



How can we color
With no adjacent colors
This Australia map?

Project 2: Constraint Satisfaction Problem

Introduction

A constraint satisfaction problem (CSP) is a problem specified such that a solution is an assignment of values to variables that is valid given constraints on the assignment and the variables' domains. CSPs are very powerful because a single unchanging set of algorithms can be used to solve any problem specified as a CSP, and many problems can be intuitively specified as CSPs. In this project you will be implementing a CSP solver and improving it with heuristic and inference algorithms.

For this project, you are provided with an autograder as well as many test cases that specify CSP problems. These test cases are specified within the `csps` directory. Functions are provided in `Testing` to parse these files into the objects your algorithms will work with. The files follow a simple format you can understand by inspection, so if your code does not work looking at the problems themselves and manually checking your logic is the best debugging strategy.

Files you will edit

BinaryCSP.py Your entire CSP implementation will be within this file

Files you will not edit

Testing.py Helper functions for invoking CSP problems
autograder.pyc A custom autograder to check your code with

Evaluation: Your code will be autograded for technical correctness, using the same autograder and test cases you are provided with. Please *do not* change the names of any provided functions or classes within the code, or you will wreak havoc on the autograder. *However*, the correctness of your

implementation -- not the autograder's judgements -- will be the final judge of your score. If your score is not perfect, we will review your code individually to ensure you receive a fair amount of partial credit. Likewise, any 'shortcuts' that make your code pass the autograder will not pass.

Academic Dishonesty: We will be checking your code against other submissions in the class for logical redundancy. If you copy someone else's code and submit it with minor changes, we will know. These cheat detectors are quite hard to fool, so please don't try. We trust you all to submit your own work only; *please* don't let us down. Likewise, *do not* attempt to write your code specifically to pass the autograder's tests. Either copying or trying to cheat the autograder will be considered violations of the student honor code.

Getting Help: You are not alone! If you find yourself stuck on something, contact the course staff for help. Office hours and Piazza are there for your support; please use them. If you can't make our office hours, let us know and we will schedule more. We want these projects to be rewarding and instructional, not frustrating and demoralizing. But, we don't know when or how to help unless you ask.

Before You Begin: A Note on Structure

All of the necessary structure for assignments and csp problems is provided for you in `BinaryCSP.py`. While you do not need to implement these structures, it is important to understand how they work.

Almost every function you will be implementing will take in a *ConstraintSatisfactionProblem* and an *Assignment*. The *ConstraintSatisfactionProblem* object serves only as a representation of the problem, and it not intended to be changed. It holds three things: a dictionary from variables to their domains (*varDomains*), a list of binary constraints(*binaryConstraints*), and a list of unary constraints (*unaryConstraints*). An *Assignment* is constructed from a *ConstraintSatisfactionProblem* and is intended to be updated as you search for the solution. It holds a dictionary from variables to their domains (*varDomains*) and a dictionary from variables to their assigned values (*assignedValues*). Notice that the *varDomains* in *Assignment* is meant to be updated, while the *varDomains* in *ConstraintSatisfactionProblem* should be left alone.

A new assignment should never be created. All changes to the assignment through the recursive backtracking and the inference methods that you will be implementing are designed to be reversible. This prevents the need to create multiple assignment objects, which becomes very space-consuming.

The constraints in the csp are represented by two classes: *BinaryConstraint* and *UnaryConstraint*. Both of these store the variables affected and have an *isSatisfied* function that takes in the value(s) and returns `False` if the constraint is broken. You will only be working with binary constraints, as *eliminateUnaryConstraints* has been implemented of for you. Two useful methods for binary

constraints include *affects*, which takes in a variable and returns True if the constraint has any impact on the variable, and *otherVariable*, which takes in one variable of the binaryConstraint and returns the other variable affected.

Question 1 (4 points): Recursive Backtracking

In this question you will be creating the basic recursive backtracking framework for solving a constraint satisfaction problem. First implement the function *consistent*. This function indicates whether a given value would be possible to assign to a variable without violating any of its constraints. You only need to consider the constraints in `csp.binaryConstraints` that affect this variable and have the other affected variable already assigned.

Once this is done implement *recursiveBacktracking*. There is pseudocode for this in the book (Section 6.3). This function is designed to take in a problem definition and a partial assignment. When finished, the assignment should either be a complete solution to the CSP or indicate failure. Note that for now your implementation will not use any inferences (what forward checking finds), as this will be implemented later.

To test and debug, use both the autograder and the functions in `Testing.py`. Be aware that individual questions and tests can be set for grading to the autograder with `-q` and `-t` respectively.

Question 2 (2 points): Variable Selection

While the recursive backtracking method eventually finds a solution for a constraint satisfaction problem, this basic solution will take a very long time for larger problems. Fortunately, there are heuristics that can be used to make it faster. One place to include heuristics is in selecting which variable to consider next.

Implement *minimumRemainingValueHeuristic*. This follows the minimum remaining value heuristic to select a variable with the fewest options left and uses the degree heuristic to break ties. Both of these heuristics are explained in the book, or in the lecture notes (Section 6.3.1).

Question 3 (2 points): Value Ordering

Another way to use heuristics is to optimize a constraint satisfaction problem solver is to attempt values in a different order.

Implement *leastConstrainingValuesHeuristic*. This takes in a variable and determines the order in which the possible values should be attempted according to the least constraining values heuristic, which prefers values that eliminate the fewest possibilities from other variables. As you might expect at this point, the book has an explanation of the heuristic (Section 6.3.1).

Question 4 (4 points): Forward Checking

While heuristics help to determine what to attempt next, there are times when a particular search path is doomed to fail long before all of the values have been tried. Inferences are a way to identify impossible assignments early on by looking at how a new value assignment affects other variables. Each inference made by an algorithm involves one possible value being removed from one variable. It should be noted that when these inferences are made they must be kept track of so that they can later be reversed if a particular assignment fails. See section 6.2 in the book for more information on inferences.

Implement *forwardChecking*. This is a very basic inference-making function. When a value is assigned, all variables connected to the assigned variable by a binary constraint are considered. If any value in those variables is inconsistent with that constraint and the newly assigned value, then the inconsistent value is removed. See the book for a more detailed explanation. (Section 6.3.2)

Once this is done implement *recursiveBacktrackingWithInferences*. This will mostly be the same as your solution from q1, but slightly changed to include the use of inferences.

Question 5 (4 points): Maintaining Arc Consistency

There are other methods for making inferences than can detect inconsistencies earlier than forward checking. One of these is the Maintaining Arc Consistency algorithm, or the MAC algorithm. First implement *revise*. This is a helper function that is responsible for determining inconsistent values in a variable (Section 6.2.2, Figure 6.3).

Then implement *maintainArcConsistency*. The MAC algorithm starts off very similar to forward checking in that it removes inconsistent values from variables connected to the newly assigned variable. The difference is that it uses a queue to propagate these changes to other related variables. The book contains pseudocode and more information about the algorithm (Section 6.3.2, refers to AC3 method in Section 6.2.2).

Question 6 (2 points): Preprocessing

Another step to making a constraint satisfaction solver more efficient is to perform preprocessing. This can eliminate impossible values before the recursive backtracking even starts. One method to do this is to use the AC3 algorithm.

Implement AC3 (Section 6.2.2). This algorithm is almost identical to MAC in that it propagates inferences in order to remove as many values as possible. The major difference is that instead of starting with one variable, AC3 begins with all variables already in the queue. Thus, reusing MAC is a fast way to implement this AC3. It also does not need to track the inferences that are made, because

if the assignment fails at any point then there is no prior state to back up to. This means that there is no solution to the CSP.

Question 7 (2 points Extra Credit): Novel CSP Problem

Create a new CSP problem in the same format as all the existing problems within the `csp` directory. The problem should not be trivial or arbitrary, but rather a meaningful problem such as the chess queens problem or a similarly interesting problem. This may also involve creating new types of constraints, which is done by extending *BinaryConstraint* and *UnaryConstraint* and implementing *isSatisfied* differently. (See *BadValueConstraint* and *NotEqualConstraint* for examples.) Submit your problem along with your code within a directory called "Extra". You will receive full credit if the problem is well specified, not trivial, and is solved by our CSP implementation.

Submission

Zip only the files you altered for this assignment as a `.zip` or `.tar.gz` and submit it on T-Square before the due date. You have a strict one hour window after the due date to handle any last minute technical issues, after which you will not be able to submit and receive a 0 for the project.