Goal The purpose of this assignment is to write a program to solve the 8-puzzle problem (and its natural generalizations) using the A^* search algorithm.

The Problem The 8-puzzle problem is a puzzle invented and popularized by Noyes Palmer Chapman in the 1870s. It is played on a 3-by-3 grid with 8 square blocks labeled 1 through 8 and a blank square. Your goal is to rearrange the blocks so that they are in order. You are permitted to slide blocks horizontally or vertically into the blank square. The following shows a sequence of legal moves from an initial board position (left) to the goal position (right).

4 7	1 2 8	3 5 6	=>	1 4 7	2 8	3 5 6	=>	1 4 7	2 8	3 5 6	=>	1 4 7	2 5 8	3 6	=>	1 4 7	2 5 8	3 6
in	iti	al														g	oal	

Best-First Search Now, we describe a solution to the problem that illustrates a general artificial intelligence methodology known as the A^* search algorithm. We define a *search node* of the game to be a board, the number of moves made to reach the board, and the previous search node. First, insert the initial search node (the initial board, 0 moves, and a null previous search node) into a priority queue. Then, delete from the priority queue the search node with the minimum priority, and insert onto the priority queue all neighboring search nodes (those that can be reached in one move from the dequeued search node). Repeat this procedure until the search node dequeued corresponds to a goal board. The success of this approach hinges on the choice of *priority function* for a search node. We consider two priority functions:

- *Hamming priority function*. The sum of the Hamming distance (number of tiles in the wrong position), plus the number of moves made so far to get to the search node. Intuitively, a search node with a small number of tiles in the wrong position is close to the goal, and we prefer a search node that have been reached using a small number of moves.
- *Manhattan priority function.* The sum of the Manhattan distance (sum of the vertical and horizontal distance) from the tiles to their goal positions, plus the number of moves made so far to get to the search node.

For example, the Hamming and Manhattan priorities of the initial search node below are 5 and 10, respectively.

8 1 3	1 2 3	1 2	3 4	56	5 7	8	1	2	3	4	5	6	7	8
4 2 7 6 5	4 5 G 7 8	1 1	0 0	1 1	. 0	1	1	2	0	0	2	2	0	3
initial	goal	Ha	mming :	= 5 +	• 0			Man	hat	tan	=	10	+ 0	

We make a key observation: To solve the puzzle from a given search node on the priority queue, the total number of moves we need to make (including those already made) is at least its priority, using either the Hamming or Manhattan priority function. (For Hamming priority, this is true because each tile that is out of place must move at least once to reach its goal position. For Manhattan priority, this is true because each tile must move its Manhattan distance from its goal position. Note that we do not count the blank square when computing the Hamming or Manhattan priorities.) Consequently, when the goal board is dequeued, we have discovered not only a sequence of moves from the initial board to the goal board, but one that makes the fewest number of moves. Challenge for the mathematically inclined: prove this fact.

A Critical Optimization Best-first search has one annoying feature: search nodes corresponding to the same board are enqueued on the priority queue many times. To reduce unnecessary exploration of useless search nodes, when considering the neighbors of a search node, don't enqueue a neighbor if its board is the same as the board of the previous search node.

8 4 7	1 6	3 2 5	8 4 7	1 2 6	2	3 5	8 4 7	1 2 6	3 5	8 4 7	1	3 2 5		8 4 7	1 2 6	3 5
or	evi	ous	sear	ch	IJ	node	ne	igh	bor	ne (di	igh sal	bor low)	nei	igh	boı

A Second Optimization To avoid recomputing the Hamming/Manhattan distance of a board (or, alternatively, the Hamming/Manhattan priority of a solver node) from scratch each time during various priority queue operations, compute it at most once per object; save its value in an instance variable; and return the saved value as needed. This caching technique is broadly applicable: consider using it in any situation where you are recomputing the same quantity many times and for which computing that quantity is a bottleneck operation.

Game Tree One way to view the computation is as a game tree, where each search node is a node in the game tree and the children of a node correspond to its neighboring search nodes. The root of the game tree is the initial search node; the internal nodes have already been processed; the leaf nodes are maintained in a priority queue; at each step, the A* algorithm removes the node with the smallest priority from the priority queue and processes it (by adding its children to both the game tree and the priority queue).



Detecting Unsolvable Puzzles Not all initial boards can lead to the goal board by a sequence of legal moves, including the two below:

1	2	3	1	2	3	4
4	5	6	5	6	7	8
8	7		9	10	11	12
			13	15	14	

To detect such situations, use the fact that boards are divided into two equivalence classes with respect to reachability: those that lead to the goal board; and those that cannot lead to the goal board. Moreover, we can identify in which equivalence class a board belongs without attempting to solve it.

• Odd board size. Given a board, an *inversion* is any pair of tiles i and j where i < j but i appears after j when considering the board in row-major order (row 0, followed by row 1, and so forth).

	1 2 3 4 5 6 => 8 7	$\begin{array}{cccccccccccccccccccccccccccccccccccc$	1 2 3 4 6 => 8 5 7	1 2 3 4 6 => 8 5 7	1 2 3 8 4 6 5 7
row-major order:	1 2 3 4 5 6 8 7	1 2 3 4 5 6 8 7	1 2 3 4 6 8 5 7	1 2 3 4 6 8 5 7	1 2 3 8 4 6 5 7
	inversions = 1 (8-7)	inversions = 1 (8-7)	inversions = 3 (6-5 8-5 8-7)	inversions = 3 (6-5 8-5 8-7)	inversions = 5 (8-4 8-6 8-5 8-7 6-5)

If the board size n is an odd integer, then each legal move changes the number of inversions by an even number. Thus, if a board has an odd number of inversions, then it cannot lead to the goal board by a sequence of legal moves because the goal board has an even number of inversions (zero).

The converse is also true: if a board has an even number of inversions, then it can lead to the goal board by a sequence of legal moves.

2 3 1 4 2 5 4 7 2 5 4 7 5 5 4 5 4 6 8 8 1 3 4 2 5 7 8 6 row-major order: 1 3 4 2 5 7 8 6 1 2 3 4 5 7 8 6 1 2 3 4 5 7 8 6 1 2 3 4 5 6 7 8 4 2 inversions = 4 inversions = inversions 2 inversions inversions = 0(3-2 4-2 7-6 8-6) (3-2 4-2 7-6 8-6) $(7-6 \ 8-6)$ $(7-6 \ 8-6)$

• Even board size. If the board size n is an even integer, then the parity of the number of inversions is not invariant. However, the parity of the number of inversions plus the row of the blank square is invariant: each legal move changes this sum by an even number. If this sum is even, then it cannot lead to the goal board by a sequence of legal moves; if this sum is odd, then it can lead to the goal board by a sequence of legal moves.

$\begin{array}{cccccccccccccccccccccccccccccccccccc$	1 2 3 4 => 5 6 8 9 10 7 11 13 14 15 12	=> 5 6 7 8 9 10 11 13 14 15 12	1 2 3 4 => 5 6 7 8 9 10 11 13 14 15 12	1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
blank row = 1 inversions = 6 sum = 7	blank row = 1 inversions = 6 sum = 7	blank row = 2 inversions = 3 sum = 5	blank row = 2 inversions = 3 	blank row = 3 inversions = 0

Problem 1. (*Board Data Type*) Implement an immutable data type called **Board** to represent a board in an *n*-puzzle, supporting the following API:

<pre>Board(int[][] tiles)</pre>	constructs a board from an $n \times n$ array; tiles[i][j] is the tile at row i and column j , with 0 denoting the blank tile
<pre>int size()</pre>	returns the size of this board size
<pre>int tileAt(int i, int j)</pre>	returns the tile at row i and column j
int hamming()	returns Hamming distance between this board and the goal board
int manhattan()	returns the Manhattan distance between this board and the goal board
boolean isGoal()	returns true if this board is the goal board, and false otherwise
<pre>boolean isSolvable()</pre>	returns true if this board solvable, and false otherwise
<pre>Iterable<board> neighbors()</board></pre>	returns an iterable object containing the neighboring boards of this board
<pre>boolean equals(Object other)</pre>	returns true if this board is the same as other, and false otherwise
String toString()	returns a string representation of this board

Performance Requirements

- The constructor should run in time $T(n) \sim n^2$, where n is the board size.
- The size(), tileAt(), hamming(), manhattan(), and isGoal() methods should run in time $T(n) \sim 1$.
- The isSolvable() method should run in time $T(n) \sim n^2 \log n^2$.
- The neighbors() and equals() methods should run in time $T(n) \sim n^2$.

```
\times ~/workspace/eight_puzzle
```

```
$ java Board data/puzzle05.txt
The board (3-puzzle):
    1
 4
       3
    2
       6
 7
    5
       8
Hamming = 5, Manhattan = 5, Goal? false, Solvable? true
Neighboring boards:
 4
    1
      3
 7
    2
      6
    5 8
  _ _ _ _ _ _ _ _ _
    1
       3
 4
    2
      6
 7
    5
      8
       _ _ _ _
 4
   1
       3
 2
       6
7
   58
   _ _ _ _ _ _ _ _
$ java Board data/puzzle4x4-unsolvable1.txt
The board (4-puzzle):
3
   2
       4 8
    6
         12
 1
5 10 7 11
9 13 14 15
Hamming = 12, Manhattan = 13, Goal? false, Solvable? false
Neighboring boards:
   2
3
      4 8
 1
    6
      7 12
```

5	10		11
9	13	14	15
3	10	17	10
3	2		8
1	6	4	12
5	10	7	11
9	13	14	15
3	2	4	8
1	6	12	-
5	10	- 2	11
0	10	1	1 -
9	13	14	15
3	2	4	8
1		6	12
5	10	7	11
9	13	14	15

Problem 2. (Solver Data Type) Implement an immutable data type called Solver that uses the A^* algorithm to solve the 8-puzzle and its generalizations. The data type should support the following API:

Solver(Board board)	finds a solution to the initial board using the A^* algorithm
int moves()	returns the minimum number of moves needed to solve the initial board
<pre>Iterable<board> solution()</board></pre>	returns a sequence of boards in a shortest solution of the initial board

Corner Cases

• The constructor should throw a NullPointerException("board is null") if *board* is null and an IllegalArgumentException("board is unsolvable") if *board* is unsolvable.

```
\times ~/workspace/eight_puzzle
$ java Solver data/puzzle05.txt
Solution (5 moves):
    1
 4
        3
     2
        6
    5
 7
        8
     _ _
        _ _ _ _
     1
         3
 4
    2
        6
 7
    5
        8
_ _ _
    _ _ _ _ _ _
 1
         3
    2
 4
         6
 7
     5
         8
     _ _
        - - - -
 _
        3
 1
     2
 4
         6
 7
    58
    _ _ _
        _ _ _
_ _ _
         3
 1
     2
    5
 4
         6
```

```
7 8
------
1 2 3
4 5 6
7 8
------
$ java Solver data/puzzle4x4-unsolvable1.txt
Unsolvable puzzle
```

Data and Test Programs The data directory contains a number of sample input files representing boards of different sizes; for example

```
x ~/workspace/eight_puzzle
$ more data/puzzle04.txt
3
0 1 3
4 2 5
7 8 6
```

The program SolverVisualizer accepts the name of an input file as command-line argument, and using your Board and Solver data types graphically solves the sliding block puzzle defined by the file

```
x ~/workspace/eight_puzzle
$ java SolverVisualizer data/puzzle04.txt
```

	datapuzzie04.tet	
1	2	3
4	5	6
7	8	
44 moves	10	ethattan distancer 0

The program PuzzleChecker accepts the names of an input files as command-line arguments, creates an initial board from each file, and writes to standard output: the filename, minimum number of moves to reach the goal board from the initial board, and the time (in secs) taken; if the initial board is unsolvable, a "-" is written for the number of moves and time taken

\times ~/workspace/eight_puzzle		
<pre>\$ java PuzzleChecker data/puzzle*.tx</pre>	t	
filename	moves	time
data/puzzle00.txt	0	0.00
data/puzzle01.txt	1	0.00
data/puzzle02.txt	2	0.00
data/puzzle03.txt	3	0.00
data/puzzle04.txt	4	0.00
data/puzzle05.txt	5	0.00
data/puzzle47.txt	47	9.41
data/puzzle48.txt	48	2.13
data/puzzle49.txt	49	19.63

data/puzzle4x4-unsolvable1.txt		
data/puzzle50.txt	50	12.31

Files to Submit

- 1. Board.java
- $2. \ {\tt Solver.java}$
- 3. notes.txt

Before you submit your files, make sure:

- You do not use concepts from sections beyond *Sorting Applications*.
- Your code is clean, well-organized, uses meaningful variable names, includes useful comments, and is efficient.
- You edit the sections (#1 mandatory, #2 if applicable, and #3 optional) in the given notes.txt file as appropriate. In section #1, for each problem, state its goal in your own words and describe your approach to solve the problem along with any issues you encountered and if/how you managed to solve those issues.

Acknowledgements This assignment is an adaptation of the Eight Puzzle assignment developed at Princeton University by Robert Sedgewick and Kevin Wayne.