Dynamic Programming for Surprise Examination Problems

Mathias Winther Madsen
m.w.madsen@uva.nl

Institute for Logic, Language, and Computation
University of Amsterdam

April 10, 2014

The surprise examination paradox is a logical puzzle based on the fact that sentences like "You will be surprised at the value of X" can convey information about X and thus indirectly change your levels of surprise.

Such a negative feedback loop between the truth of the sentence and the level of surprise can lead to some rather curious phenomena. A classic example concerns the case in which a teacher who announces to the class:

1. On one of the next five days, you will receive an exam.
2. On the day you receive the exam, it will come as a surprise to you.

When "surprise" is formalized in the framework of epistemic logic, this promise is, contrary to everyday experience, impossible to keep: Whatever set of days the teacher restricts the exam to, the surprise element is always spoiled on the last of those days, since the students will be able to infer by then that the exam is due. The students can thus always eliminate the last day from such a set of possible exam days until no alternatives remain. Thus, no strategy on the part of the teacher can fulfill these two criteria.

This conflict between logical analysis and everyday experience has prompted a number of prominent philosophers to suggest ways of resolving the paradox, either reanalyzing the teacher's announcement in ways that gives it well-defined truth conditions, or by exposing flaws in the reasoning of the students [5, 4].

My talk will present a different tack on the problem. In a probabilistic framework, the notion of "surprisal" is open to more fine-grained quantification [1], and the paradox takes on a different flavor. The logical problem of model construction can be replaced by the problem of picking a probability distribution which maximizes surprise even when the distribution is revealed in advance [6]. The ideas that I present here have been anticipated in unpublished form by Karl Narveson and reported by Timothy Chow [3]. I will, however, expand and generalize the idea enough that I think it will be worth a note of its own.

1 Graph Representations

In order to operationalize the idea of a probabilistic surprise examination problem, I will model a surprise examination as a walk on a finite directed graph. I will assume that the graph is equipped with "start" and "finish" nodes, and that a a subset of the edges constitute a specially designated set of "critical transitions" (shown as fully drawn lines in Fig. 1).

In this scheme, the teacher's choices in a specific realization of the problem are represented as a walk through the graph. A probability distribution over this set of choices can then be represented as a distribution over the possible paths. Such distributions can also be specified by giving a conditional probability distributions for each node. Any interaction between earlier choices and later options (e.g., you can't spend your money twice) can be modeled by having different nodes represent different states of the world at particular times, with node connections restricted as appropriate.

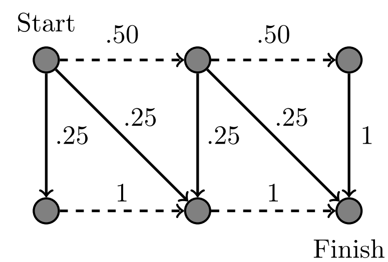

Figure 1: A surprise examination graph, including a probability distribution over paths.

The concept of "critical" edges is intended to capture the idea that certain actions (like having the exam) are so intrinsically important that they can generate surprise on the part of the students, and more surprise when they have low probability. The non-critical actions (like not having the exam) are those that will fail to generate surprise however improbable they are. The selection of what "counts" in terms of criticality can thus significantly change the nature of the problem faced by the teacher.

As an example of such a representation, the classical five-day surprise examination problem can be translated into a graph as shown in Figure 2. The graph has five distinct time points and two levels, corresponding to a state where the exam has not yet been held or a state where it has. Since the exam will, by assumption, only be held once, the arrows only lead downwards and never back up. Since it further must be held before Friday, it also has to transition to the Finish node on the last day.

2 Surprisal Measures

Having represented the skeleton of the surprise examination problem, we need to settle on a quantitative measure of "surprise." It is natural to require such a statistic to be a decreasing function of the probabilities, but this still leaves a large number of options. A conservative choice would be the negative logarithm

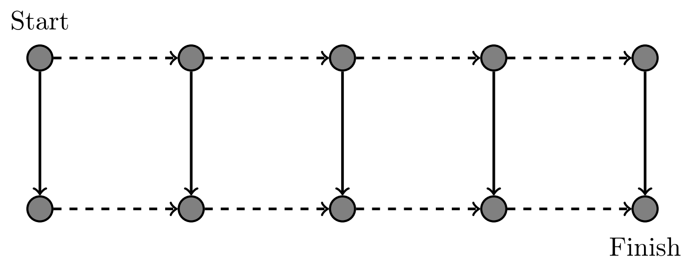

Figure 2: The surprise examination problem in its original form.

of the probability,

$$s(x) = \log \frac{1}{p(x)}.$$

This is the measure occasionally called "surprisal" in information theory, and it corresponds to optimal codeword length. If we disregard the issue of non-critical events, the expected value of the surprisal is the entropy, and it measures the most austere use of output symbols per input symbol we can achieve when encoding a text.

However, we are not narrowly interested in coding here, so other legitimate surprisal functions exist as well. Each of them has an expected value which is an interesting descriptive statistic in its own right:

1. $s(x) = 1/p(x)$: The counting measure; its expected value is the number of values x can take.
2. $s(x) = 1/p(x)^r$: Generalized counting measures, more sensitive to deviations from the uniform distribution when r is large.
3. $s(x) = 1 - p(x)$: The recurrence measure; its expected value is the probability that two independent samples of x will have different values.

Each of these measures of surprise have characteristic properties and interesting asymptotic relationships. However, I will use $s(x) = -\log p(x)$ without further argument, simply because it is the most conventional.

3 Relation to Entropy

Given that we choose $s(x) = -\log p(x)$ our surprisal measure, there is a natural simple connection between maximum entropy problems and surprise examination problems. Seeing this connection requires a few steps.

First, surprisal is a negative logarithm of a probability; since probabilities are smaller than 1, this means that surprisals are larger than 0. It follows that both entropies and solutions to maximum surprisal problems are non-negative.

Second, joint probabilities like $p(x_1, x_2, x_3)$ can, by the chain rule of probability theory, be decomposed in any order, e.g.,

$$p(x_1, x_2, x_3) \;=\; p(x_1)\, p(x_2 \,|\, x_1)\, p(x_3 \,|\, x_1, x_2).$$

Consequently, the surprisal at a joint event $x = (x_1, x_2, x_3)$, which we might think of as a path through a graph, can be decomposed into a sum of conditional surprisal values, one per node visited:

$$\log \frac{1}{p(x_1, x_2, x_3)} \;=\; \log \frac{1}{p(x_1)} + \log \frac{1}{p(x_2 \,|\, x_1)} + \log \frac{1}{p(x_3 \,|\, x_1, x_2)}.$$

We can thus see the surprisal associated with a path either as a surprisal produced by revealing the whole path at once, or a grand total accumulated node by node:

$$s(X) \;=\; s(X_1, X_2, X_3) \;=\; s(X_1) + s(X_2 \,|\, X_1) + s(X_3 \,|\, X_1, X_2).$$

Translating this identity into the world of expected values rather than individual samples, we get the chain rule of entropy: $H(X_1, X_2) = H(X_1) + H(X_2 \,|\, X_1)$, and its inductive generalizations.

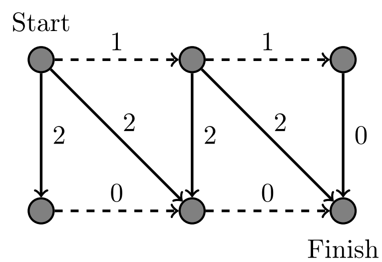

Figure 3: Surprisal values for the transitions in the graph above (Fig. 1).

Lastly, the set of critical edges between the nodes in the graph is a subset of the set of all edges; as a consequence, you can turn a sample entropy into a surprise examination score by weeding out the non-critical transition surprisals from the sum. Thus, if we let

$$S(X) = \sum_{x \in C} p(x)\, \log \frac{1}{p(x)}$$

be the surprisal associated with the critical transitions only, then

$$0 \leq S(X) \leq H(X).$$

As a further consequence, this means that the maximum surprisal that the teacher can obtain is smaller than the maximum entropy on the same graph:

$$0 \leq S_{\max} \leq H_{\max}.$$

This is most easily seen by remembering that the random path X^* that achieves the maximum surprisal on the critical edges produces a weakly larger surprisal if we count all the edges, $S(X^*) \leq H(X^*)$; and since any entropy is smaller than the maximum entropy, the solution to the surprise examination problem is smaller than the solution to the maximum entropy problem.

4 Informal Examples

As I stated in the beginning of this note, the surprise examination problem is due to a conflict between two opposing forces: The fact that the teacher wants to do surprising things, and the fact that things that are done frequently are less surprising.

The simplest illustration of this phenomenon in the case in which there are two options available, one critical and one non-critical. If the teacher chooses to perform the critical action with probability p, the payoff from this action will be $-\log p$, and the average payoff will be $-p \log p$. The optimal choice from the perspective of the teacher is thus to balance the increasing function p against the decreasing function $-\log p$. This is achieved for $p^* = 1/e \approx .37$, as illustrated in Figure 4.

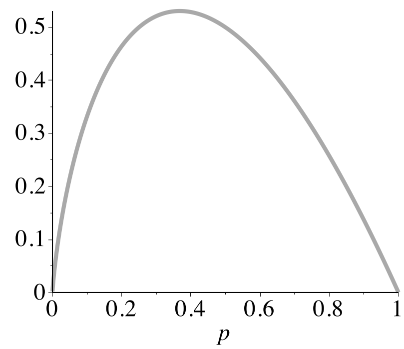

Figure 4: A graph of $-p \log p$.

If we use logarithms to the base of 2 — that is, if we measure the surprisal in bits — then the expected level of surprise at this optimum is about .53 bits. This is different from and smaller than the corresponding value for the maximum entropy distribution, which is achieved at $p^* = 1/2$ and has an expected surprisal of 1 bit. This result is thus consistent with the comments made in the previous section.

Scaling up slightly, consider the surprise examination problem with three possible examination days. Let p be the probability that the exam is placed on the first day, and let q be the probability that it is placed on the second day given that it was not placed on the first. Then the average surprisal about the timing of exam will be

$$s(p, q) \;=\; p \log \frac{1}{p} + (1-p) q \log \frac{1}{q}.$$

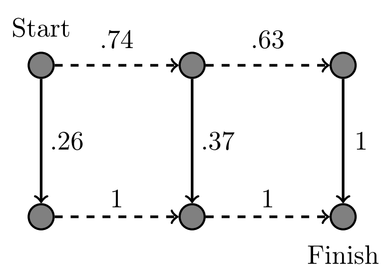

Figure 5: The solution to the three-day surprise examination problem.

Here, p and $(1-p)q$ are the probabilities of having the exam on the first and the second day, respectively; and $-\log p$ and $\log q$ are the corresponding surprisal levels conditioned on the knowledge of the students at that time. We could also have included a third term with weight $(1-p)(1-q)$, but this would have value 0 since there is no conditional uncertainty about the date of the exam on the last day.

As before, the choice between having the exam the second day or having it the third is essentially the choice between a critical event and one that will produce no surprise. In optimum, we thus have $q^* = 1/e$. The surprisal function can thus be reduced to

$$s(p, q^*) \;=\; p \log \frac{1}{p} + (1-p)\frac{\log e}{e},$$

and the maximum of this function is achieved at

$$p^* \;=\; \frac{1}{e^{1+1/e}} \;\approx\; .26.$$

The highest achievable examination surprise in a two-day week is thus about $s(.26, .37) \approx .90$. By contrast, the maximum entropy distribution puts equal weight on all the three paths through the graph, corresponding to the parameter settings $p^* = 1/3$ and $q^* = 1/2$ and an entropy of $H = \log 3 \approx 1.58$.

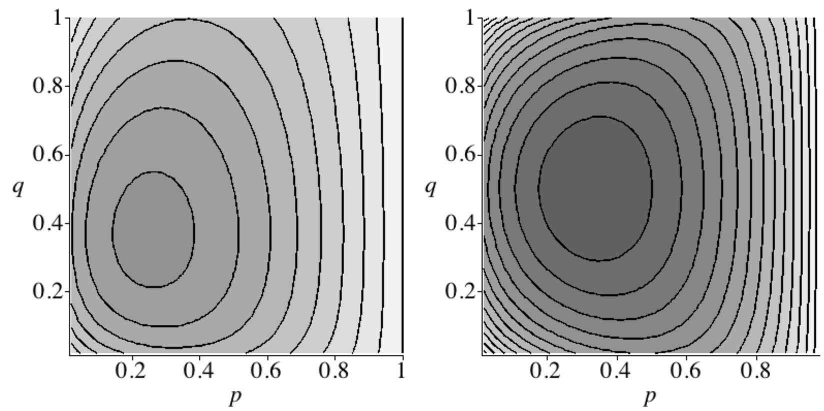

Figure 6: Left: Average surprisal for a two-day surprise examination problem, considering only exams as critical events. Right: Average surprisal when all events are considered critical (i.e., entropy).

5 Solution Method

Surprise examination problems have a so-called "optimal substructure" which means that they can be solved by means of dynamic programming [2]. This is a kind of bottom-up recursion which can exploit repetitions in search trees to solve nominally exponential problems in polynomial time. In the present case, it simply means that we solve the problem backwards, working from finish to start.

More specifically, $S_{\max}(X \mid m)$, the highest surprise that the teacher can achieve starting in state m, depends only on

1. the transition probabilities $p_{m,n}$ from the present state m to the later states n;

2. the highest surprisal that can be achieved at those later states, $S_{\max}(X \mid n)$.

The value of $S_{\max}(X \mid m)$ at any specific node m thus only depends on nodes that are further ahead in the future. We can therefore solve the surprise examination problem by working upwards from the latest, most local subproblems, keeping track of partial results and making sure that we never solve the same problem twice. This is the same procedure that many philosophers know in the form of backwards induction in centipede games.

The recurrence step in this computation is defined by the equality

$$\begin{aligned} S_{\max}(X \mid m) &= \sum_{n \in V(m)} p_{m,n} \left(\mathbb{I}_C(m,n)\, s(p_{m,n}) + S_{\max}(X \mid n) \right) \\ &= \sum_{n \in V(m)} p_{m,n} S_{\max}(X \mid n) + \sum_{n \in C(m)} p_{m,n} s(p_{m,n}). \end{aligned}$$

where $V(m)$ is the set of states accessible from m, and $C(m) \subseteq V(m)$ is the set of states that can be reached by a critical transition. The base case is the deterministic halting at the Finish node, $S_{\max}(X \mid \text{Finish}) = 0$.

In each level of recursion, we thus have to solve a problem of the form

$$\max_p \left(\sum_{n \in V} p_n S_n + \sum_{n \in C} p_n \log \frac{1}{p_n} \right) \text{ subject to } \sum_{n \in V} p_n = 1$$

This problem can be solved using Lagrange multipliers, but there is a couple of technicalities to be aware of:

1. All non-critical transitions except the one with the highest expected surprisal can fortunately be disregarded; this is so because the the teacher's payoff does not depend negatively on the transition probabilities to non-critical states.

2. When some accessible states have a large expected surprisal, the problem may have a corner solution, i.e., $p^*_{m,n} = 1$ or $p^*_{m,n} = 0$ for some n. This happens necessarily when $S_{\max}(X \mid n)$ exceeds $1/\ln 2 \approx 1.44$, but can also happen in other cases.

However, when the complications due to corner solutions do not obtain, the solution to the recurrence step is

$$p_{m,n} = \frac{2^{S_n - S^*}}{e},$$

where $S_n = S_{\max}(X \mid n)$, and S^* is the largest expected surprisals that can be achieved among the non-critical states from m.

I emphasize again, however, that this formula will fail to yield a correct solution when there is a transition for which $p_{n,m} = 0$ or $p_{n,m} = 1$. In the former case, such solutions can be found by first eliminating the coordinate in question and then using the solution formula.

As an example of the results this solution formula gives, an optimal allocation of probability mass is shown in Figure 7. Notice again that this figure illustrates a single recursive step only, not a complete solution to a problem.

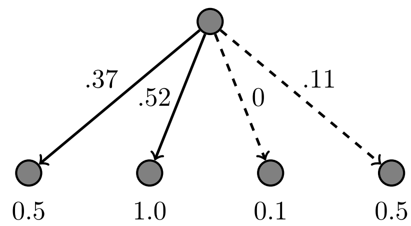

Figure 7: An optimal conditional distribution with four potential next states.

6 Conclusion

The surprise examination problem does not have a logical solution, and this can cause some confusion, because we all know that it of course has a solution. Reformulating the problem in terms of probabilities makes the tensions inherent in the paradox tangible and, more importantly, computable. I have here tried to show how our intuitions about tradeoffs and conflicts in situations like the surprise examination problem are, if not "solved," then at least open to analysis in statistical terms.

References

[1] Fred Attneave. *Applications of Information Theory to Psychology: A Summary of Basic Concepts, Methods, and Results.* Henry Holt and Company, New York, 1959.

[2] Richard Bellman. On the theory of dynamic programming. *Proceedings of the National Academy of Sciences of the United States of America*, 38(8):716, 1952.

[3] Timothy Y. Chow. The Surprise Examination or Unexpected Hanging Paradox. *American Mathematical Monthly*, pages 41–51, 1998.

[4] David Kaplan and Richard Montague. A Paradox Regained. *Notre Dame Journal of Formal Logic*, 1(3):79–90, 1960.

[5] Willard Van Orman Quine. On a So-Called Paradox. *Mind*, 62(245):65–67, 1953.

[6] H. von Stackelberg. *Marktform und Gleichgewicht.* Julius Springer, Vienna, 1934.