if α contains an empty clause then return false;
    l ← choose-literal(α); // We decided to choose x1
    return DPLL(α ∧ {l}) or DPLL(α ∧ {¬l});
```

repeat until no further transformations possible

The order of choosing literals is important - it usually defines the size of the search tree!

DPLL Algorithm: choosing literals

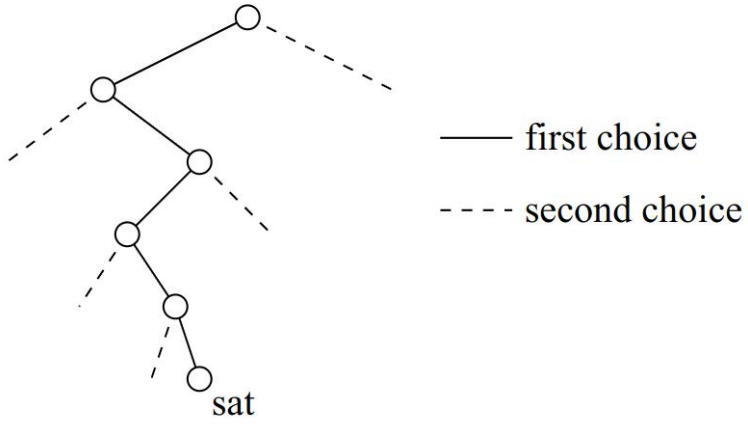

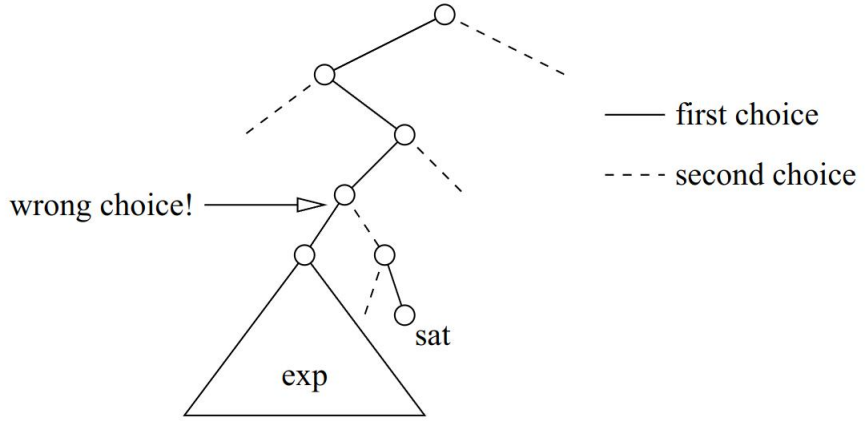

The order of choosing literals is important - it usually defines the size of the search tree!

Image from https://www.diag.uniroma1.it/~liberato/ar/dpll/dpll.html

DPLL Algorithm: choosing literals

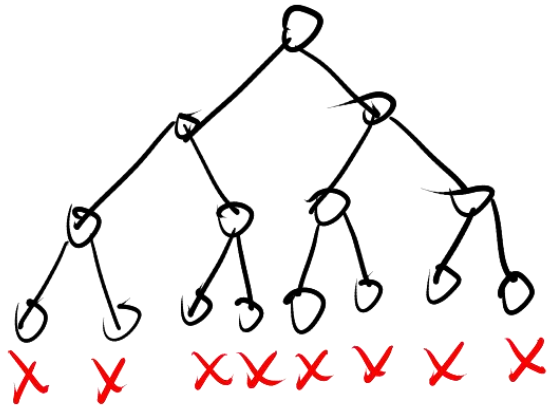

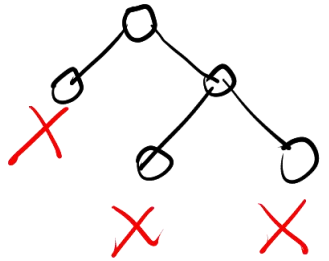

Proving unsatifiability is even harder. To explore the search tree faster, we want to find conflicts earlier. Roughly speaking, more clauses lead to more conflicts.

Practical Hardness of 3-SAT

Random SAT problems with **m** clauses, **n** variables
What if m >> n? What if n >> m?

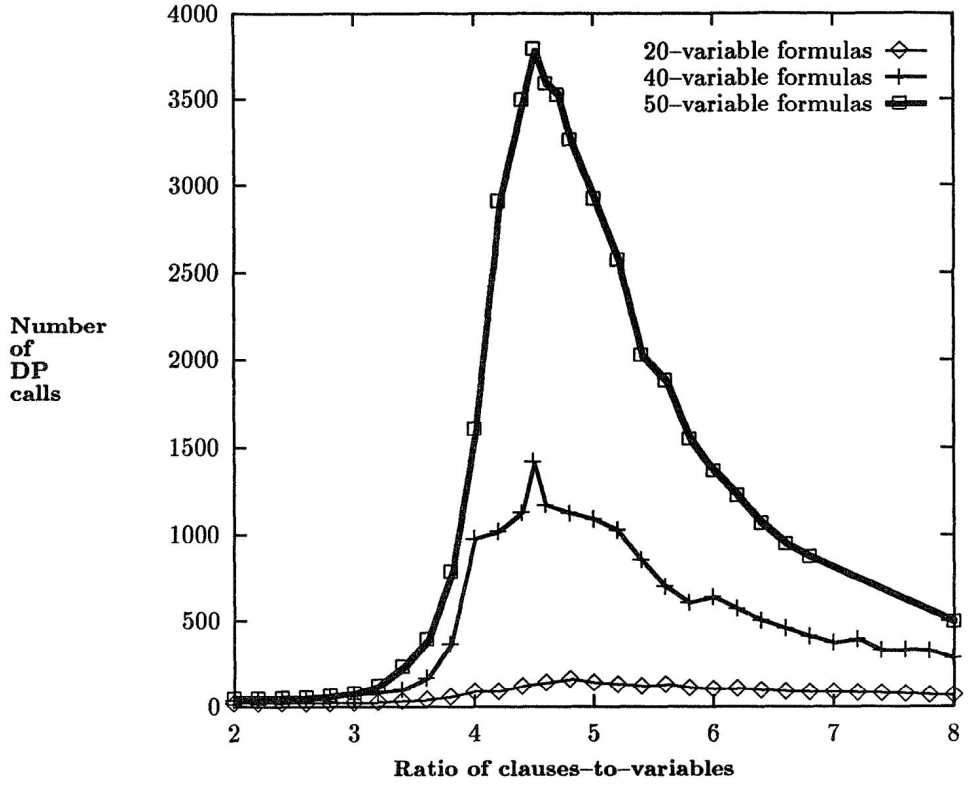

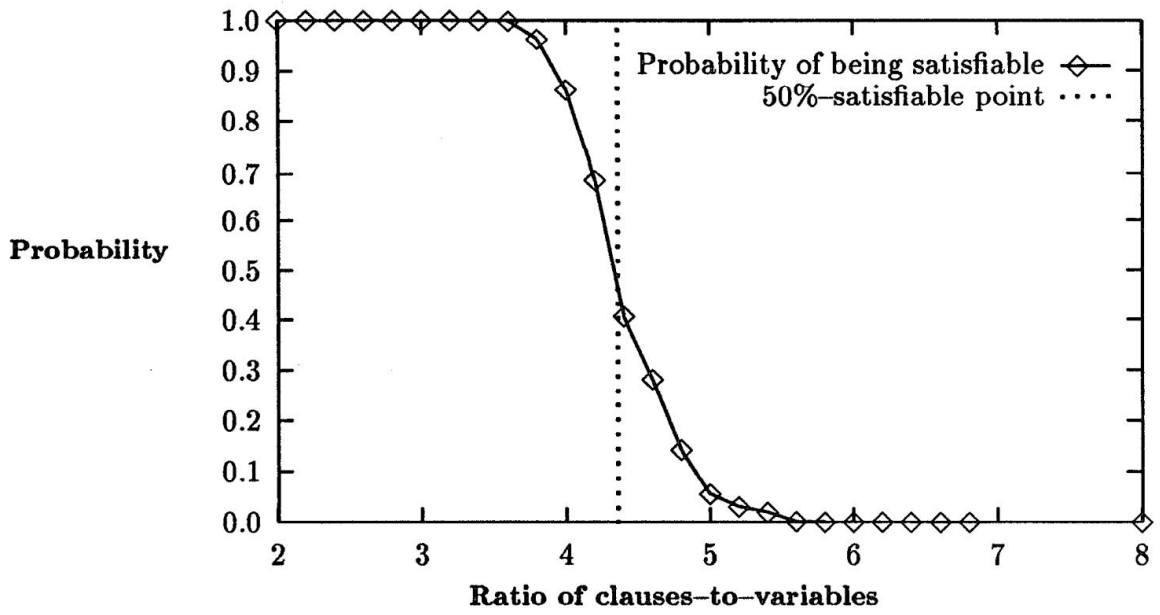

(50 variable formulas)

Mitchell, D., Selman, B., & Levesque, H. (1992). Hard and easy distributions of SAT problems.

Discuss: how can machine learning help to solve SAT?

$$\alpha := (\neg x_1 \vee x_2) \wedge (\neg x_3 \vee x_4) \\ \wedge (\neg x_2 \vee \neg x_5 \vee \neg x_6) \wedge (\neg x_5 \vee x_6) \\ \wedge (x_5 \vee x_7) \wedge (x_1 \vee x_6 \vee \neg x_7)$$

Discuss: how can machine learning help to solve SAT?

End-to-End solvers (e.g., input SAT instances, output sat/unsat results)

Helps part of the DPLL (e.g., variable selection, variable initialization, restarts policy, etc)

Class project ideas!

Checkout survey papers if you are interested.

Guo, Wenxuan, et al. "Machine learning methods in solving the boolean satisfiability problem." Machine Intelligence Research 20.5 (2023): 640-655.

End-to-End solving SAT problems: NeuroSAT

Input: literals (l) and clauses (c)

Use **graph neural networks** to connect literals to clauses

Each node has an embedding feature

message passing on the graph

$$\alpha: = (x_1 \vee x_2) \wedge (\neg x_1 \vee \neg x_2)$$

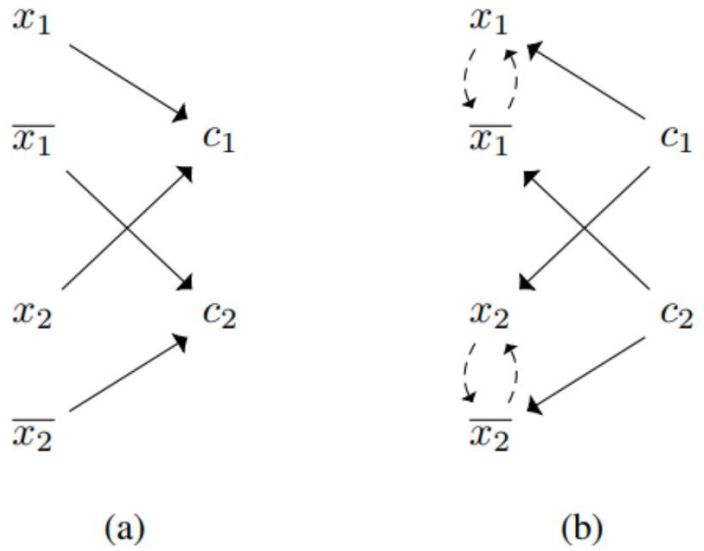

(a) Stage 1 (b) Stage 2

Selsam, Daniel, et al. "Learning a SAT Solver from Single-Bit Supervision." International Conference on Learning Representations.

End-to-End solving SAT problems: NeuroSAT

Each node has an **embedding feature**.

d: feature dimension, m: number of clauses; n: number of literals

$$L^{(t)} \in \mathbb{R}^{2n \times d} \qquad C^{(t)} \in \mathbb{R}^{m \times d}$$

Several NNs during message passing (M: **adjacency** matrix):

MLPs: $(\mathbf{L}_{\text{msg}}, \mathbf{C}_{\text{msg}}, \mathbf{L}_{\text{vote}})$ LSTMs: $(\mathbf{L}_{\text{u}}, \mathbf{C}_{\text{u}})$

$$(C^{(t+1)}, C_h^{(t+1)}) \leftarrow \mathbf{C}_{\text{u}}([C_h^{(t)}, M^\top \mathbf{L}_{\text{msg}}(L^{(t)})])$$

$$(L^{(t+1)}, L_h^{(t+1)}) \leftarrow \mathbf{L}_{\text{u}}([L_h^{(t)}, \text{Flip}(L^{(t)}), M\mathbf{C}_{\text{msg}}(C^{(t+1)})])$$

updated features and hidden states

Hidden states from previous time-step

Aggregate features from neighbors

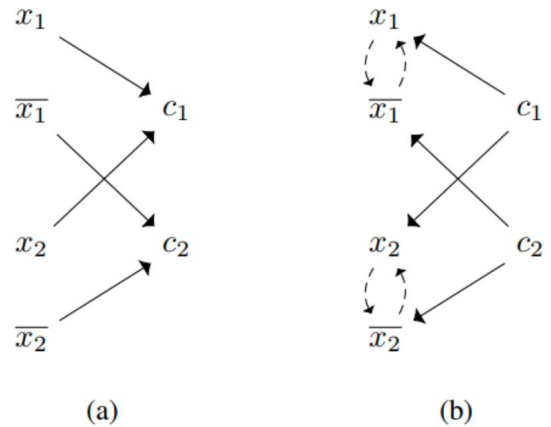

(a) Stage 1 (b) Stage 2

End-to-End solving SAT problems: NeuroSAT

$$(C^{(t+1)}, C_h^{(t+1)}) \leftarrow \mathbf{C}_{\text{u}}([C_h^{(t)}, M^\top \mathbf{L}_{\text{msg}}(L^{(t)})])$$

$$(L^{(t+1)}, L_h^{(t+1)}) \leftarrow \mathbf{L}_{\text{u}}([L_h^{(t)}, \text{Flip}(L^{(t)}), M\mathbf{C}_{\text{msg}}(C^{(t+1)})])$$

updated features and hidden states

Hidden states from previous time-step

Aggregate features from neighbors

After message passing for T steps, aggregate all literal features to produce a variable voting:

$$L_*^{(T)} \leftarrow \mathbf{L}_{\text{vote}}(L^{(T)}) \in \mathbb{R}^{2n}$$

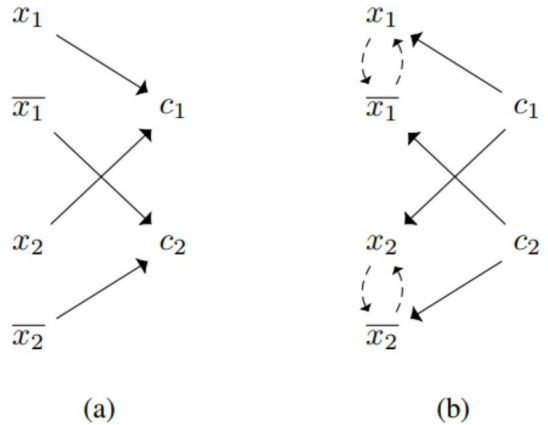

(a) Stage 1 (b) Stage 2

End-to-End solving SAT problems: NeuroSAT

Average vote: y is a simple average of $L_*^{(T)} \leftarrow \mathbf{L}_{\text{vote}}(L^{(T)}) \in \mathbb{R}^{2n}$

Training with randomly generated SAT instances denoted as SR(**n**), sat/unsat labels from a SAT solver

blue: vote for unsat; red: vote for sat

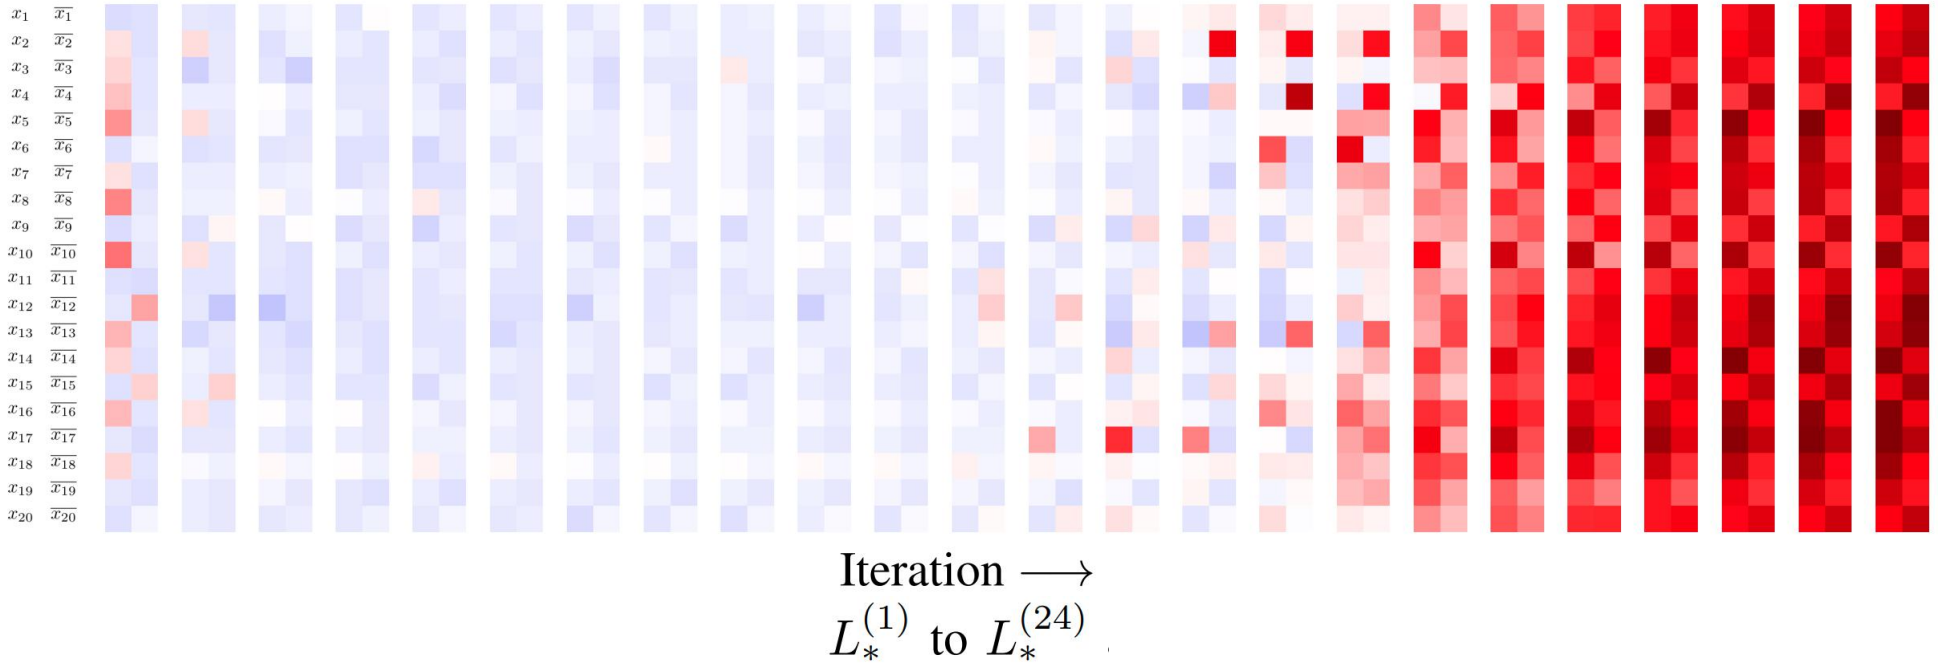

End-to-End solving SAT problems: NeuroSAT

Average vote: y is a simple average of $L_*^{(T)} \leftarrow \mathbf{L}_{\text{vote}}(L^{(T)}) \in \mathbb{R}^{2n}$

Training with randomly generated SAT instances denoted as SR(**n**), sat/unsat labels from a SAT solver

Trained on:	$\mathbf{SR}(\mathbf{U}(10, 40))$
Trained with:	26 iterations
Tested on:	$\mathbf{SR}(40)$
Tested with:	26 iterations
Overall test accuracy:	85%
Accuracy on *unsat* problems:	96%
Accuracy on *sat* problems:	73%
Percent of *sat* problems solved:	**70%**

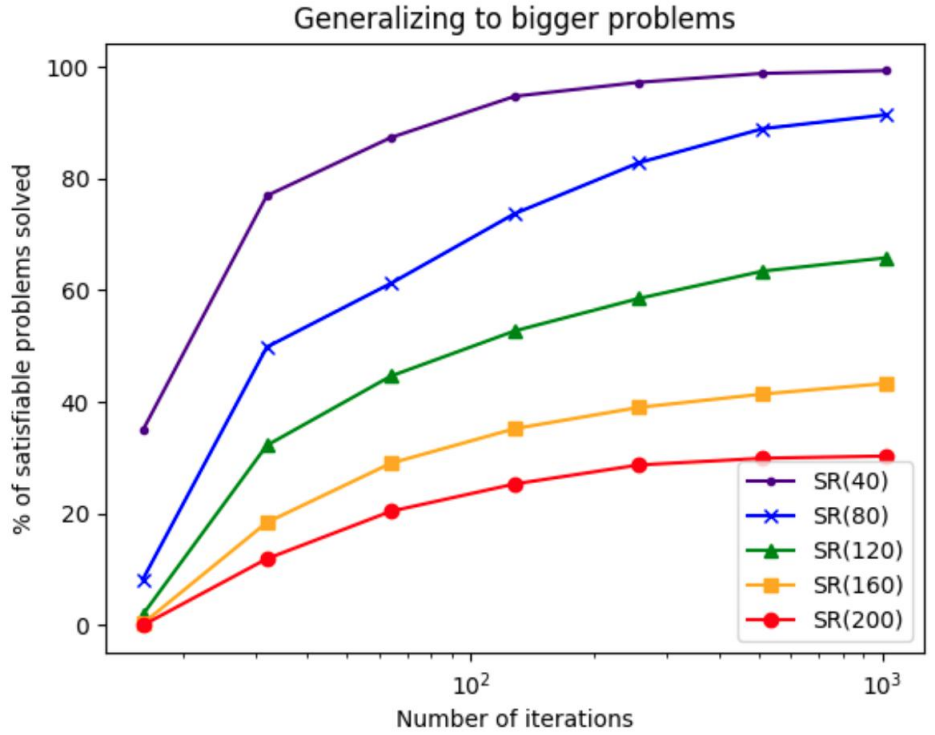

Assignments

- HW1 will be released on Feb 2 (due date after 2 weeks)
- Check Canvas for HW1

- Keep thinking about class projects! Form teams (2 people).

- Next lecture: SMT solving