Foundations of Artificial Intelligence

M. Helmert, T. Keller
S. Eriksson
Spring Term 2020

University of Basel
Computer Science

Exercise Sheet 13

Due: May 27, 2020

Important: For submission, consult the rules at the end of the exercise. Non-adherence to the rules will lead to your submission not being corrected.

All exercises on this sheet are bonus exercises.

The goal of this exercise is to compare the different problem solving methods covered in this course by looking at how they solve *propositional satisfiability* in the form of *3SAT*, i.e., decide satisifiability of a formula φ in disjunctive normal form where each clause contains three literals. We use the following notation:

- $\gamma \in \varphi$: clause γ occurs in φ
- $|\varphi|$: the number of clauses in φ
- $V = \{v_i \mid 1 \leq i \leq |V|\}$: the set of variables occuring in φ

Combinatorial optimization We define $P = \langle C, S, \min, v\rangle$ with

- $C = \{\langle c_1, \dots, c_{|V|}\rangle \mid c_i \in \{\mathbf{T}, \mathbf{F}\} \textit{ for all } 1 \leq i \leq |V|\}$,
- $S = \{c = \langle c_1, \dots, c_{|V|}\rangle \mid c \in C, \mathcal{I}(c) \models \varphi\}$,
- $v(c) = |\{\gamma \mid \gamma \in \varphi, \mathcal{I}(c) \not\models \gamma\}|$,

where $\mathcal{I}(c) = \{v_i \mapsto c_i \mid v_i \in V\}$ is an interpretation of V that assigns c_i to the corresponding v_i. The neighbors of $c \in C$ are those c' which differ to c in the assignment of exactly one c_i.

Constraint Network We consider a *ternary constraint network*, where the constraints are relations between three variables. We define $\mathcal{C} = \langle V, \mathrm{dom}, (R_{xyz})\rangle$, where

- V is the set of variables occuring in φ,
- $\mathrm{dom}(v) = \{\mathbf{T}, \mathbf{F}\}$, and
- $R_{xyz} = \{\langle v_x, v_y, v_z\rangle \mid \{x \mapsto v_x, y \mapsto v_y, z \mapsto v_z\} \models \gamma$ for all $\gamma \in \varphi$ over $x, y, z\}$.

Planning We define a STRIPS planning task $\Pi = \langle V', I, G, A\rangle$:

- $V' = \{u_i \mid v_i \in V\} \cup \{s_i \mid \gamma_i \in \varphi\}$, where u_i denotes whether v_i is still unassigned, and s_i denotes whether clause γ_i is satisfied by the current assignment.
- $I = \{u_i \mid v \in V\}$, i.e., all variables are still unassigned and no clause is satisfied yet.
- $G = \{s_i \mid \gamma_i \in \varphi\}$, i.e., all clauses must be satisfied.
- $A = \{\texttt{assign-v-i-val} \mid v_i \in V, \mathrm{val} \in \{\mathbf{T}, \mathbf{F}\}\}$, where for each action $a = \texttt{assign-v-i-val}$ we have

$$\begin{aligned} \mathit{pre}(a) &= \{u_i\} & \mathit{add}(a) &= \{s_i \mid \{v_i \mapsto \mathrm{val}\} \models \gamma_i\} \\ \mathit{cost}(a) &= 1 & \mathit{del}(a) &= \{u_i\}. \end{aligned}$$

Exercise 13.1 (4 bonus marks)

Compare the state spaces of DPLL, 3SAT as combinatorial optimization problem, 3SAT as ternary constraint network, and 3SAT as planning task. How many states does the states space of each formalization (at most) contain? Is it a tree?

Exercise 13.2 (1+1 bonus marks)

Compare the solutions found by hill climbing (for COP), simple backtracking (for constraint networks), DPLL and A* with an admissible heuristic (for planning):

(a) Which algorithms are guaranteed to find a solution if one exists?

(b) Which solution property (in terms of assignment to V) does A* optimize?

Exercise 13.3 (2 bonus marks)

Consider a generalized form of arc consistency for ternary constraint networks where a variable x is arc-consistent to a constraint R_{xyz} iff for each value in $\text{dom}(x)$ there is a tuple $\langle v_x, v_y, v_z \rangle \in R_{xyz}$ with $v_y \in \text{dom}(y)$ and $v_z \in \text{dom}(z)$. A constraint is arc consistent if all variables occuring in it are arc-consistent to it.
Compare *BacktrackingWithInference* using this generalized form of arc consistency, variable order "minimum remaining values" followed by "alphabetical" and value order $\mathbf{T} \prec \mathbf{F}$ with DPLL with splitting order "alphabetical" and value order $\mathbf{T} \prec \mathbf{F}$.

Exercise 13.4 (1+1 bonus marks)

Consider 3SAT as planning task Π (as defined above).

(a) What is $h^{\max}(\mathrm{I})$? *Hint: it is the same for any* φ

(b) We can build a collection of landmarks $\mathcal{L}$ with a landmark l for each clause $\gamma \in \varphi$ such that l contains all actions denoting variable assignments that satisfy γ. When we use the hitting set heuristic on $\mathcal{L}$ we can extract all actions that are part of the hitting set. Why does this set of actions not necessarily reflect a satisfying assignment for φ?

Exercise 13.5 (2 bonus marks)

Compare how good the four different algorithms are suited for solving 3SAT. Which one would you choose? List one advantage of each algorithm, and one disadvange of each one you would not choose.
Hint: the solution to the other parts of this exercise can and should be used here!

Submission rules:

- Upload a single PDF file (ending .pdf). If you want to submit handwritten parts, include their scans in the single PDF. Put the names of all group members on top of the first page. Use page numbers or put your names on each page. Make sure your PDF has size A4 (fits the page size if printed on A4).
- Only upload one submission per group. Do not upload several versions, i.e., if you need to resubmit, use the same file name again so that the previous submission is overwritten.