Artificial Intelligence 1 Winter Semester 2024/25

– Lecture Notes – Part II: General Problem Solving

Prof. Dr. Michael Kohlhase

Professur für Wissensrepräsentation und -verarbeitung Informatik, FAU Erlangen-Nürnberg Michael.Kohlhase@FAU.de

2025-02-06

This document contains Part II of the course notes for the course "Artificial Intelligence 1" held at FAU Erlangen-Nürnberg in the Winter Semesters 2016/17 ff. This part introduces search-based methods for general problem solving using atomic and factored representations of states.

Concretely, we discuss the basic techniques of search-based symbolic AI. First in the shape of classical and heuristic search and adversarial search paradigms. Then in constraint propagation, where we see the first instances of inference-based methods. Other parts of the lecture notes can be found at http://kwarc.info/teaching/AI/notes-*.pdf.

Contents

6	\mathbf{Pro}	blem Solving and Search 5
	6.1	Problem Solving
	6.2	Problem Types
	6.3	Search
	6.4	Uninformed Search Strategies
		6.4.1 Breadth-First Search Strategies
		6.4.2 Depth-First Search Strategies
		6.4.3 Further Topics
	6.5	Informed Search Strategies
		6.5.1 Greedy Search
		6.5.2 Heuristics and their Properties
		6.5.3 A-Star Search
		6.5.4 Finding Good Heuristics
	6.6	Local Search
7	Adv	versarial Search for Game Playing 49
	7.1	Introduction
	7.2	Minimax Search
	7.3	Evaluation Functions
	7.4	Alpha-Beta Search
	7.5	Monte-Carlo Tree Search (MCTS)
	7.6	State of the Art
	7.7	Conclusion
8	Cor	straint Satisfaction Problems 85
0	81	Constraint Satisfaction Problems: Motivation 85
	8.2	The Waltz Algorithm 90
	8.3	CSP: Towards a Formal Definition 93
	8.4	Constraint Networks: Formalizing Binary CSPs
	8.5	CSP as Search 98
	8.6	Conclusion & Preview 103
	0.0	
9	Cor	105 straint Propagation
	9.1	Introduction
	9.2	Constraint Propagation/Inference
	9.3	Forward Checking
	9.4	Arc Consistency
	9.5	Decomposition: Constraint Graphs, and Three Simple Cases
	9.6	Cutset Conditioning
	9.7	Constraint Propagation with Local Search
	9.8	Conclusion & Summary

CONTENTS

Chapter 6

Problem Solving and Search

In this chapter, we will look at a class of algorithms called search algorithms. These are algorithms that help in quite general situations, where there is a precisely described problem, that needs to be solved. Hence the name "General Problem Solving" for the area.

6.1 Problem Solving

A Video Nugget covering this section can be found at https://fau.tv/clip/id/21927. Before we come to the search algorithms themselves, we need to get a grip on the types of problems themselves and how we can represent them, and on what the various types entail for the problem solving process.

The first step is to classify the problem solving process by the amount of knowledge we have available. It makes a difference, whether we know all the factors involved in the problem before we actually are in the situation. In this case, we can solve the problem in the abstract, i.e. make a plan before we actually enter the situation (i.e. offline), and then when the problem arises, only execute the plan. If we do not have complete knowledge, then we can only make partial plans, and have to be in the situation to obtain new knowledge (e.g. by observing the effects of our actions or the actions of others). As this is much more difficult we will restrict ourselves to offline problem solving.



Definition 6.1.1. In offline problem solving an agent computing an action sequence based complete knowledge of the environment.
 Remark 6.1.2. Offline problem solving only works in fully observable, deterministic, static, and episodic environments.
 Definition 6.1.3. In online problem solving an agent computes one action at a time based on incoming perceptions.
 This Semester: We largely restrict ourselves to offline problem solving. (easier)

We will use the following problem as a running example. It is simple enough to fit on one slide and complex enough to show the relevant features of the problem solving algorithms we want to talk about.



Given this example to fortify our intuitions, we can now turn to the formal definition of problem formulation and their solutions.



6.1. PROBLEM SOLVING





Observation: The formulation of problems from **??** uses an atomic (black-box) state representation. It has enough functionality to construct the state space but nothing else. We will come back to this in slide **??**.

Remark 6.1.10. Note that search problems formalize problem formulations by making many of the implicit constraints explicit.

Structure Overview: Search Problem

▷ The structure overview for search problems:



We will now specialize **??** to deterministic, fully observable environments, i.e. environments where actions only have one – assured – outcome state.



6.2 Problem Types

Note that the definition of a search problem is very general, it applies to many many real-world problems. So we will try to characterize these by difficulty. A Video Nugget covering this section can be found at https://fau.tv/clip/id/21928.



⊳ Definit	▷ Definition 6.2.3. A search problem is called a contingency problem, iff								
⊳ the cont	environment is non deterr ingencies)	(solution can b	ranch, depen	ding on					
⊳ the s	state space is unknown ons)	(like a baby,	agent has to lea	arn about sta	tes and				
FAU	Michael Kohlhase: Artificial Intelligence 1		125	2025-02-06					

We will explain these problem types with another example. The problem \mathcal{P} is very simple: We have a vacuum cleaner and two rooms. The vacuum cleaner is in one room at a time. The floor can be dirty or clean.

The possible states are determined by the position of the vacuum cleaner and the information, whether each room is dirty or not. Obviously, there are eight states: $S = \{1, 2, 3, 4, 5, 6, 7, 8\}$ for simplicity.

The goal is to have both rooms clean, the vacuum cleaner can be anywhere. So the set \mathcal{G} of goal states is $\{7,8\}$. In the single-state version of the problem, [right, suck] shortest solution, but [suck, right, suck] is also one. In the multiple-state version we have

 $[right\{2,4,6,8\}, suck\{4,8\}, left\{3,7\}, suck\{7\}]$



Example: Vacuum-Cleaner World (continued)

▷ Contingency Problem:



In the contingency version of \mathcal{P} a solution is the following:

 $[suck\{5,7\}, right \to \{6,8\}, suck \to \{6,8\}, suck\{5,7\}]$

etc. Of course, local sensing can help: narrow $\{6,8\}$ to $\{6\}$ or $\{8\}$, if we are in the first, then suck.

Single-state problem formula	tion
▷ Defined by the following four items	5
1. Initial state:	(e.g. Arad)
2. Successor function $S_a(s)$: (e.g.	$S_{goZer} = \{(Arad, Zerind), (goSib, Sibiu), \dots\}$
3. Goal test:	(e.g. $x = Bucharest$ (explicit test)) noDirt(x) (implicit test)
4. Path cost (optional):(e.g. sum of	f distances, number of operators executed, etc.)
▷ Solution: A sequence of actions lea	ading from the initial state to a goal state.
FAU Michael Kohlhase: Artificial Intelligence 1	128 2025-02-06 CONTRACTOR

"Path cost": There may be more than one solution and we might want to have the "best" one in a certain sense.



6.2. PROBLEM TYPES

FAU Michael Kohlhase: Artificial Intelligence 1 129 2025-02-06

"State": e.g., we don't care about tourist attractions found in the cities along the way. But this is problem dependent. In a different problem it may well be appropriate to include such information in the notion of state.

"Realizability": one could also say that the abstraction must be sound wrt. reality.



How many states are there? N factorial, so it is not obvious that the problem is in NP. One needs to show, for example, that polynomial length solutions do always exist. Can be done by combinatorial arguments on state space graph (really ?).

Some rule-books give a different goal state for the 8-puzzle: starting with 1, 2, 3 in the top row and having the hold in the lower right corner. This is completely irrelevant for the example and its significance to AI-1.





e

2025-02-06



6.3 Search

Fau

A Video Nugget covering this section can be found at https://fau.tv/clip/id/21956.

133

Michael Kohlhase: Artificial Intelligence 1



6.3. SEARCH



Fau

Michael Kohlhase: Artificial Intelligence 1

134

2025-02-06



Let us now think a bit more about the implementation of tree search algorithms based on the ideas discussed above. The abstract, mathematical notions of a search problem and the induced tree search algorithm gets further refined here.

Implementation: States vs. nodes



It is very important to understand the fundamental difference between a state in a search problem, a node search tree employed by the tree search algorithm, and the implementation in a search tree node. The implementation above is faithful in the sense that the implemented data structures contain all the information needed in the tree search algorithm.

So we can use it to refine the idea of a tree search algorithm into an implementation.

Implementation of Search Algorithms					
Definition 6.2.4 (Involumented Tree Second Algorithms)					
> Definition 0.3.4 (Implemented Tree Search Algorithm).					
procedure Tree Search (problem,strategy)					
<pre>fringe := insert(make_node(initial_state(problem)))</pre>					
Іоор					
if empty(fringe) fail end if					
node := first(fringe,strategy)					
if GoalTest(node) return node					
<pre>else fringe := insert(expand(node,problem))</pre>					
end if					
end loop					
end procedure					
The fringe is the set of search tree nodes not yet expanded in tree search.					
\triangleright Idea: We treat the fringe as an abstract data type with three accessors: the					
▷ binary function first retrieves an element from the fringe according to a strategy.					
▷ binary function insert adds a (set of) search tree node into a fringe.					
ightarrow unary predicate empty to determine whether a fringe is the empty set.					
\triangleright The strategy determines the behavior of the fringe (data structure) (see below)					
Michael Kohlhase: Artificial Intelligence 1 137 2025-02-06					

6.4. UNINFORMED SEARCH STRATEGIES

Note: The pseudocode in **??** is still relatively underspecified – leaves many implementation details unspecified. Here are the specifications of the functions used without.

- •
- make node constructs a search tree node from a state.
- initial state accesses the initial state of a search problem.
- State returns the state associated with its aregument.
- GoalNode checks whether its argument is a goal node
- expand = creates new search tree nodes by for all successor states.

Essentially, only the first function is non-trivial (as the strategy argument shows) In fact it is the only place, where the strategy is used in the algorithm.

An alternative implementation would have been to make the fringe a queue, and insert order the fringe as the strategy sees fit. Then first can just return the first element of the queue. This would have lead to a different signature, possibly different runtimes, but the same overall result of the algorithm.

Search strategies ▷ **Definition 6.3.5.** A strategy is a function that picks a node from the fringe of a search tree. (equivalently, orders the fringe and picks the first.) ▷ Definition 6.3.6 (Important Properties of Strategies). completeness does it always find a solution if one exists? number of nodes generated/expanded time complexity maximum number of nodes in memory space complexity optimality does it always find a least cost solution? > Time and space complexity measured in terms of: b maximum branching factor of the search tree dminimal graph depth of a solution in the search tree maximum graph depth of the search tree (may be ∞) mComplexity means here always worst-case complexity! FAU e Michael Kohlhase: Artificial Intelligence 1 138 2025-02-06

Note that there can be infinite branches, see the search tree for Romania.

6.4 Uninformed Search Strategies

Video Nuggets covering this section can be found at https://fau.tv/clip/id/21994 and https://fau.tv/clip/id/21995.

Uninformed search strategies

Definition 6.4.1. We speak of an uninformed search algorithm, if it only uses the information available in the problem definition.

⊳ Next:	Frequently used search algorithm	ns				
⊳ Bre	adth first search					
▷ Uniform cost search						
⊳ Depth first search						
⊳ Dep	oth limited search					
⊳ Iter	ative deepening search					
Fau	Michael Kohlhase: Artificial Intelligence 1	139	2025-02-06			

The opposite of uninformed search is informed or heuristic search that uses a heuristic function that adds external guidance to the search process. In the Romania example, one could add the heuristic to prefer cities that lie in the general direction of the goal (here SE).

Even though heuristic search is usually much more efficient, uninformed search is important nonetheless, because many problems do not allow to extract good heuristics.

6.4.1 Breadth-First Search Strategies





We will now apply the breadth first search strategy to our running example: Traveling in Romania. Note that we leave out the green dashed nodes that allow us a preview over what the search tree will look like (if expanded). This gives a much cleaner picture we assume that the readers already have grasped the mechanism sufficiently.





Breadth-first search: Properties Completeness Yes (if b is finite) $1+b+b^2+b^3+\ldots+b^d$, so $\mathcal{O}(b^d)$, i.e. exponential Time complexity in d \triangleright Space complexity $\mathcal{O}(b^d)$ (fringe may be whole level) Yes (if cost = 1 per step), not optimal in general Optimality > **Disadvantage:** Space is the big problem (can easily generate nodes at $500MB/sec \cong 1.8TB/h$) > Optimal?: No! If cost varies for different steps, there might be better solutions below the level of the first one. > An alternative is to generate all solutions and then pick an optimal one. This works only, if m is finite. FAU

The next idea is to let cost drive the search. For this, we will need a non-trivial cost function: we will take the distance between cities, since this is very natural. Alternatives would be the driving time, train ticket cost, or the number of tourist attractions along the way.

142

C

2025-02-06

Of course we need to update our problem formulation with the necessary information.

Michael Kohlhase: Artificial Intelligence 1





Note that we must sum the distances to each leaf. That is, we go back to the first level after the third step.

Uniform-cost searc	ch: Properties			
Completeness	Yes (if step costs	$\geq \epsilon > 0$)		
Time complexity number of nodes with path cost less than that of opti-				ti-
	mal solution			
Space complexity	ditto			
Optimality	Yes			
Michael Kohlhase:	Artificial Intelligence 1	145	2025-02-06	CONTRACTOR AND A CONTRACT OF A

If step cost is negative, the same situation as in breadth first search can occur: later solutions may be cheaper than the current one.

If step cost is 0, one can run into infinite branches. UCS then degenerates into depth first search, the next kind of search algorithm we will encounter. Even if we have infinite branches, where the sum of step costs converges, we can get into trouble, since the search is forced down these infinite paths before a solution can be found.

Worst case is often worse than BFS, because large trees with small steps tend to be searched first. If step costs are uniform, it degenerates to BFS.

6.4.2 Depth-First Search Strategies

Depth-first Search

- ▷ **Idea:** Expand deepest unexpanded node.
- ▷ Definition 6.4.6. Depth-first search (DFS) is the strategy where the fringe is organized as a (LIFO) stack i.e. successors go in at front of the fringe.
- ▷ Definition 6.4.7. Every node that is pushed to the stack is called a backtrack point. The action of popping a non-goal node from the stack and continuing the search with the new top element of the stack (a backtrack point by construction) is called backtracking, and correspondingly the DFS algorithm backtracking search.

6.4. UNINFORMED SEARCH STRATEGIES









Depth-first search: Properties

	Completeness	Yes: if search tree finite
		No: if search tree contains infinite paths or
		loops
	Time complexity	$\mathcal{O}(b^m)$
		(we need to explore until max depth m in any
\triangleright		case!)
	Space complexity	$\mathcal{O}(bm)$ (i.e. linear space)
		(need at most store m levels and at each level
		at most <i>b</i> nodes)
	No (there can be many better solutions in the	
		unexplored part of the search tree)

 \triangleright **Disadvantage:** Time terrible if m much larger than d.

▷ Advantage: Time may be much less than breadth first search if solutions are dense.

6.4. UNINFORMED SEARCH STRATEGIES





25

CHAPTER 6. PROBLEM SOLVING AND SEARCH

CC State Blands Resistance

2025-02-06



▷ Consequence: IDS used in practice for search spaces of large, infinite, or unknown depth.

152

Note: To find a solution (at depth d) we have to search the whole tree up to d. Of course since we do not save the search state, we have to re-compute the upper part of the tree for the next level. This seems like a great waste of resources at first, however, IDS tries to be complete without the space penalties.

However, the space complexity is as good as DFS, since we are using DFS along the way. Like in BFS, the whole tree on level d (of optimal solution) is explored, so optimality is inherited from there. Like BFS, one can modify this to incorporate uniform cost search behavior.

As a consequence, variants of IDS are the method of choice if we do not have additional information.



 \triangleright **Example 6.4.12.** IDS may fail to be be optimal at step sizes > 1.

Michael Kohlhase: Artificial Intelligence 1

FAU



6.4.3 Further Topics



Uninformed Search Summary

> Tree/Graph Search Algorithms: Systematically explore the state tree/graph

induced by a search problem in search of a goal state. Search strategies only differ by the treatment of the fringe.

Criterion	Breadth first	Uniform cost	first	lterative deepening
Completeness	Yes ¹	Yes ²	No	Yes
Time complexity	b^d	$pprox b^d$	b^m	b^{d+1}
Space complexity	b^d	$pprox b^d$	bm	bd
Optimality	Yes*	Yes	No	Yes*
Conditions	¹ b finite	$^2 0 < \epsilon \le$	cost	



Informed Search Strategies 6.5



FAL

6.5. INFORMED SEARCH STRATEGIES

other unir	oformed algorithms.						
⊳ Next Ste	▷ Next Step: Introduce additional knowledge about the problem (heuristic search)						
⊳ Best-f ⊳ Iterati	irst-, A^* -strategies ve improvement algorithms.	(guide	the search by h	euristics)			
▷ Definitio external ir	on 6.5.1. A search algorithm is not part of	s called informed, the search proble	iff it uses some em – to guide th	e form of e search.			
Fau	Michael Kohlhase: Artificial Intelligence 1	157	2025-02-06	CONTRACTOR OF CO			

6.5.1 Greedy Search

A Video Nugget covering this subsection can be found at https://fau.tv/clip/id/22015.

Best-first search					
Idea: Order the fringe by estimated "desirability" (Expand model)	ost desirable				
Definition 6.5.2. An evaluation function assigns a desirability value to of the search tree.	o each node				
▷ Note: A evaluation function is not part of the search problem, but must be added externally.					
▷ Definition 6.5.3. In best first search, the fringe is a queue sorted in decreasing order of desirability.					
\triangleright Special cases: Greedy search, A^* search					
Michael Kohlhase: Artificial Intelligence 1 158 2025-02-06	CONSTRUCTION OF THE SECOND				

This is like UCS, but with an evaluation function related to problem at hand replacing the path cost function.

If the heuristic is arbitrary, we expect incompleteness! Depends on how we measure "desirability". Concrete examples follow.

Greedy search

- ▷ **Idea:** Expand the node that *appears* to be closest to the goal.
- \triangleright **Definition 6.5.4.** A heuristic is an evaluation function h on states that estimates the cost from n to the nearest goal state. We speak of heuristic search if the search algorithm uses a heuristic in some way.
- \triangleright **Note:** All nodes for the same state must have the same *h*-value!
- \triangleright **Definition 6.5.5.** Given a heuristic *h*, greedy search is the strategy where the fringe is organized as a queue sorted by increasing *h* value.
- ▷ **Example 6.5.6.** Straight-line distance from/to Bucharest.



In greedy search we replace the *objective* cost to *construct* the current solution with a heuristic or *subjective* measure from which we think it gives a good idea how far we are from a solution. Two things have shifted:

- we went from internal (determined only by features inherent in the search space) to an external/heuristic cost
- instead of measuring the cost to build the current partial solution, we estimate how far we are from the desired goal

Romania with Straight-Line Distances

 \triangleright Example 6.5.7 (Informed Travel). $h_{SLD}(n) = straight - line \ distance \ to \ Bucharest$

Arad	366	Mehadia	241	Bucharest	0	Neamt	234
Craiova	160	Oradea	380	Drobeta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193	Fragaras	176	Sibiu	253
Giurgiu	77	Timisoara	329	Hirsova	151	Urziceni	80
lasi	226	Vaslui	199	Lugoj	244	Zerind	374



Greedy Search: Romania



Let us fortify our intuitions with another example: navigation in a simple maze. Here the states are the cells in the grid underlying the maze and the actions navigating to one of the adjoining cells. The initial and goal states are the left upper and right lower corners of the grid. To see the influence of the chosen heuristic (indicated by the red number in the cell), we compare the search induced goal distance function with a heuristic based on the Manhattan distance. Just follow the greedy search by following the heuristic gradient.





Not surprisingly, the first maze is searchless, since we are guided by the perfect heuristic. In cases, where there is a choice, the this has no influence on the length (or in other cases cost) of the solution.

In the "good case" example, greedy search performs well, but there is some limited backtracking needed, for instance when exploring the left lower corner 3×3 area before climbing over the second wall.

In the "bad case", greedy search is led down the lower garden path, which has a dead end, and does not lead to the goal. This suggests that there we can construct adversary examples – i.e. example mazes where we can force greedy search into arbitrarily bad performance.





Remark 6.5.12. Greedy search is similar to UCS. Unlike the latter, the node evaluation function has nothing to do with the nodes explored so far. This can prevent nodes from being enumerated systematically as they are in UCS and BFS.

For completeness, we need repeated state checking as the example shows. This enforces complete enumeration of the state space (provided that it is finite), and thus gives us completeness.

Note that nothing prevents from *all* nodes being searched in worst case; e.g. if the heuristic function gives us the same (low) estimate on all nodes except where the heuristic mis-estimates the distance to be high. So in the worst case, greedy search is even worse than BFS, where d (depth of first solution) replaces m.

The search procedure cannot be optimal, since actual cost of solution is not considered.

For both, completeness and optimality, therefore, it is necessary to take the actual cost of partial solutions, i.e. the path cost, into account. This way, paths that are known to be expensive are avoided.

6.5.2 Heuristics and their Properties

A Video Nugget covering this subsection can be found at https://fau.tv/clip/id/22019.

Heuristic Functions

- ▷ **Definition 6.5.13.** Let Π be a search problem with states S. A heuristic function (or short heuristic) for Π is a function $h: S \to \mathbb{R}_0^+ \cup \{\infty\}$ so that h(s) = 0 whenever s is a goal state.
- \triangleright h(s) is intended as an estimate the distance between state s and the nearest goal state.
- ▷ **Definition 6.5.14.** Let Π be a search problem with states S, then the function $h^*: S \to \mathbb{R}^+_0 \cup \{\infty\}$, where $h^*(s)$ is the cost of a cheapest path from s to a goal state, or ∞ if no such path exists, is called the goal distance function for Π .
- ▷ **Notes:**
 - harpoon harbor har

CHAPTER 6. PROBLEM SOLVING AND SEARCH

- \triangleright Return value ∞ : To indicate dead ends, from which the goal state can't be reached anymore.
- \triangleright The distance estimate depends only on the state *s*, not on the node (i.e., the path we took to reach *s*).

Michael Kohlhase: Artificial Intelligence 1

164

2025-02-06

2025-02-06

Where does the word "Heuristic" come from? ▷ Ancient Greek word ευρισκειν (= "I find") (aka. ευρεκα!) ▷ Popularized in modern science by George Polya: "How to solve it" [Pól73] ▷ Same word often used for "rule of thumb" or "imprecise solution method".

Heuristic Functions: The Eternal Trade-Off

- \triangleright "Distance Estimate"? (*h* is an arbitrary function in principle)
 - ▷ In practice, we want it to be accurate (aka: informative), i.e., close to the actual goal distance.
 - $_{\vartriangleright}$ We also want it to be fast, i.e., a small overhead for computing h.
 - ▷ These two wishes are in contradiction!
- ▷ Example 6.5.15 (Extreme cases).
 - $\rhd h=0:$ no overhead at all, completely un-informative.
 - $h = h^*$: perfectly accurate, overhead $\hat{=}$ solving the problem in the first place.
- ▷ Observation 6.5.16. We need to trade off the accuracy of h against the overhead for computing it.

166

FAU

Properties of Heuristic Functions

Michael Kohlhase: Artificial Intelligence 1

 \triangleright Definition 6.5.17. Let Π be a search problem with states S and actions A. We say that a heuristic h for Π is admissible if $h(s) \leq h^*(s)$ for all $s \in S$.

We say that h is consistent if $h(s) - h(s') \leq c(a)$ for all $s \in S$, $a \in A$, and $s' \in \mathcal{T}(s, a)$.

\triangleright In other words ...:

- \triangleright *h* is admissible if it is a lower bound on goal distance.
- \triangleright *h* is consistent if, when applying an action *a*, the heuristic value cannot decrease by more than the cost of *a*.

6.5. INFORMED SEARCH STRATEGIES





6.5.3 A-Star Search

A Video Nugget covering this subsection can be found at https://fau.tv/clip/id/22020.

A* Search: Evaluation Function

Idea: Avoid expanding paths that are already expensive(make use of actual cost) The simplest way to combine heuristic and path cost is to simply add them. 35



This works, provided that h does not overestimate the true cost to achieve the goal. In other words, h must be *optimistic* wrt. the real cost h^* . If we are too pessimistic, then non-optimal solutions have a chance.






To extend our intuitions about informed search algorithms to A^* -search, we take up the maze examples from above again. We first show the good maze with Manhattan distance again.

Additional Observations (Not Limited to Path Planning)
Example 6.5.23 (Greedy best-first search, "good case").



To compare it to A^* -search, here is the same maze but now with the numbers in red for the evaluation function f where h is the Manhattan distance.

Additional Observations (Not Limited to Path Planning)																	
⊳ Examp	le 6.	5.24	4 (<i>A</i>	* (g	+h), "g	good	cas	e").								
I	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		
1	18		22	22	22	22	22	_	24	24	24	24	24	24	24		
2	18		20	20	20		22	22	22	22	22	22	22	22	22		
3	18	18	18		20		22	22	22		L						
4	18	18	18		20	20	20		22	22		24	24	24	24		
5	18	18	18		20	20		24	22	22	22	22		24	24		
																G	
⊳ In <i>A</i> cann	* wi ot de	th a ecrea	cons ase n	sister nore	nt he thar	eurist n <i>g</i> ii	tic, g	ı + h ases)	alw	ays i	ncre	ases	mor	notoi	nicall	y (h	
\triangleright We need more search, in the "right upper half". This is typical: Greedy best first search tends to be faster than A^* .																	
Fau	Mic	hael Ko	hlhase: /	Artificial	Intellige	nce 1				174			2	2025-02-	06	COMERCIAL DESERVED]

Let's now consider the "bad maze" with Manhattan distance again.





And we compare it to A^* -search; again the numbers in red are for the evaluation function f.



Finally, we compare that with the goal distance function for the "bad maze". Here we see that the lower garden path is under-estimated by the evaluation function f, but still large enough to keep the search out of it, thanks to the admissibility of the Manhattan distance.

Additional Observations (Not Limited to Path Planning) \triangleright Example 6.5.27 (A^* (g + h) using h^*).



A^* search: *f*-contours

 \triangleright **Intuition:** A^* -search gradually adds "*f*-contours" (areas of the same *f*-value) to the search.



A^* search: Properties

 \triangleright Properties or A^* -search:

Completeness	Yes (unless there are infinitely many nodes n
	with $f(n) \leq f(0)$)
Time complexity	Exponential in [relative error in h $ imes$ length of
	solution]
Space complexity	Same as time (variant of BFS)
Optimality	Yes

6.5. INFORMED SEARCH STRATEGIES

 $\triangleright A^*\text{-search expands all (some/no) nodes with } f(n) < h^*(n)$ $\triangleright \text{ The run-time depends on how well we approximated the real cost } h^* \text{ with } h.$ $\blacksquare Michael Kohlhase: Artificial Intelligence 1 179 2025-02-06 \blacksquare 179 2025-02-06 = 179 2025-02-$

6.5.4 Finding Good Heuristics

A Video Nugget covering this subsection can be found at https://fau.tv/clip/id/22021. Since the availability of admissible heuristics is so important for informed search (particularly for A^* -search), let us see how such heuristics can be obtained in practice. We will look at an example, and then derive a general procedure from that.



Actually, the crucial difference between the heuristics h_1 and h_2 is that – not only in the example configuration above, but for all configurations – the value of the latter is larger than that of the former. We will explore this next.



We now try to generalize these insights into (the beginnings of) a general method for obtaining

admissible heuristics.



Relaxation means to remove some of the constraints or requirements of the original problem, so that a solution becomes easy to find. Then the cost of this easy solution can be used as an optimistic approximation of the problem.



42

6.6 Local Search

Video Nuggets covering this section can be found at https://fau.tv/clip/id/22050 and https://fau.tv/clip/id/22051.





Local Search: Iterative improvement algorithms

▷ Definition 6.6.7. The traveling salesman problem (TSP is to find shortest trip through set of cities such that each city is visited exactly once.





In order to understand the procedure on a more intuitive level, let us consider the following scenario: We are in a dark landscape (or we are blind), and we want to find the highest hill. The search procedure above tells us to start our search anywhere, and for every step first feel around, and then take a step into the direction with the steepest ascent. If we reach a place, where the next step would take us down, we are finished.

Of course, this will only get us into local maxima, and has no guarantee of getting us into global ones (remember, we are blind). The solution to this problem is to re-start the search at random (we do not have any information) places, and hope that one of the random jumps will get us to a slope that leads to a global maximum.

Example Hill Climbing with 8 Queens



Recent work on hill climbing algorithms tries to combine complete search with randomization to escape certain odd phenomena occurring in statistical distribution of solutions.













Genetic algorithms (continued)

- ▷ **Problem:** Genetic algorithms require states encoded as strings.
- ▷ Crossover only helps iff substrings are meaningful components.
- ▷ Example 6.6.17 (Evolving 8 Queens). First crossover



Chapter 7

Adversarial Search for Game Playing

A Video Nugget covering this chapter can be found at https://fau.tv/clip/id/22079.

7.1 Introduction

Video Nuggets covering this section can be found at https://fau.tv/clip/id/22060 and https://fau.tv/clip/id/22061.



▷ What do you think?

- ▷ Playing a game well clearly requires a form of "intelligence".
- ▷ Games capture a pure form of competition between opponents.
- \triangleright Games are abstract and precisely defined, thus very easy to formalize.
- \triangleright Game playing is one of the oldest sub-areas of AI (ca. 1950).
- \triangleright The dream of a machine that plays chess is, indeed, *much* older than Al!



"Game" Playing? *Which* Games?

- $\triangleright \ldots$ sorry, we're not gonna do soccer here.
- ▷ Definition 7.1.3 (Restrictions). A game in the sense of Al-1 is one where
 - ▷ Game state discrete, number of game state finite.
 - ▷ Finite number of possible moves.
 - \triangleright The game state is fully observable.
 - \triangleright The outcome of each move is deterministic.
 - \triangleright Two players: Max and Min.
 - \triangleright Turn-taking: It's each player's turn alternatingly. Max begins.
 - \triangleright Terminal game states have a utility u. Max tries to maximize u, Min tries to minimize u.
 - \triangleright In that sense, the utility for Min is the exact opposite of the utility for Max ("zero sum").
 - ▷ There are no infinite runs of the game (no matter what moves are chosen, a terminal state is reached after a finite number of moves).

FAU Michael Kohlhase: Artificial Intelligence 1 198 2025-02-06

7.1. INTRODUCTION





(A Brief Note On) Formalization

- \triangleright **Definition 7.1.4.** An adversarial search problem is a search problem $\langle S, A, T, I, G \rangle$, where
 - 1. $S = S^{Max} \uplus S^{Min} \uplus G$ and $A = A^{Max} \uplus A^{Min}$
 - 2. For $a \in \mathcal{A}^{\text{Max}}$, if $s \xrightarrow{a} s'$ then $s \in \mathcal{S}^{\text{Max}}$ and $s' \in (\mathcal{S}^{\text{Min}} \cup \mathcal{G})$.
 - 3. For $a \in \mathcal{A}^{\mathrm{Min}}$, if $s \xrightarrow{a} s'$ then $s \in \mathcal{S}^{\mathrm{Min}}$ and $s' \in (\mathcal{S}^{\mathrm{Max}} \cup \mathcal{G})$.

together with a game utility function $u: \mathcal{G} \to \mathbb{R}$. (the "score" of the game)

- ▷ Remark: A round of the game one move Max, one move Min is often referred to as a "move", and individual actions as "half-moves" (we don't in Al-1)

CHAPTER 7. ADVERSARIAL SEARCH FOR GAME PLAYING

FAU	Michael Kohlhase: Artificial Intelligence 1	201	2025-02-06	STATE ATCHING RESERVED			
Why Gam	nes are Hard to Solve:						
\triangleright What is a	a "solution" here?						
$\triangleright \begin{array}{l} \textbf{Definition} \\ \textbf{A strategy} \\ \sigma^X(s) = 0 \end{array}$	on 7.1.6. Let Θ be an adversariate of T for X is a function $\sigma^X : S^X - a$.	al search problem, $ ightarrow \mathcal{A}^X$ so that a is	, and let $X \in \{ \mathrm{Mat} \$ s applicable to s v	$\max, Min\}.$			
⊳ We don't	\triangleright We don't know how the opponent will react, and need to prepare for all possibilities.						
\triangleright Definition 7.1.7. A strategy is called optimal if it yields the best possible utility for X assuming perfect opponent play (not formalized here).							
⊳ Problem (state/sea	: In (almost) all games, comp arch tree too huge)	outing an optimal	strategy is infeas	sible.			
⊳ Solution	: Compute the next move "on	demand", given	the current state	instead.			
FAU	Michael Kohlhase: Artificial Intelligence 1	202	2025-02-06	COMERCIALIS RESERVED			

Why Games are hard to solve II

- \triangleright **Example 7.1.8.** Number of reachable states in chess: 10^{40} .
- \triangleright Example 7.1.9. Number of reachable states in go: 10^{100} .
- ▷ It's even worse: Our algorithms here look at search trees (game trees), no duplicate pruning.
- ▷ Example 7.1.10.
 - \triangleright Chess without duplicate pruning: $35^{100} \simeq 10^{154}$.
 - \triangleright Go without duplicate pruning: $200^{300}\simeq 10^{690}.$

FAU

203

e

2025-02-06

How To Describe a Game State Space?

Michael Kohlhase: Artificial Intelligence 1

- ▷ Like for classical search problems, there are three possible ways to describe a game: blackbox/API description, declarative description, explicit game state space.
- ▷ **Question:** Which ones do humans use?
 - $_{\rm \triangleright}$ Explicit \approx Hand over a book with all 10^{40} moves in chess.
 - $_{\triangleright}$ Blackbox \approx Give possible chess moves on demand but don't say how they are generated.

▷ Answer: Declarative!
 With "game description language"
 [^] natural language.





7.2 Minimax Search

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22061.

53





Minimax: Outline

 \triangleright We max, we min, we max, we min ...

1. Depth first search in game tree, with Max in the root.

- 2. Apply game utility function to terminal positions.
- 3. Bottom-up for each inner node n in the search tree, compute the utility $\hat{u}(n)$ of n as follows:
 - \triangleright If it's Max's turn: Set $\hat{u}(n)$ to the maximum of the utilities of n's successor nodes.
 - \triangleright If it's Min's turn: Set $\hat{u}(n)$ to the minimum of the utilities of n's successor nodes.

209

4. Selecting a move for Max at the root: Choose one move that leads to a successor node with maximal utility.

FAU

Michael Kohlhase: Artificial Intelligence 1



The Minimax Algorithm: Pseudo-Code

```
▷ Definition 7.2.2. The minimax algorithm (often just called minimax) is given by
the following functions whose argument is a state s \in S^{\text{Max}}, in which Max is to
move.
function Minimax-Decision(s) returns an action
v := \text{Max}-\text{Value}(s)
return an action yielding value v in the previous function call
function Max-Value(s) returns a utility value
if Terminal-Test(s) then return u(s)
v := -\infty
for each a \in \text{Actions}(s) do
v := \max(v, \text{Min}-\text{Value}(\text{ChildState}(s, a)))
return v
function Min-Value(s) returns a utility value
```

2025-02-06





56

7.2. MINIMAX SEARCH



57







Minimax, Pro and Contra

▷ Minimax advantages:

> Minimax is the simplest possible (reasonable) search algorithm for games.

7.3. EVALUATION FUNCTIONS



7.3 Evaluation Functions

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22064. We now address the problem that minimax is infeasible in practice. As so often, the solution is to eschew optimal strategies and to approximate them. In this case, instead of a computed utility function, we estimate one that is easy to compute: the evaluation function.

Evaluation Eurotions for Minimax
▷ Problem: Search tree are too big to search through in minimax.
\triangleright Solution: We impose a search depth limit (also called horizon) d , and apply an evaluation function to the cut-off states, i.e. states s with $dp(s) = d$.
\triangleright Definition 7.3.1. An evaluation function f maps game states to numbers:
$\triangleright f(s)$ is an estimate of the actual value of s (as would be computed by unlimited-depth minimax for s).
\triangleright If cut-off state is terminal: Just use \hat{u} instead of f .
 ▷ Analogy to heuristic functions (cf. ??): We want f to be both (a) accurate and (b) fast.
\rhd Another analogy: (a) and (b) are in contradiction \rightsquigarrow need to trade-off accuracy against overhead.
\triangleright In typical game playing algorithms today, f is inaccurate but very fast. (usually no good methods known for computing accurate f)
Michael Kohlhase: Artificial Intelligence 1 214 2025-02-06

Example Revisited: Minimax With Depth Limit d = 2



Example Chess



Linear Evaluation Functions

- ▷ **Problem:** How to come up with evaluation functions?
- ▷ **Definition 7.3.2.** A common approach is to use a weighted linear function for f, i.e. given a sequence of features $f_i: S \to \mathbb{R}$ and a corresponding sequence of weights $w_i \in \mathbb{R}$, f is of the form $f(s):=w_1 \cdot f_1(s) + w_2 \cdot f_2(s) + \cdots + w_n \cdot f_n(s)$
- ▷ **Problem:** How to obtain these weighted linear functions?
 - \triangleright Weights w_i can be learned automatically.
- (learning agent)
- \triangleright The features f_i , however, have to be designed by human experts.
- ▷ **Note:** Very fast, very simplistic.
- \triangleright **Observation:** Can be computed incrementally: In transition $s \xrightarrow{a} s'$, adapt f(s) to f(s') by considering only those features whose values have changed.

7.3. EVALUATION FUNCTIONS

FAU Michael Kohlhase: Artificial Intelligence 1 217 2025-02-06

This assumes that the features (their contribution towards the actual value of the state) are independent. That's usually not the case (e.g. the value of a rook depends on the pawn structure).



So, How Deeply to Search? \triangleright **Goal:** In given time, search as deeply as possible. ▷ **Problem:** Very difficult to predict search running time. (need an anytime algorithm) ▷ **Solution:** Iterative deepening search. \triangleright Search with depth limit $d = 1, 2, 3, \ldots$ ▷ When time is up: return result of deepest completed search. ▷ Definition 7.3.4 (Better Solution). The quiescent search algorithm uses a dynamically adapted search depth d: It searches more deeply in unquiet positions, where value of evaluation function changes a lot in neighboring states. ▷ Example 7.3.5. In quiescent search for chess: ▷ piece exchange situations ("you take mine, I take yours") are very unquiet ightarrow Keep searching until the end of the piece exchange is reached. FAU C Michael Kohlhase: Artificial Intelligence 1 219 2025-02-06

7.4 Alpha-Beta Search

We have seen that evaluation functions can overcome the combinatorial explosion induced by minimax search. But we can do even better: certain parts of the minimax search tree can be safely ignored, since we can prove that they will only sub-optimal results. We discuss the technique of alphabeta-pruning in detail as an example of such pruning methods in search algorithms.





7.4. ALPHA-BETA SEARCH







67



Alpha-Beta Search: Pseudocode ▷ **Definition 7.4.4.** The alphabeta search algorithm is given by the following pseudocode function Alpha-Beta-Search (s) returns an action $v := Max - Value(s, -\infty, +\infty)$ return an action yielding value v in the previous function call function Max–Value(s, α, β) returns a utility value if Terminal–Test(s) then return u(s) $v := -\infty$ for each $a \in Actions(s)$ do $v := \max(v, \operatorname{Min}-\operatorname{Value}(\operatorname{ChildState}(s, a), \alpha, \beta))$ $\alpha := \max(\alpha, v)$ if $v \ge \beta$ then return v / * Here: $v \ge \beta \Leftrightarrow \alpha \ge \beta * /$ return \overline{v} function Min–Value(s, α , β) returns a utility value if Terminal–Test(s) then return u(s) $v := +\infty$ for each $a \in Actions(s)$ do $v := \min(v, \mathsf{Max} - \mathsf{Value}(\mathsf{ChildState}(s, a), \alpha, \beta))$ $\beta := \min(\beta, v)$ if $v \leq \alpha$ then return v /* Here: $v \leq \alpha \Leftrightarrow \alpha \geq \beta */$ return v $\hat{=}$ Minimax (slide 211) + α/β book-keeping and pruning. Fau Michael Kohlhase: Artificial Intelligence 1 225 2025-02-06

Note: Note that α only gets assigned a value in Max-nodes, and β only gets assigned a value in Min-nodes.

7.4. ALPHA-BETA SEARCH








7.4. ALPHA-BETA SEARCH









How Much Pruning Do We Get?

- \triangleright Choosing the best moves first yields most pruning in alphabeta search.
 - ▷ The maximizing moves for Max, the minimizing moves for Min.
- \triangleright **Observation:** Assuming game tree with branching factor *b* and depth limit *d*:
 - \triangleright Minimax would have to search b^d nodes.
 - \triangleright Best case: If we always choose the best moves first, then the search tree is reduced to $b^{\frac{d}{2}}$ nodes!
 - Practice: It is often possible to get very close to the best case by simple moveordering methods.
- \triangleright Example 7.4.5 (Chess).

e

2025-02-06

- ▷ Move ordering: Try captures first, then threats, then forward moves, then backward moves.
- \triangleright From 35^d to $35^{\frac{d}{2}}$. E.g., if we have the time to search a billion (10^9) nodes, then minimax looks ahead d = 6 moves, i.e., 3 rounds (white-black) of the game. Alpha-beta search looks ahead 6 rounds.

228

FAU Michael Kohlhase: Artificial Intelligence 1

7.5 Monte-Carlo Tree Search (MCTS)

Video Nuggets covering this section can be found at https://fau.tv/clip/id/22259 and https://fau.tv/clip/id/22262.

We will now come to the most visible game-play program in recent times: The AlphaGo system for the game of go. This has been out of reach of the state of the art (and thus for alphabeta search) until 2016. This challenge was cracked by a different technique, which we will discuss in this section.







This looks only at a fraction of the search tree, so it is crucial to have good guidance *where to go*, i.e. which part of the search tree to look at.



The sampling goes middle, left, right, right, left, middle. Then it stops and selects the highestaverage action, 60, left. After first sample, when values in initial state are being updated, we have the following "expansions" and "avg. reward fields": small number of expansions favored for exploration: visit parts of the tree rarely visited before, what is out there? avg. reward: high values favored for exploitation: focus on promising parts of the search tree.



This is the exact same search as on previous slide, but incrementally building the search tree, by always keeping the first state of the sample. The first three iterations middle, left, right, go to show the tree extension; do point out here that, like the root node, the nodes added to the tree have expansions and avg reward counters for every applicable action. Then in next iteration right, after 30 leaf node was found, an important thing is that the averages get updated *along the entire path*, i.e., not only in the root as we did before, but also in the nodes along the way. After all six iterations have been done, as before we select the action left, value 60; but we keep the part of the tree below that action, "saving relevant work already done before".

How to Guide the Search in MCTS? How to sample?: What exactly is "random"? Classical formulation: balance exploitation vs. exploration. Exploitation: Prefer moves that have high average already (interesting regions of state space) Exploration: Prefer moves that have not been tried a lot yet (don't overlook other, possibly better, options) UCT: "Upper Confidence bounds applied to Trees" [KS06].

7.5. MONTE-CARLO TREE SEARCH (MCTS)

- ▷ Inspired by Multi-Armed Bandit (as in: Casino) problems.
- ▷ Basically a formula defining the balance. Very popular (buzzword).
- ▷ Recent critics (e.g. [FD14]): Exploitation in search is very different from the Casino, as the "accumulated rewards" are fictitious (we're only thinking about the game, not actually playing and winning/losing all the time).

Michael Kohlhase: Artificial Intelligence 1 233 2025-02-06





 \triangleright RL policy network p_{ρ} : Reinforcement learning, self-play ("learn to win").

2025-02-06

▷ Value network v_{θ} : Use self-play games with p_{ρ} as training data for game-position evaluation v_{θ} ("predict which player will win in this state").

235

Comments on the Figure:

Michael Kohlhase: Artificial Intelligence 1

FAU

- a A fast rollout policy p_{π} and supervised learning (SL) policy network p_{σ} are trained to predict human expert moves in a data set of positions. A reinforcement learning (RL) policy network p_{ρ} is initialized to the SL policy network, and is then improved by policy gradient learning to maximize the outcome (that is, winning more games) against previous versions of the policy network. A new data set is generated by playing games of self-play with the RL policy network. Finally, a value network v_{θ} is trained by regression to predict the expected outcome (that is, whether the current player wins) in positions from the self-play data set.
- b Schematic representation of the neural network architecture used in AlphaGo. The policy network takes a representation of the board position s as its input, passes it through many convolutional layers with parameters σ (SL policy network) or ρ (RL policy network), and outputs a probability distribution $p_{\sigma}(a|s)$ or $p_{\rho}(a|s)$ over legal moves a, represented by a probability map over the board. The value network similarly uses many convolutional layers with parameters θ , but outputs a scalar value $v_{\theta}(s')$ that predicts the expected outcome in position s'.



Comments on the Figure:

a Each simulation traverses the tree by selecting the edge with maximum action value Q, plus a bonus u(P) that depends on a stored prior probability P for that edge.

- b The leaf node may be expanded; the new node is processed once by the policy network p_{σ} and the output probabilities are stored as prior probabilities P for each action.
- c At the end of a simulation, the leaf node is evaluated in two ways:
 - using the value network v_{θ} ,
 - and by running a rollout to the end of the game

with the fast rollout policy $p \pi$, then computing the winner with function r.

d Action values Q are updated to track the mean value of all evaluations $r(\cdot)$ and $v_{\theta}(\cdot)$ in the subtree below that action.

AlphaGo, Conclusion?: This is definitely a great achievement!

- "Search + neural networks" looks like a great formula for general problem solving.
- expect to see lots of research on this in the coming decade(s).
- The AlphaGo design is quite intricate (architecture, learning workflow, training data design, neural network architectures, ...).
- How much of this is reusable in/generalizes to other problems?
- Still lots of human expertise in here. Not as much, like in chess, about the game itself. But rather, in the design of the neural networks + learning architecture.

7.6 State of the Art

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22250.

State of the Art > Some well-known board games: ⊳ Chess: Up next. > Othello (Reversi): In 1997, "Logistello" beat the human world champion. Best computer players now are clearly better than best human players. ▷ Checkers (Dame): Since 1994, "Chinook" is the offical world champion. In 2007, it was shown to be unbeatable: Checkers is solved. (We know the exact value of, and optimal strategy for, the initial state.) ▷ Go: In 2016, AlphaGo beat the Grandmaster Lee Sedol, cracking the "holy grail" of board games. In 2017, "AlphaZero" - a variant of AlphaGo with zero prior knowledge beat all reigning champion systems in all board games (including AlphaGo) 100/0 after 24h of self-play. > Intuition: Board Games are considered a "solved problem" from the AI perspective. Fau C 2025-02-06 Michael Kohlhase: Artificial Intelligence 1 237

Computer Chess: "Deep Blue" beat Garry Kasparov in 1997



7.7 Conclusion

Summary

- ▷ Games (2-player turn-taking zero-sum discrete and finite games) can be understood as a simple extension of classical search problems.
- ▷ Each player tries to reach a terminal state with the best possible utility (maximal vs. minimal).

7.7. CONCLUSION

- \triangleright Minimax searches the game depth-first, max'ing and min'ing at the respective turns of each player. It yields perfect play, but takes time $\mathcal{O}(b^d)$ where b is the branching factor and d the search depth.
- ▷ Except in trivial games (Tic-Tac-Toe), minimax needs a depth limit and apply an evaluation function to estimate the value of the cut-off states.
- ▷ Alpha-beta search remembers the best values achieved for each player elsewhere in the tree already, and prunes out sub-trees that won't be reached in the game.
- Monte Carlo tree search (MCTS) samples game branches, and averages the findings. AlphaGo controls this using neural networks: evaluation function ("value network"), and action filter ("policy network").

Michael Kohlhase: Artificial Intelligence 1 241 2025-02-06

Suggested Reading:

- Chapter 5: Adversarial Search, Sections 5.1 5.4 [RN09].
 - Section 5.1 corresponds to my "Introduction", Section 5.2 corresponds to my "Minimax Search", Section 5.3 corresponds to my "Alpha-Beta Search". I have tried to add some additional clarifying illustrations. RN gives many complementary explanations, nice as additional background reading.
 - Section 5.4 corresponds to my "Evaluation Functions", but discusses additional aspects relating to narrowing the search and look-up from opening/termination databases. Nice as additional background reading.
 - I suppose a discussion of MCTS and AlphaGo will be added to the next edition ...

Chapter 8

Constraint Satisfaction Problems

In the last chapters we have studied methods for "general problem", i.e. such that are applicable to all problems that are expressible in terms of states and "actions". It is crucial to realize that these states were atomic, which makes the algorithms employed (search algorithms) relatively simple and generic, but does not let them exploit the any knowledge we might have about the *internal* structure of states.

In this chapter, we will look into algorithms that do just that by progressing to factored states representations. We will see that this allows for algorithms that are many orders of magnitude more efficient than search algorithms.

To give an intuition for factored states representations we, we present some motivational examples in ?? and go into detail of the Waltz algorithm, which gave rise to the main ideas of constraint satisfaction algorithms in ??. ?? and ?? define constraint satisfaction problems formally and use that to develop a class of backtracking/search based algorithms. The main contribution of the factored states representations is that we can formulate advanced search heuristics that guide search based on the structure of the states.

Constraint Satisfaction Problems: Motivation 8.1

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22251.

A (Constraint Satisfaction) Problem > Example 8.1.1 (Tournament Schedule). Who's going to play against who, when
and where? and where?



Constraint Satisfaction Problems (CSPs)

- ▷ Standard search problem: state is a "black box" any old data structure that supports goal test, eval, successor state, ...
- \triangleright **Definition 8.1.2.** A constraint satisfaction problem (CSP) is a triple $\langle V, D, C \rangle$ where
 - 1. V is a finite set V of variables,
 - 2. an V-indexed family $(D_v)_{v \in V}$ of domains, and
 - 3. for some subsets $\{v_1, \ldots, v_k\} \subseteq V$ a constraint $C_{\{v_1, \ldots, v_k\}} \subset D_{v_1} \times \ldots \times D_{v_k}$.
 - A variable assignment $\varphi \in (v \in V) \rightarrow D_v$ is a solution for C, iff $\langle \varphi(v_1), \ldots, \varphi(v_k) \rangle \in C_{\{v_1, \ldots, v_k\}}$ for all $\{v_1, \ldots, v_k\} \subseteq V$.
- \triangleright **Definition 8.1.3.** A CSP γ is called satisfiable, iff it has a solution: a total variable assignment φ that satisfies all constraints.
- ▷ Definition 8.1.4. The process of finding solutions to CSPs is called constraint solving.
- \triangleright Remark 8.1.5. We are using factored representation for world states now!
- ▷ Allows useful *general-purpose* algorithms with more power than standard tree search algorithm.

FAU

Michael Kohlhase: Artificial Intelligence 1

243

2025-02-06

Another Constraint Satisfaction Problem



Example 8.1.6 (SuDoKu). Fill the cells with row/column/block-unique digits



Bundesliga Constraints

 \triangleright Variables: $v_{Avs.B}$ where A and B are teams, with domains $\{1, \ldots, 34\}$: For each match, the index of the weekend where it is scheduled.



How to Solve the Bundesliga Constraints?

- ightarrow 306 nested for-loops (for each of the 306 matches), each ranging from 1 to 306. Within the innermost loop, test whether the current values are (a) a permutation and, if so, (b) a legal Bundesliga schedule.
 - Estimated running time: End of this universe, and the next couple billion ones after it ...
- \triangleright Directly enumerate all permutations of the numbers $1, \ldots, 306$, test for each whether it's a legal Bundesliga schedule.
 - Estimated running time: Maybe only the time span of a few thousand universes.

▷ View this as variables/constraints and use backtracking (this chapter)

▷ **Executed running time**: About 1 minute.

More Constraint Satisfaction Problems

▷ How do they actually do it?: Modern computers and CSP methods: fractions of a second. 19th (20th/21st?) century: Combinatorics and manual work.

▷ Try it yourself: with an off-the shelf CSP solver, e.g. Minion [Min]

Michael Kohlhase: Artificial Intelligence 1

247

2025-02-06

8.1. CONSTRAINT SATISFACTION PROBLEMS: MOTIVATION



- 1. U.S. Major League Baseball, 30 teams, each 162 games. There's one crucial additional difficulty, in comparison to Bundesliga. Which one? Travel is a major issue here!! Hence "Traveling Tournament Problem" in reference to the TSP.
- 2. This particular scheduling problem is called "car sequencing", how to most efficiently get cars through the available machines when making the final customer configuration (non-standard/flexible/custom extras).
- 3. Another common form of scheduling ...
- 4. The problem of assigning radio frequencies so that all can operate together without noticeable interference. Variable domains are available frequencies, constraints take form of $|x y| > \delta_{xy}$, where delta depends on the position of x and y as well as the physical environment.

Our Agenda for This Topic								
Our treatment of the topic "Constraint Satisfaction Problems" consists of Chap- ters 7 and 8. in [RN03]								
> This Chapter: Basic definitions and concepts; naïve backtracking search.								
Sets up the framework. Backtracking underlies many successful algorithms for solving constraint satisfaction problems (and, naturally, we start with the sim- plest version thereof).								
> Next Chapter: Constraint propagation and decomposition methods.								
Constraint propagation reduces the search space of backtracking. Decomposi- tion methods break the problem into smaller pieces. Both are crucial for efficiency in practice.								
Michael Kohlhase: Artificial Intelligence 1 249 2025-02-06								



8.2 The Waltz Algorithm

We will now have a detailed look at the problem (and innovative solution) that started the field of constraint satisfaction problems.

Background:

Adolfo Guzman worked on an algorithm to count the number of simple objects (like children's blocks) in a line drawing. David Huffman formalized the problem and limited it to objects in general position, such that the vertices are always adjacent to three faces and each vertex is formed from three planes at right angles (trihedral). Furthermore, the drawings could only have three kinds of lines: object boundary, concave, and convex. Huffman enumerated all possible configurations of lines around a vertex. This problem was too narrow for real-world situations, so Waltz generalized it to include cracks, shadows, non-trihedral vertices and light. This resulted in over 50 different line labels and thousands of different junctions. [ILD]



8.2. THE WALTZ ALGORITHM



18 Legal Kinds of Junctions

▷ **Observation 8.2.2.** There are only 18 "legal" kinds of junctions:



91

CHAPTER 8. CONSTRAINT SATISFACTION PROBLEMS









8.3 CSP: Towards a Formal Definition

We will now work our way towards a definition of CSPs that is formal enough so that we can define the concept of a solution. This gives use the necessary grounding to talk about algorithms later. A Video Nugget covering this section can be found at https://fau.tv/clip/id/22277.

```
Types of CSPs
▷ Definition 8.3.1. We call a CSP discrete, iff all of the variables have countable domains; we have two kinds:

▷ finite domains
○ e.g., Boolean CSPs
○ solvability = Boolean satisfiability ~> NP complete)
▷ infinite domains (e.g. integers, strings, etc.)
▷ e.g., job scheduling, variables are start/end days for each job
▷ need a "constraint language", e.g., StartJob<sub>1</sub> + 5 ≤ StartJob<sub>3</sub>
▷ linear constraints decidable, nonlinear ones undecidable

▷ Definition 8.3.2. We call a CSP continuous, iff one domain is uncountable.
▷ Example 8.3.3. Start/end times for Hubble Telescope observations form a continuous CSP.
▷ Theorem 8.3.4. Linear constraints solvable in poly time by linear programming methods.
```

CHAPTER 8. CONSTRAINT SATISFACTION PROBLEMS

▷ Theorem 8.3.5. There cannot be optimal algorithms for nonlinear constraint systems.

Michael Kohlhase: Artificial Intelligence 1

257

2025-02-06

Types of Constraints							
▷ We classify the constraints by the number of variables they involve.							
\triangleright Definition 8.3.6. Unary constraints involve a single variable, e.g., $SA \neq green$.							
\triangleright Definition 8.3.7. Binary constraints involve pairs of variables, e.g., $SA \neq WA$.							
\triangleright Definition 8.3.8. Higher-order constraints involve $n = 3$ or more variables, e.g., cryptarithmetic column constraints.							
The number n of variables is called the order of the constraint.							
Definition 8.3.9. Preferences (soft constraint) (e.g., red is better than green) are often representable by a cost for each variable assignment ~ constrained opti- mization problems.							
Michael Kohlhase: Artificial Intelligence 1 258 2025-02-06							
Non-Binary Constraints, e.g. Send More Money							
> Example 8.3.10 (Send More Money). A student writes home:							
$\begin{array}{cccccccccccccccccccccccccccccccccccc$							

 \triangleright Variables: S, E, N, D, M, O, R, Y, each with domain $\{0, \ldots, 9\}$.

⊳ Constraints:

- 1. all variables should have different values: $S \neq E, S \neq N, \ldots$
- 2. first digits are non-zero: $S \neq 0$, $M \neq 0$.
- 3. the addition scheme should work out: i.e. $1000 \cdot S + 100 \cdot E + 10 \cdot N + D + 1000 \cdot M + 100 \cdot O + 10 \cdot R + E = 10000 \cdot M + 1000 \cdot 0 + 100 \cdot N + 10 \cdot E + Y.$

BTW: The solution is $S \mapsto 9, E \mapsto 5, N \mapsto 6, D \mapsto 7, M \mapsto 1, O \mapsto 0, R \mapsto 8, Y \mapsto 2 \sim$ parents send $10652 \in$

▷ **Definition 8.3.11.** Problems like the one in ?? are called crypto-arithmetic puzzles.

FAU

Michael Kohlhase: Artificial Intelligence 1

259

2025-02-06

Encoding Higher-Order Constraints as Binary ones

94

8.3. CSP: TOWARDS A FORMAL DEFINITION

▷ **Problem:** The last constraint is of order 8. (n = 8 variables involved)▷ **Observation 8.3.12.** We can write the addition scheme constraint column wise using auxiliary variables, i.e. variables that do not "occur" in the original problem. $D + E = Y + 10 \cdot X_1$ $\begin{array}{rcl} S & E & N & D \\ X_1 + N + R & = & E + 10 \cdot X_2 \\ X_2 + E + O & = & N + 10 \cdot X_3 \end{array} \qquad \begin{array}{rcl} S & E & N & D \\ + & M & O & R & E \\ \hline M & O & N & E & Y \end{array}$ $X_3 + S + M = O + 10 \cdot M$ These constraints are of order ≤ 5 . \triangleright General Recipe: For $n \geq 3$, encode $C(v_1, \ldots, v_{n-1}, v_n)$ as $C(p_1(x), \dots, p_{n-1}(x), v_n) \wedge v_1 = p_1(x) \wedge \dots \wedge v_{n-1} = p_{n-1}(x)$ ▷ **Problem:** The problem structure gets hidden. (search algorithms can get confused) FAU e Michael Kohlhase: Artificial Intelligence 1 260 2025-02-06



Real-world CSPs

▷ Example 8.3.16 (Assignment problems). e.g., who teaches what class



8.4 Constraint Networks: Formalizing Binary CSPs

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22279.

Constraint Networks (Formalizing binary CSPs) \triangleright **Definition 8.4.1.** A constraint network is a triple $\gamma := \langle V, D, C \rangle$, where $\triangleright V$ is a finite set of variables, $\triangleright D := \{D_v \mid v \in V\}$ the set of their domains, and $\triangleright C := \{C_{uv} \subseteq D_u \times D_v \mid u, v \in V \text{ and } u \neq v\} \text{ is a set of constraints with } C_{uv} =$ C_{vu}^{-1} . We call the undirected graph $\langle V, \{(u,v) \in V^2 \mid C_{uv} \neq D_u \times D_v \} \rangle$, the constraint graph of γ . ▷ We will talk of CSPs and mean constraint networks. ▷ **Remarks:** The mathematical formulation gives us a lot of leverage: $\triangleright C_{uv} \subseteq D_u \times D_v \cong$ possible assignments to variables u and v> Relations are the most general formalization, generally we use symbolic formulations, e.g. "u = v" for the relation $C_{uv} = \{(a,b) \mid a = b\}$ or " $u \neq v$ ". \triangleright We can express unary constraints C_u by restricting the domain of $v: D_v := C_v$. FAU Michael Kohlhase: Artificial Intelligence 1 263 2025-02-06

Example: SuDoKu as a Constraint Network

Example 8.4.2 (Formalize SuDoKu). We use the added formality to encode SuDoKu as a constraint network, not just as a CSP as ??.

2	5			3		9		1
	1				4			
4		7				2		8
		5	2					
				9	8	1		
	4				3			
			3	6			7	2
	7							3
9		3				6		4

- \triangleright Variables: $V = \{v_{ij} \mid 1 \le i, j \le 9\}$: $v_{ij} = \text{cell in row } i \text{ column } j$.
- \triangleright Domains For all $v \in V$: $D_v = D = \{1, \dots, 9\}.$
- \triangleright Unary constraint: $C_{v_{ij}} = \{d\}$ if cell i, j is pre-filled with d.

Note that the ideas are still the same as ??, but in constraint networks we have a language to formulate things precisely.

264

Fau

Constraint Networks (Solutions)

 \triangleright Let $\gamma := \langle V, D, C \rangle$ be a constraint network.

Michael Kohlhase: Artificial Intelligence 1

- ▷ **Definition 8.4.3.** We call a partial function $a: V \rightarrow \bigcup_{u \in V} D_u$ a variable assignment if $a(u) \in D_u$ for all $u \in \text{dom}(a)$.
- ▷ **Definition 8.4.4.** Let $C := \langle V, D, C \rangle$ be a constraint network and $a: V \rightarrow \bigcup_{v \in V} D_v$ a variable assignment. We say that *a* satisfies (otherwise violates) a constraint C_{uv} , iff $u, v \in \operatorname{dom}(a)$ and $(a(u), a(v)) \in C_{uv}$. *a* is called consistent in *C*, iff it satisfies all constraints in *C*. A value $w \in D_u$ is legal for a variable *u* in *C*, iff $\{(u,w)\}$ is a consistent assignment in *C*. A variable with illegal value under *a* is called conflicted.
- \triangleright Example 8.4.5. The empty assignment ϵ is (trivially) consistent in any constraint network.
- ▷ **Definition 8.4.6.** Let f and g be variable assignments, then we say that f extends (or is an extension of) g, iff $dom(g) \subset dom(f)$ and $f|_{dom(g)} = g$.
- \triangleright **Definition 8.4.7.** We call a consistent (total) assignment a solution for γ and γ itself solvable or satisfiable.

FAU

Michael Kohlhase: Artificial Intelligence 1

265

CC STATE BOARD AND A STREET AND

2025-02-06

2025-02-06

How it all fits together

▷ Lemma 8.4.8. Higher-order constraints can be transformed into equi-satisfiable

2025-02-06

binary constraints using auxiliary variables.

- ▷ Corollary 8.4.9. Any CSP can be represented by a constraint network.
- \triangleright In other words The notion of a constraint network is a refinement of a CSP.
- \triangleright So we will stick to constraint networks in this course.
- Observation 8.4.10. We can view a constraint network as a search problem, if we take the states as the variable assignments, the actions as assignment extensions, and the goal states as consistent assignments.
- ▷ Idea: We will explore that idea for algorithms that solve constraint networks.

Michael Kohlhase: Artificial Intelligence 1 266

8.5 CSP as Search

We now follow up on ?? to use search algorithms for solving constraint networks.

The key point of this section is that the factored states representations realized by constraint networks allow the formulation of very powerful heuristics. A Video Nugget covering this section can be found at https://fau.tv/clip/id/22319.

Standard search formulation (incremental) ▷ **Idea:** Every constraint network induces a single state problem. \triangleright Definition 8.5.1 (Let's do the math). Given a constraint network $\gamma := \langle V, D, C \rangle$, then $\Pi_{\gamma} := \langle S_{\gamma}, \mathcal{A}_{\gamma}, \mathcal{T}_{\gamma}, \mathcal{G}_{\gamma} \rangle$ is called the search problem induced by γ , iff \triangleright State S_{γ} are variable assignments \triangleright Action \mathcal{A}_{γ} : extend $\varphi \in \mathcal{S}_{\gamma}$ by a pair $x \mapsto v$ not conflicted with φ . \triangleright Transition model $\mathcal{T}_{\gamma}(a,\varphi) = \varphi, x \mapsto v$ (extended assignment) \triangleright Initial state \mathcal{I}_{γ} : the empty assignment ϵ . \triangleright Goal states \mathcal{G}_{γ} : the total, consistent assignments \triangleright What has just happened?: We interpret a constraint network γ as a search problem Π_{γ} . A solution to Π_{γ} induces a solution to γ . ▷ **Idea:** We have algorithms for that: e.g. tree search. \triangleright **Remark:** This is the same for all CSPs! \odot \sim fail if no consistent assignments exist (not fixable!) FAU C Michael Kohlhase: Artificial Intelligence 1 267 2025-02-06

Standard search formulation (incremental)

 \triangleright **Example 8.5.2.** A search tree for $\Pi_{Australia}$:

8.5. CSP AS SEARCH



Backtracking Search

▷ Assignments for different variables are independent!
 ▷ e.g. first WA = red then NT = green vs. first NT = green then WA = red
 ▷ ~ we only need to consider assignments to a single variable at each node
 ▷ ~ b = d and there are dⁿ leaves.
 ▷ Definition 8.5.3. Depth first search for CSPs with single-variable assignment extensions actions is called backtracking search.
 ▷ Backtracking search is the basic uninformed algorithm for CSPs.
 ▷ It can solve the *n*-queens problem for ≈ n, 25.

Backtracking Search (Implementation)
▷ Definition 8.5.4. The generic backtracking search algorithm:
<pre>procedure Backtracking—Search(csp) returns solution/failure</pre>
return Recursive—Backtracking (Ø. csp)
····· ···· ··· ··· ··· ··· ··· ··· ···
procedure Recursive—Backtracking (assignment) returns soln/failure
if assignment is complete then return assignment
var := Select-Unassigned-Variable(Variables[csp], assignment, csp)
foreach value in Order–Domain–Values(var, assignment, csp) do
if value is consistent with assignment given Constraints[csp] then
add {var = value} to assignment
result := Recursive-Backtracking(assignment,csp)

CHAPTER 8. CONSTRAINT SATISFACTION PROBLEMS









Where in ?? does the most constraining variable play a role in the choice? SA (only possible choice), NT (all choices possible except WA, V, T). Where in the illustration does most constrained variable play a role in the choice? NT (all choices possible except T), Q (only Q and WA

possible).



8.6 Conclusion & Preview



Suggested Reading: p

• Chapter 6: Constraint Satisfaction Problems, Sections 6.1 and 6.3, in [RN09].

- Compared to our treatment of the topic "Constraint Satisfaction Problems" (?? and ??), RN covers much more material, but less formally and in much less detail (in particular, my slides contain many additional in-depth examples). Nice background/additional reading, can't replace the lectures.
- Section 6.1: Similar to our "Introduction" and "Constraint Networks", less/different examples, much less detail, more discussion of extensions/variations.
- Section 6.3: Similar to my "Naïve Backtracking" and "Variable- and Value Ordering", with less examples and details; contains part of what we cover in **??** (RN does inference first, then backtracking). Additional discussion of *backjumping*.

Chapter 9

Constraint Propagation

In this chapter we discuss another idea that is central to symbolic AI as a whole. The first component is that with the factored states representations, we need to use a representation language for (sets of) states. The second component is that instead of state-level search, we can graduate to representation-level search (inference), which can be much more efficient that state level search as the respective representation language actions correspond to groups of state-level actions.

9.1 Introduction

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22321.



 \triangleright Example 9.1.3. Constraint network γ :

Illustration: Decomposition

CHAPTER 9. CONSTRAINT PROPAGATION





9.2 Constraint Propagation/Inference

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22326.





108




- ▷ **Simple:** Constraint propagation as a pre-process:
 - ▷ When: Just once before search starts.
 - Effect: Little running time overhead, little pruning power. (not considered here)
- ▷ **More Advanced:** Constraint propagation during search:
 - ▷ When: At every recursive call of backtracking.
 - **Effect**: Strong pruning power, may have large running time overhead.
- ▷ **Search vs. Inference:** The more complex the inference, the *smaller* the number of search nodes, but the *larger* the running time needed at each node.
- \triangleright Idea: Encode variable assignments as unary constraints (i.e., for a(v) = d, set the unary constraint $D_v = \{d\}$), so that inference reasons about the network restricted to the commitments already made in the search.
- FAU

Michael Kohlhase: Artificial Intelligence 1

284



9.3 Forward Checking

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22326.



9.3. FORWARD CHECKING



Note: It's a bit strange that we start with d' here; this is to make link to arc consistency – coming up next – as obvious as possible (same notations u, and d vs. v and d').

Forward Checking: Discussion

- \triangleright **Definition 9.3.4.** An inference procedure is called sound, iff for any input γ the output γ' have the same solutions.
- ▷ Lemma 9.3.5. Forward checking is sound.

Proof sketch: Recall here that the assignment a is represented as unary constraints inside γ .

- \triangleright Corollary 9.3.6. γ and γ' are equivalent.
- \triangleright Incremental computation: Instead of the first for-loop in ??, use only the inner one every time a new assignment a(v) = d' is added.
- ▷ Practical Properties:
 - ▷ Cheap but useful inference method.
 - ▷ Rarely a good idea to not use forward checking (or a stronger inference method subsuming it).
- ▷ **Up next:** A stronger inference method (subsuming forward checking).

 $\blacktriangleright \text{ Definition 9.3.7. Let } p \text{ and } q \text{ be inference procedures, then } p \text{ subsumes } q \text{, if } p(\gamma) \sqsubseteq q(\gamma) \text{ for any input } \gamma.$ $\blacksquare Michael Kohlhase: Artificial Intelligence 1 287 2025-02-06$

9.4 Arc Consistency

Video Nuggets covering this section can be found at https://fau.tv/clip/id/22350 and https://fau.tv/clip/id/22351.



Arc Consistency: Definition

- \triangleright Definition 9.4.2 (Arc Consistency). Let $\gamma := \langle V, D, C \rangle$ be a constraint network.
 - 1. A variable $u \in V$ is arc consistent relative to another variable $v \in V$ if either $C_{uv} \notin C$, or for every value $d \in D_u$ there exists a value $d' \in D_v$ such that $(d,d') \in C_{uv}$.
 - 2. The constraint network γ is arc consistent if every variable $u \in V$ is arc consistent relative to every other variable $v \in V$.

The concept of arc consistency concerns both levels.

- \triangleright **Intuition:** Arc consistency $\hat{=}$ for every domain value and constraint, at least one value on the other side of the constraint "works".
- \triangleright **Note** the asymmetry between u and v: arc consistency is directed.

FAU Michael Kohlhase: Artificial Intelligence 1

ntelligence 1

289



Arc Consistency: Example

- \triangleright Definition 9.4.5 (Arc Consistency). Let $\gamma := \langle V, D, C \rangle$ be a constraint network.
 - 1. A variable $u \in V$ is arc consistent relative to another variable $v \in V$ if either $C_{uv} \notin C$, or for every value $d \in D_u$ there exists a value $d' \in D_v$ such that $(d,d') \in C_{uv}$.
 - 2. The constraint network γ is arc consistent if every variable $u \in V$ is arc consistent relative to every other variable $v \in V$.

The concept of arc consistency concerns both levels.

 \triangleright Example 9.4.6.



Enforcing Arc Consistency for One Pair of Variables

 \triangleright **Definition 9.4.9 (Revise).** Revise is an algorithm enforcing arc consistency of u relative to v

 $\begin{array}{l} \text{function } \operatorname{Revise}(\gamma, u, v) \text{ returns modified } \gamma \\ \text{for each } d \in D_u \text{ do} \\ \text{if there is no } d' \in D_v \text{ with } (d, d') \in C_{uv} \text{ then } D_u := D_u \backslash \{d\} \\ \text{return } \gamma \end{array}$

- ▷ Lemma 9.4.10. If d is maximal domain size in γ and the test " $(d,d') \in C_{uv}$?" has time complexity $\mathcal{O}(1)$, then the running time of $\operatorname{Revise}(\gamma, u, v)$ is $\mathcal{O}(d^2)$.
- \triangleright Example 9.4.11. Revise (γ, v_3, v_2)



- \triangleright Lemma 9.4.13. If γ has n variables, m constraints, and maximal domain size d, then the time complexity of $AC1(\gamma)$ is $\mathcal{O}(md^2nd)$.
- \triangleright Proof sketch: $\mathcal{O}(md^2)$ for each inner loop, fixed point reached at the latest once all *nd* variable values have been removed.
- ▷ **Problem:** There are redundant computations.
- ▷ **Question:** Do you see what these redundant computations are?
- \triangleright **Redundant computations:** u and v are revised even if theirdomains haven't changed since the last time.

 \triangleright Better algorithm avoiding this: AC 3

FAU

(coming up)



AC-3: Example

 \triangleright Example 9.4.15. $y \operatorname{div} x = 0$: $y \operatorname{modulo} x$ is 0, i.e., y is divisible by x

9.4. ARC CONSISTENCY







AC-3: Runtime

- ▷ Theorem 9.4.16 (Runtime of AC-3). Let $\gamma := \langle V, D, C \rangle$ be a constraint network with m constraints, and maximal domain size d. Then AC 3(γ) runs in time $\mathcal{O}(md^3)$.
- \triangleright *Proof:* by counting how often Revise is called.



9.5 Decomposition: Constraint Graphs, and Three Simple Cases

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22353.



Problem Structure



"Decomposition" 1.0: Disconnected Constraint Graphs
> Theorem 9.5.2 (Disconnected Constraint Graphs). Let γ := (V, D, C) be a constraint network. Let a_i be a solution to each connected component γ_i of the constraint graph of γ. Then a := U_ia_i is a solution to γ.
> Proof:

a satisfies all C_{uv} where u and v are inside the same connected component.
The latter is the case for all C_{uv}.
If two parts of γ are not connected, then they are independent.

> Example 9.5.3. Color Tasmania separately in Australia

• Example 9.5.4 (Doing the Numbers).
> γ with n = 40 variables, each domain size k = 2. Four separate connected components each of size 10.

▷ Reduction of worst-case when using decomposition:

 \triangleright No decomposition: 2^{40} . With: $4 \cdot 2^{10}$. Gain: $2^{28} \approx 280.000.000$.

Definition 9.5.5. The process of decomposing a constraint network into components is called decomposition. There are various decomposition algorithms.

Michael Kohlhase: Artificial Intelligence 1

Fau

300

SCOLUMN STREET





Nearly tree-structured CSPs

- ▷ Definition 9.5.8. Conditioning: instantiate a variable, prune its neighbors' domains.
- ⊳ Example 9.5.9.



304

2025-02-06

Michael Kohlhase: Artificial Intelligence 1







9.6 Cutset Conditioning

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22354.





9.7 Constraint Propagation with Local Search

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22355.



Performance of min-conflicts

- \triangleright Given a random initial state, can solve *n*-queens in almost constant time for arbitrary *n* with high probability (e.g., n = 10,000,000)
- \triangleright The same appears to be true for any randomly-generated CSP *except* in a narrow range of the ratio

 $R = \frac{\text{number of constraints}}{\text{number of variables}}$



9.8 Conclusion & Summary

A Video Nugget covering this section can be found at https://fau.tv/clip/id/22356.



Topics We Didn't Cover Here

- \triangleright Path consistency, *k*-consistency: Generalizes arc consistency to size *k* subsets of variables. Path consistency $\hat{=}$ 3-consistency.
- ▷ Tree decomposition: Instead of instantiating variables until the leaf nodes are trees, distribute the variables and constraints over sub-CSPs whose connections form a tree.
- ▷ Backjumping: Like backtracking search, but with ability to back up across several

9.8. CONCLUSION & SUMMARY

levels (to a previous variable assignment identified to be responsible for failure).

- No-Good Learning: Inferring additional constraints based on information gathered during backtracking search.
- ▷ Local search: In space of total (but not necessarily consistent) assignments. (E.g., 8 queens in ??)
- \triangleright Tractable CSP: Classes of CSPs that can be solved in P.

Michael Kohlhase: Artificial Intelligence 1

- Global Constraints: Constraints over many/all variables, with associated specialized inference methods.
- Constraint Optimization Problems (COP): Utility function over solutions, need an optimal one.

313

Suggested Reading:

EAU

- Chapter 6: Constraint Satisfaction Problems in [RN09], in particular Sections 6.2, 6.3.2, and 6.5.
 - Compared to our treatment of the topic "constraint satisfaction problems" (?? and ??), RN covers much more material, but less formally and in much less detail (in particular, our slides contain many additional in-depth examples). Nice background/additional reading, can't replace the lectures.
 - Section 6.3.2: Somewhat comparable to our "inference" (except that equivalence and tightness are not made explicit in RN) together with "forward checking".
 - Section 6.2: Similar to our "arc consistency", less/different examples, much less detail, additional discussion of path consistency and global constraints.
 - Section 6.5: Similar to our "decomposition" and "cutset conditioning", less/different examples, much less detail, additional discussion of tree decomposition.

CHAPTER 9. CONSTRAINT PROPAGATION

Bibliography

- [FD14] Zohar Feldman and Carmel Domshlak. "Simple Regret Optimization in Online Planning for Markov Decision Processes". In: Journal of Artificial Intelligence Research 51 (2014), pp. 165–205.
- [ILD] 7. Constraints: Interpreting Line Drawings. URL: https://www.youtube.com/watch? v=l-tzjenXrvI&t=2037s (visited on 11/19/2019).
- [KS06] Levente Kocsis and Csaba Szepesvári. "Bandit Based Monte-Carlo Planning". In: Proceedings of the 17th European Conference on Machine Learning (ECML 2006). Ed. by Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou. Vol. 4212. LNCS. Springer-Verlag, 2006, pp. 282–293.
- [Met+53] N. Metropolis et al. "Equations of state calculations by fast computing machines". In: Journal of Chemical Physics 21 (1953), pp. 1087–1091.
- [Min] *Minion Constraint Modelling*. System Web page at http://constraintmodelling. org/minion/. URL: http://constraintmodelling.org/minion/.
- [Pól73] George Pólya. How to Solve it. A New Aspect of Mathematical Method. Princeton University Press, 1973.
- [RN03] Stuart J. Russell and Peter Norvig. Artificial Intelligence: A Modern Approach. 2nd ed. Pearso n Education, 2003. ISBN: 0137903952.
- [RN09] Stuart Russell and Peter Norvig. Artificial Intelligence: A Modern Approach. 3rd. Prentice Hall Press, 2009. ISBN: 0136042597, 9780136042594.
- [Sil+16] David Silver et al. "Mastering the Game of Go with Deep Neural Networks and Tree Search". In: Nature 529 (2016), pp. 484-503. URL: http://www.nature.com/nature/ journal/v529/n7587/full/nature16961.html.
- [Wal75] David Waltz. "Understanding Line Drawings of Scenes with Shadows". In: The Psychology of Computer Vision. Ed. by P. H. Winston. McGraw-Hill, 1975, pp. 1–19.

BIBLIOGRAPHY