A simple Minesweeper algorithm

Mike Sheppard*

October 9, 2023

Abstract

This paper rephrases the Minesweeper problem in terms of information theory and entropy optimization, allowing greater flexibility in applications and a simplified algorithm. The goal is to find the most probable distribution of mine locations, given the information that is currently available. All information must be taken into account and no additional information not given may be used. The solution is optimal in the context of information theory.

1 Introduction

Minesweeper is a popular logic game where a set of hidden mines are placed across a two dimensional grid. A player takes turns selecting squares trying to deduce the location of all the mines without accidently selecting a mine itself which ends the game. Each square can either be: (1) empty - revealing any connected empty squares, (2) a number - representing how many mines are adjacent to that square, (3) a mine - ending the game. The player can place flags on any number of squares they believe contain mines. The game is won when all non-mine squares have been revealed. An example of an easy level game in mid-play is shown to the right, with six out of the ten hidden mines being able to be determined with certainty at current play. (Hint: top left unselected square is one mine). Flags can be placed on those squares to remind the player where they believe the mines are located - the flags can be removed or changed at any time if prior reasoning may be incorrect.

Many academic papers have been written about Minesweeper, its properties, and algorithms on how to play; including a full website dedicated to collecting those references [1]. The papers include algorithms such as Boolean Satisfiability Problem, combinatorial reasoning, cellular automaton, genetic programming, graph theory, neural networks, and the use of quantum logic gates to name a

*Author contact information : https://www.linkedin.com/in/mike-sheppard-4762314/

few. All with varying degrees of complexity and application versus theory. It is known that Minesweeper is NP-Complete.

This paper rephrases the Minesweeper problem in terms of information theory and entropy optimization, allowing greater flexibility in applications and a simplified algorithm. The goal is to find the most probable distribution of mine locations, given the information that is currently available. All information must be taken into account and no additional information not given may be used. The solution is optimal in the context of information theory.

2 Algorithm

Let $H(X) = -\sum_{i=1}^{n} p_i \log_2 p_i$ represent the Shannon entropy of some discrete probability distribution with $\sum_{i=1}^{n} p_i = 1$ [2]. According to Jaynes, the probability distribution which best represents the current state of knowledge about a system is the one with largest entropy [3]. If the constraints are given in the form of expected values then Jaynes' principle of maximum entropy is the uniquely correct method for inductive inference and maximizing any function but entropy will lead to inconsistency [4]. Any other distribution, other than the maximum entropy distribution, would contain additional information or assumptions not given to us which may or may not be correct.

Let p_i represent the probability the i^{th} square contains a mine, with $p_i = 1$ being certain a mine is present. For example, the game board above has six squares in which $p_i = 1$ can be deduced by logical reasoning alone. Alternatively, some other squares can also be determined to be clear of any mines, in which case $p_i = 0$. Any other scenario the square is in some uncertain region where the likelihood it contains a mine is determined by the information given by the game, but cannot be logically deduced as certain one way or another. The principle of maximum entropy can be used to determine the most probable distribution of mine locations consistent with the given board. The displayed counts of mines turn into expected value constraints across the probabilities.

Notation: Let $i = \{1, \ldots, n\}$ represent the n unselected squares in some user-defined order. Let $j = \{0, \ldots, m\}$ represent the $m + 1$ constraints on the mine locations. These include the zeroth constraint of total number of mines across all global unselected squares, as well as the m numbers shown yielding local information. The zeroth constraint can be ignored if total number of mines is unknown or not given. For each j there is a set, Ω_j, which are the relevant indices of i that contribute to mine count shown in number c_j.

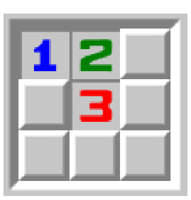

Consider the near-trivial game to the left, used for illustration purposes. There are $m = 3$ revealed numbers, with values $c = \{1, 2, 3\}$, along with $c_0 = 3$ representing the three hidden mines (not labeled). Let the order of the $n = 6$ unselected squares go from left to right, top to bottom. Then, for example, for $j = 2$ we have $c_2 = 2$ and $\Omega_2 = \{1, 2, 3\}$ representing the constraint $p_1 + p_2 + p_3 = 2$. The $m + 1$ constraints are $\sum_{i \in \Omega_j} p_i = c_j$ for $j = \{0, \ldots, m\}$.

As each square individually can be considered a Bernoulli random variable of probability of being a mine, the correct optimization is the sum of Shannon entropies across the n Bernoulli distributions as a collective. As all unselected squares are equal in importance, the total entropy to be optimized is the unweighted sum of the individual Bernoulli distribution entropies. Let $H_b(p) = -p\log_2(p) - (1-p)\log_2(1-p)$ represent the binary entropy function. Then,

$$\begin{aligned} \max_{p} \quad & \sum_{i=1}^{n} H_b(p_i) \\ \text{s.t.} \quad & \sum_{i\in\Omega_j} p_i = c_j \quad j = \{0,\ldots,m\} \end{aligned} \tag{1}$$

is the proper optimization collectively across the n Bernoulli distributions subject to the $m+1$ constraints. This can be solved using standard optimization solvers, such as gradient descent. However, for this particular problem a simpler algorithm is possible. By relabelling the complementary probability $1-p_i$ as a new variable q_i, along with additional constraint $p_i + q_i = 1$, the binary entropy function is now the sum of two singular Shannon entropy terms,

$$\begin{aligned} \max_{p,q} \quad & \sum_{i=1}^{n} (-p_i\log_2(p_i)) + (-q_i\log_2(q_i)) & \\ \text{s.t.} \quad & \sum_{i\in\Omega_j} p_i = c_j & j = \{0,\ldots,m\} \\ & \sum_{i\in\Omega_j} q_i = |\Omega_j| - c_j & j = \{0,\ldots,m\} \\ & p_i + q_i = 1 & i = \{1,\ldots,n\} \end{aligned} \tag{2}$$

The inclusion of newly additional $m+1$ constraints on q_i's represent the complementary constraint for each j. For example, if out of 8 unselected squares there are 3 known mines, then the sum of q_i's would represent the constraint of how many squares are not mines, in this case 5. At first this seems to complicate the problem, going from n unknowns and $m+1$ constraints to $2n$ unknowns and $2(m+1)+n$ constraints. However, in this formulation the optimization is the sum of individual Shannon entropy terms with linear constraints being subsets of inclusion or exclusion for each probability, all coefficients being zero or one. This means a simplified version of generalized iterative scaling [5] can be used, with slight extension for the constraints going simultaneously across n probability distributions, $\{p_i, q_i\}$ for $i = \{1,\ldots,n\}$.

This yields the relatively simple algorithm below, using only proportional fitting and basic arithmetic operations.

Minesweeper - Entropy optimization

For each $j = \{0, \ldots, m\}$ of known constraints on number of mines and location, let c_j be the number of mines indicated and Ω_j be the set of $i = \{1, \ldots, n\}$ unselected squares relevant to constraint j.

Initialize

— $\{p_i, q_i\} \leftarrow \{1, 1\}$ for $i = \{1, \ldots, n\}$

Iterate until convergence

— **For** $j = \{0, \ldots, m\}$

— — **If** $\sum_{i \in \Omega_j} p_i \neq c_j$

— — — $p_i \leftarrow p_i \left(\frac{c_j}{\sum_{i \in \Omega_j} p_i} \right)$ for $i \in \Omega_j$

— — **If** $\sum_{i \in \Omega_j} q_i \neq |\Omega_j| - c_j$

— — — $q_i \leftarrow q_i \left(\frac{|\Omega_j| - c_j}{\sum_{i \in \Omega_j} q_i} \right)$ for $i \in \Omega_j$

— $\{p_i, q_i\} \leftarrow \{p_i, q_i\}/(p_i + q_i)$ for $i = \{1, \ldots, n\}$

return $\{p_i, q_i\}$

In words the algorithm is quite simple: First initialize all pairs $\{p_i, q_i\}$ to one. Then for each constraint if the sum of the corresponding probabilities is a multiplicative factor off from the known given number, multiply all the corresponding probabilities by that same factor so it now matches the given sum. This is true for p_i's with number of mines and q_i's with number of non-mines. Once all constraints have been updated, re-normalize each pair of probabilities and repeat until convergence.

The convergence criteria is up to the user, as machine-precision may not be necessary. It could be until third decimal place doesn't change, or until the set of unselected squares with matching lowest probability doesn't change - in that case it is the set of safest squares that can be played which is of importance and not necessarily their probability. Any square given $p_i = 1$ can be flagged immediately and any set of squares with $p_i = 0$ is certain to be free of mines.

The output $\{p_i, q_i\}$ is identical to the maximum entropy solution and contains all information across the constraints without imposing any other assumption. The squares with the lowest p, or highest q, should be selected for next turn.

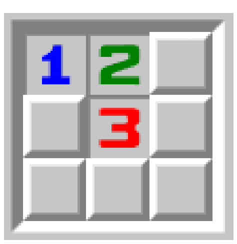

As an example, the output of the algorithm for the following puzzle is

$p = \{0.5, 1, 0.5, 0.333, 0.333, 0.333\}$

$q = \{0.5, 0, 0.5, 0.666, 0.666, 0.666\}$

There is a 100% certainty the second unselected square (under the blue 1) is a mine, with 33% chance of the bottom three. The second square should be flagged as being a mine, and one of the bottom three selected for next turn with $q = 0.666$ probability being safe.

3 Combinatorics

The rationale behind the optimzation and algorithm can also be explained differently using the combinatorial argument of Wallis to Jaynes [6].

Consider using discrete fully intact indivisible mines. The full-mines are randomly placed in the unselected squares and any combination not consistent with the numbers shown on the board are ignored. Among the valid placements some assignments will have a higher number of ways of randomly being allocated than others, their combinatorial multiplicity.

Next consider the mines being divisible, in this case half-mines. Twice as many half mines are randomly placed, with maximum two half-mines per square, and those assignments consistent with the board are kept. Some valid assignments will have higher multiplicities than others. Continue dividing the mines into thirds, fourths, etc. repeating the random assignments and keeping track of those consistent with the board constraints. In the limit the mines are infinitely divisible and we have reached the continuum. The most probable result is the one which maximizes the multiplicity in that limit. Instead of maximizing the multiplicity directly a monotonic function is maximized, the Shannon entropy. This is equivalent to Jaynes' principle of maximum entropy.

4 Conclusion

In this paper we present a new algorithm for Minesweeper which in its simplicity solves for the maximum entropy solution consistent with all constraints given on the board. In terms of information theory it solves for the most probable distribution of mine locations. The algorithm is general to go beyond two dimensional grids, and is also applicable to variants such as non-rectangular, 3D, Hexagonal, or Triangular to name a few.

While other procedures exists which solve the entropy optimization problem faster, the algorithm presented here is the easiest to program and understand. Even for expert level of Minesweeper with 99 mines across a board of dimensions 16 x 30, the iteration to convergence should be short in time.

References:

1. "Minesweeper Math Papers" last accessed October 8, 2023 https://minesweepergame.com/math-papers.php

2. (a) Shannon, Claude E., July 1948. "A Mathematical Theory of Communication". Bell System Technical Journal. 27 (3): 379–423
(b) Shannon, Claude E., Oct. 1948. "A Mathematical Theory of Communication". Bell System Technical Journal. 27 (4): 623–656

3. (a) Jaynes, E. T. (1957). "Information Theory and Statistical Mechanics" Physical Review. Series II. 106 (4): 620–630.
(b) Jaynes, E. T. (1957). "Information Theory and Statistical Mechanics II" Physical Review. Series II. 108 (2): 171–190

4. Shore, John, and Rodney Johnson. "Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy." IEEE Transactions on information theory 26.1 (1980): 26-37.

5. Darroch, J.N.; Ratcliff, D. (1972). "Generalized iterative scaling for log-linear models". The Annals of Mathematical Statistics. 43 (5): 1470–1480.

6. Jaynes, E. T. (2003) "Probability Theory: The Logic of Science", Cambridge University Press

Minesweeper Online: https://minesweeper.online/
(used for game board image screenshots for illustration)