


CS30 - Assignment 9

Due at Sunday, April 10 at 11:59pm



<https://xkcd.com/74/>

For this assignment, you will be using best first search to find solutions to sudoku problems! As always (and particularly for this assignment), read through the entire handout before you start working.

For this assignment, you may (and in fact, I'd encourage you to) work with a partner. If you work with a partner, you only need to submit one file, just make sure both of your names are at the top of it. **If you choose to do work with a partner, you must both be there whenever you're working on the assignment and you should work on each of the problems collaboratively.** In particular, what is not appropriate is to divide the problems up and then combine your answers into a single submission (or, of course, just have one person do all the work :).

1 Starter code

I have provided you with some code to get you started this time that includes the search method implementation, functions for generating sudoku problems (three different ones) and also a few methods in the class.

To get started, download the file:

<http://www.cs.pomona.edu/~dkauchak/classes/cs30/assignments/assign9/assign9-starter.txt>

and save as a file named with your first name and last name followed by `assign9.py`. You should then edit this file with your implementation.

This file has a number of things in it that will be useful:

- At the top, I've started the class definition already for you called `SudokuState`, which you'll be filling in.
- I've written four methods for you in this class:
 - `get_subgrid`: The sudoku board is a 9 by 9 grid that is subdivided into nine 3 by 3 subgrids. This method takes a row and column for an entry in the 9 by 9 grid and returns a number 1 through 9 indicating which subgrid that row/column entry falls into, where subgrids are numbered as follows:

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

- `get_any_available_cell`: See the experiment section for how this will be used.
 - `__str__`: Generates a “pretty” version of the state and board as a string. This looks nice and can be useful for debugging, but it hides information about what’s in the cells that don’t have numbers in them.
 - `get_raw_string`: Generates a more complete version of the state and board as a string. It does include the cells that haven’t been filled in and will also be useful for debugging, though it prints a lot of information.
- I’ve included three methods for generating three different starting states: `problem1`, `problem2` and `heart` (the example from the class notes).
 - I’ve also included the same `dfs` implementation as last time as well as commented code (you’ll uncomment it eventually) that calls `dfs` to get a solution and then prints it out.

2 SudokuState class requirements

Like any search problem, we will need our class to do three key things:

- Construct the starting state
- Determine if a state is a goal state
- From a state, determine what are the “next” states that we can get to by performing an action.

2.1 Specification

To support this functionality, fill in the details for the `SudokuState` class. The class must have the following instance variables:

- `self.size`: the size of a side of the sudoku board (it will always be set to 9 for our implementation, though)
- `self.num_placed`: the number of *valid* numbers that have been put on the board. We are going to make sure that we only create states that do not violate the sudoku constraints.
- `self.board`: a list of lists (i.e. a matrix). Each entry will contain one of two things:
 1. A number 1-9. This will indicate that that entry has been filled.
 2. A list of numbers 1-9. This will indicate that the entry has NOT been filled. The numbers in the list represent valid numbers that *could* still be placed in that entry given the other numbers on the board.

Your class should have *at least* the following methods:

- `__init__`

Should take zero parameters, and construct an empty sudoku board. An empty sudoku board should be a 9 by 9 matrix (list of lists) where each entry is a list with the numbers 1 through 9. For example, the following lines would generate an empty board and print out the “pretty version” (using the `__str__` method, which doesn’t print the board entries if there isn’t a number there) as well as the “raw version” (using the `get_raw_string` method) showing the actual lists being stored in the matrix:

```
empty_state = SudokuState()
print empty_state # print the pretty version
print empty_state.get_raw_string() # print the actual data
```


- `is_filled`

Takes two parameters, a `row` and a `column` and returns `True` if that entry has been filled, i.e. has a number put into it, or `False` if that entry is still empty. Remember, entries that are filled have a number in them, while entries that have a list of possible values are empty.

Hint:

There is a function called `isinstance` that may be useful here. It takes a value and type and returns `True` if that object is of that type, e.g.

```
>>> s = "banana"
>>> isinstance(s, str)
True
>>> isinstance(s, int)
False
>>> isinstance(s, list)
False
>>>
```

- `remove_conflict`

Takes three parameters, a `row`, a `column` and a `number`. *If* the entry at that row and column hasn't been filled in yet, remove `number` from the list of possible values at this entry. This method will be called when a number has been placed in either the same row, column or subgrid as this entry and it will remove that number as a possibility.

Hint: Type `help(list)` in the python shell to see all of the methods that are available for lists or you can look at the documentation online (take a look at which ones might be helpful).

- `remove_all_conflicts`

Takes three parameters, a `row`, a `column` and a `number`. This method is called after `number` has been placed at entry `row`, `column`. It should update all of the other unfilled entries that are in the same row, column or subgrid as this entry so that they no longer have `number` as an option.

For example, if we put a 7 at (1, 2), then all unfilled entries in row 1 and column 2 and in the first subgrid would need to have 7 removed from their list of possible valid values.

Hints:

- * Make sure that you understand what this method is doing. It's making sure that all entries that are affected when we put a number on the board are updated appropriately.
- * The `get_subgrid` and `remove_conflict` methods may be useful here.

- `add_number`

Takes three parameters, a `row`, a `column` and a `number` and returns a new state (i.e. a *copy* of the current state :) with `number` added at `row`, `column`. This method is similar to the `add_queen` method from last time. It should NOT modify the current state, but should produce a new state. You will need to make sure that when you add that number you make sure that the new state stays consistent.

After you implement this method, you should be able to use any of the methods included in the starter code that generate starting boards (e.g. `heart()`).

Hint: Remember that any state that is an instance of a class (whether it be `self` or any other variable that references an object of a class) can call *any* of the methods of that class. This is an important method so make sure you test it thoroughly before moving on (the two “string” methods should be useful for testing this).

- `get_most_constrained_cell`

Takes zero parameters and returns a tuple containing the row and column of the most constrained entry in the board, i.e. the entry that is not filled in yet that has the fewest possible options remaining. Think of this function as our heuristic function that will help guide our search.

- `solution_is_possible`

Takes zero parameters and returns `True` if all entries on the board still have at least one possible value that they can take, `False` otherwise (i.e. there is at least one entry without any possible values that it can take).

- `next_states`

Takes zero parameters and returns a list of the next states that can be reached from this current state by trying to put a number *in the most constrained cell*. You should only return viable states in this list, i.e. states where a solution is still possible.

Hint: To be able to tell if a state is viable, you’ll need to first generate it and then check if a solution is still possible before adding it to the list of states. Note that this is a bit different from your n-queens problem where we test *from* the current state. This is another possible way of doing it.

- `is_goal`

Takes zero parameters and returns whether or not this state is a goal state.

Hint: Like the n-queens problem, this can be answered very easily based on information available in the state, in particular, you can do it without having to traverse all the entries in the board.

3 Where’s best first search?

If you do everything correctly, when you’ve implemented these methods you should be able to uncomment the code at the bottom of your file and solve the sample problems. You may notice, however, that we’re using the exact same `dfs` function that we used before. Why isn’t this depth first search? The key is that the `next_states` method greedily chooses which entry to fill in based on the heuristic to fill in the most constrained entry first.

4 Experiments

Once you have things working, I want you to do a couple of experiments. I have included a method called `find_any_available_cell` that serves a similar functionality to `get_most_constrained_cell`, but it just chooses a random entry (in fact, the first empty one it comes across). In your `next_states` method, you can change the code to use either of these methods for selecting the next entry and `dfs` should still find a solution. If you use the former, it performs a traditional, uninformed search, and if you use the latter, it performs a best first, informed search.

Experiment with these two variants on the three different example problems. Do you see a difference in performance (I've included the time it takes to solve the problem in the information printed out)? Include a short paragraph on your experiments in comments at the top of your file that talk about the results that you found and an explanation.

5 Hints/advice

- Make sure you understand the role that each of these methods serves. Although we only need the `is_goal` and `next_states` methods for `dfs` to work, all of the other methods I'm asking you to implement will be useful in implementing these methods or others.
- I strongly encourage you to implement your class a method at a time and then test each method individually. Since many of the methods will use previous methods, it's important that you're confident that the previous methods are correct. Here are a few examples of how you might test individual methods *as you go*, but I encourage you to also come up with your own tests:

```
- __init__
    s = SudokuState()
    print s
    print s.get_raw_string()

- remove_all_conflicts
    s = SudokuState()
    s.board[0][0] = 1 # set a one in the upper left hand corner
    s.remove_all_conflicts(0, 0, 1)
    print s
    print s.get_raw_string()

- add_number
    s = SudokuState()
    s2 = s.add_number(0, 0, 1)
    print s
    print s.get_raw_string() # make sure s hasn't changed
    print s2
    print s2.get_raw_string() # make sure s2 is the proper state
```

```
– next_states

s = SudokuState()
next = s.next_states()
for s in next:
    print s
    print s.get_raw_string()
```

- I have tried to order the required methods in the order that I would suggest implementing them.
- If your code doesn't work, you need to figure out *why*. One of the best ways to do this is to put in print statements to print out how your code is working. Try to isolate particular functionalities and then make sure they're working like you expect. Make sure, however, to remove these print statements once you fix it!
- Start early! If you hit a problem, come talk to myself or the mentors (after, of course, you've tried to debug it on your own for a little bit).

When you're done

You should have a single .py file with

- 1) the `SudokuState` class,
- 2) the starter code at the bottom uncommented and
- 3) your experiment paragraph in comments at the top, all in a file named with your first name and last name followed by `assign9.py`.

Additionally, make sure that the file is properly commented:

- You should have comments at the very beginning of the file stating your name, course, assignment number and the date.
- Each function should have an appropriate docstring.
- Your class should have an appropriate docstring.
- Include other miscellaneous comments to make things clear.

Submit your .py file online using the courses submission mechanism.

Grading

	points
init	3
is_filled	2
remove_conflict	2
remove_all_conflicts	5
add_number	2
get_most_constrained_cell	3
next_states	3
is_goal	2
experiments	3
comments/style	3
total	28