Noesis

The Journal of the Noetic Society
Number 40
July 1989

Editorial

Chris Cole
P. O. Box 9545
Newport Beach, CA 92658

The annual meeting this year is scheduled for August 5 and 6, Saturday and Sunday. The meeting will be held in a conference room at the Oakbrook Hills Hotel and Resort, 3500 Midwest Road, Oakbrook, Illinois. Members are responsible for their own transportation, room and board. Rooms are $69 per night, which is a good price for this hotel, and we are getting a free conference room based upon my representation that five to ten people will be staying in the hotel for two nights. Please don't make me a liar. For reservations call June Miell at (312) 850-5555 and be sure to mention the Noetic Society.

I will be attending the annual meeting of the National Puzzlers League in Cleveland on July 21, 22 and 23. The NPL was founded on July 4, 1883 and claims as members most of the professional word puzzlers in the country. They publish *The Enigma* monthly, filled with very hard word puzzles. I hope to distribute copies of Ron Hoeflin's Mega Test at the meeting. I will be reporting on the meeting in a future issue. If you cannot wait, you can join by sending $10 to:

David A. Rosen
207 East 27th Street, #3K
New York, NY 10016

Another word puzzle-oriented publication, which has more of a computer orientation, is *Word Ways, The Journal of Recreational Linguistics*. It is published quarterly and you can subscribe by sending $15 to:

A. Ross Eckler
Spring Valley Road
Morristown, NJ 07960

If anyone knows of analogous organizations in the mathematical or scientific sphere, please let us know.

Puzzles

(solutions on last page)

Physics

A boy, a girl and a dog are standing together on a long, straight road. Simulataneously, they all start walking in the same direction: The boy at 4 mph, the girl at 3 mph, and the dog trots back and forth between them at 10 mph. Assume all reversals of direction instantaneous.

In one hour, where is the dog and in which direction is he facing?

Logic

The moderator takes a set of 8 stamps, 4 red and 4 green, known to the logicians, and loosely affixes two to the forehead of each logician so that each logician can see all the other stamps except those 2 in the moderator's pocket and the two on her own head. He asks them in turn if they know the colors of their own stamps:

A: "No"
B: "No"
C: "No"
A: "No"
B: "Yes"

What are the colors of her stamps?

Mathematics

Prove that in any group of 6 or more people, there exists a subgroup of either 3 mutual acquaintances or 3 mutual strangers.

English

What acronyms have become common words?

Trivia

Two people are talking long distance on the phone; one is in an East- Coast state, the other is in a West-Coast state. The first asks the other "What time is it?", hears the answer, and says, "That's funny. It's the same time here!" How can that be?

Bayesian Regression - III

This article, the last in this series, will tie together several of my earlier puzzles and articles and hopefully make sense of the lot. I finish with some disturbing speculations.

Bayesian inference is based on Bayes' Theorem:

$$P(A|B) = \frac{P(B|A) * P(A)}{\text{sum over } C\ (P(B|C) * P(C))}$$

This Theorem can be proved in a variety of ways; perhaps the simplest is with a Venn diagram.

When used in Bayesian inference, this Theorem is interpreted as giving the probability of hypothesis A given observation B, based upon the "prior" probability of hypothesis A and alternative hypotheses C.

Even though Bayesian inference seems obscure, it is fundamental to correct decision making. Consider this example:

A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. Here is some data:

a) Although the two companies are equal in size, 85% of cab accidents in the city involve Green cabs and 15% involve Blue cabs.

b) A witness identified the cab in this particular accident as Blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colors 80% of the time and failed 20% of the time.

What is the probability that the cab involved in the accident was Blue rather than Green?

The probability that the color of the cab was Blue is 80%! After all, the witness is correct 80% of the time, and this time he said it was Blue!

What else need be considered? Nothing, right?

If we look at Bayes' Theorem we get a much lower probability:

P(was Blue | saw Blue) =

$$\frac{P(\text{saw Blue} \mid \text{was Blue}) * P(\text{was Blue})}{P(\text{saw Blue} \mid \text{was Blue}) * P(\text{was Blue}) + P(\text{saw Blue} \mid \text{was Green}) * P(\text{was Green})}$$

$= .8 * .15 / (.8 * .15 + .2 * .85) = 41\%$

But why should we consider statistical theorems when the problem appears so clear cut? Should we just accept the 80% figure as correct?

Well, suppose the observer was only correct 50% of the time. Then the odds that the observer incorrectly reported a Green cab as a Blue cab become much larger than the odds that the observer incorrectly reported a Blue cab as a Green cab. The asymmetry in the observer's errors is so large that our intuition demands that we reject the naive 50% figure and make some corrections for the asymmetric errors. And Bayes' Theorem is just the mathematics we need to make the correction.

When the biases get so enormous, things start getting quite a bit more in line with intuition. However, note that if the biases were comparable (i.e., the number of Green and Blue cabs roughly equal), then the results of intuition and Bayesian inference would be very close.

So Bayesian inference is fundamental to correct decision making. The problem is that Bayesian inference, relying as it does on Bayes' Theorem, is highly mathematical and therefore difficult for our intuition to handle. Consider the following:

A contestant on "Let's Make a Deal" can choose one of three doors. A prize is behind one of the doors. Suppose the contestant chooses door number 1. Monty Hall now selects one of the remaining doors (door 2, say), opens it and reveals that it is empty. Monty then offers the contestant the chance to switch to the remaining door number 3. Is there an advantage to switching?

The obvious answer is no. After all, what is the difference between the door that we initially chose and the door Monty did not choose? If Monty randomly selected between the two doors, this analysis would be correct.

$P(\text{door 3} \mid \text{not door 2}) =$

$$\frac{P(\text{not door 2} \mid \text{door 3}) * P(\text{door 3})}{\text{sum over i } (P(\text{not door 2} \mid \text{door i}) * P(\text{door i}))} =$$

1 * 1/3 / (1 * 1/3 + 0 * 1/3 + 1 * 1/3) = 1/2.

However, Monty knows which door contains the prize and always reveals a door that is empty. Therefore, Monty's action has conveyed some information, changed the prior probabilities and we should switch. Formally:

P(door 2) = 0 and P(door 3) = 2/3, so P(door 3 | not door 2) =

1 * 2/3 / (1 * 1/3 + 0 * 0 + 1 * 2/3) = 2/3.

The "Monty Hall" puzzle regularly confounds even the most educated person. In a recent rehashing of the problem on Usenet, dozens of mathematicians argued the subject before sanity was restored. Not surprisingly, statisticians came to the rescue.

Consider another familiar application of Bayes' Theorem:

An urn contains ten balls. Ten balls are sampled from the urn and replaced. All ten are white. What are the odds that the urn contains ten white balls? Formally:

P(10 white | 10 white selected) =

P(10 white selected | 10 white) * P(10 white)
----
sum over i (P(10 white selected | i white) * P(i white))

In order to compute this, we need to know P(i white), the prior probabilities that there are i white balls. Not being told how the urn was prepared, we cannot know this. Therefore, we cannot compute the odds that the urn contains ten white balls.

However, suppose we ask a different question: How many times must I sample this urn to be 95% sure that it only contains white balls? Well, the worst case would be if it contained only one black ball. My odds of not picking that ball on each sample are 90%. So how big must n be for .9 ^ n <= .05? The answer is n = 29.

This is "sample theory", and it seems like one way out of Bayesian regression. But more on this below.

Finally, we come to that thorniest of thickets: Newcomb's Problem.

You are presented with two boxes: one certainly contains $1000 and the other might contain $1 million. You can either take one box or both. You cannot change what is in the boxes. Therefore, to maximize your gain you should take both boxes.

The following argument might be made: your expected gain if you take both boxes is (nearly) $1000, whereas your expected gain if you take one box is (nearly) $1 million, therefore you should take one box. However, this argument is fallacious. In order to compute the expected gain, one would use the formulas:

E(take one) =

$0 * P(predict take both | take one) +
$1,000,000 * P(predict take one | take one)

E(take both) =

$1,000 * P(predict take both | take both) +
$1,001,000 * P(predict take one | take both)

While you are given that P(do X | predict X) is high, it is not given that P(predict X | do X) is high. Can we use Bayes' Theorem to compute it? Let's see:

P(predict X | do X) =

P(do X | predict X) * P(predict X)
----
sum over i (P(do X | predict i) * P(predict i))

But what is P(predict i)? We do not know. Therefore, like the urn problem, we cannot compute P(predict X | do X). Indeed, specifying that P(predict X | do X) is high would be equivalent to specifying that the being could use magic (or reverse causality) to fill the boxes. Therefore, the expected gain from either action cannot be determined from the information given.

This is good, since it removes a conflict between common sense (take both boxes) and expected value theory.

And finally, I finish with some speculations. Bayesian regression, the impossibility of calculating probabilities without assuming probabilities, has plagued philosophers for all of recorded history. No modern device, such as sample theory, is likely to remove the problem. Sample theory and other robust statistical methods do allow a measure of independence from our assumptions, but they do not remove them entirely. The necessity of making assumptions simply cannot be removed. The paradigm of the brain in a vat demonstrates that.

However, few of us worry about being brains in a vat. If that is all Bayesian regression threatens us with, I am not impressed. Unfortunately, I fear that the problem is much more profound than that. If we know that many of our judgements are based upon hidden assumptions, perhaps we will not be so quick to criticize ideas we do not understand. Indeed, our most basic assumptions may be a result of our circuitry.

Choose between two alternatives:

A. 89% chance of an unknown amount
10% chance of $1 million
1% chance of $1 million

B. 89% chance of an unknown amount (the same amount as in A)
10% chance of $2.5 million
1% chance of nothing

What is the rational choice? Does this choice remain the same if the unknown amount is $1 million? If it is nothing?

This is "Allais' Paradox".

Which choice is rational depends upon the subjective value of money. Many people are risk averse, and prefer the better chance of $1 million of option A. This choice is firm when the unknown amount is $1 million, but seems to waver as the amount falls to nothing. In the latter case, the risk averse person favors B because there is not much difference between 10% and 11%, but there is a big difference between $1 million and $2.5 million.

Thus the choice between A and B depends upon the unknown amount, even though it is the same unknown amount independent of the choice. This violates the "independence axiom" that rational choice between two alternatives should depend only upon how those two alternatives differ.

However, if the amounts involved in the problem are reduced to tens of dollars instead of millions of dollars, people's behavior tends to fall back in line with the axioms of rational choice. People tend to choose option B regardless of the unknown amount. Perhaps when presented with such huge numbers, people begin to calculate qualitatively. For example, if the unknown amount is $1 million the options are:

A. a fortune, guaranteed

B. a fortune, almost guaranteed

a tiny chance of nothing

Then the choice of A is rational. However, if the unknown amount is nothing, the options are:

A. small chance of a fortune ($1 million)
large chance of nothing

B. small chance of a larger fortune ($2.5 million)
large chance of nothing

In this case, the choice of B is rational. The Allais' Paradox then results from the limited ability to rationally calculate with such unusual quantities. The brain is not a calculator and rational calculations may rely on things like training, experience, and analogy, none of which would be help in this case. This hypothesis could be tested by studying the correlation between paradoxical behavior and "unusualness" of the amounts involved.

If this explanation is correct, then the Paradox amounts to little more than the observation that the brain is an imperfect rational engine. The brain appears to be a large neural network. Neural networks are good at pattern matching. Could it be that "rationality" is being simulated by our brains, with imperfect results? Perhaps the only kind of reasoning that is built in is reasoning by analogy, a notably unreliable method. Perhaps we only recognize syllogisms when they are like syllogisms we have seen before.

Add to this the disturbing possibility that our brains are partially hard-wired, such as in our visual cortex, and perhaps the degree of certainty that we attach to our beliefs deserves a second look. For as we learn from Bayesian regression, our judgements are no more certain that our assumptions. The main lessons are, I think, skepticism and humility.

Book Review

Labyrinths of Reason
by William Poundstone

William Poundstone aspires to the ranks of authors such as Martin Gardner, Douglas Hofstadter, Raymond Smullyan, Isaac Asimov, etc. These authors popularize the recent esoteric results of computability theory, formal logic, philosophy of science, etc. Surprisingly, there is a market for such books, even though most of the results seem barely accessible to the non-specialist. I fear much is lost in the translation.

Labyrinths of Reason is Poundstone's fourth book. His first two, *Big Secrets* and *Bigger Secrets* are amusing exposes of supposedly secret but in reality just partially-hidden-and-embarrassing facts like Masonic rituals, government radio frequencies, etc. The real secret stuff, like the formula for Coke or Kentucky Fried Chicken, is not revealed but only guessed at. The third book, *The Recursive Universe*, nobly attempts to sort out the issues of energy and information, but fails. In the end, you may conclude that somehow energy and information are the same.

In *Labyrinths*, Poundstone includes brains in vats, confirmation theory, satisfiability and NP-completeness, Searle's Chinese room, the Voynich manuscript, a Holmes pastiche, and the paradoxes of Hempel's raven, Goodman's grue-bleen, Poincare's doubling, Theseus' ship, Ekborn's hanging, Hollis' number, Buridan's sentence, Gettier's counter- examples, Zeno's arrow, Olber's sky, and Newcomb's demon.

For covering a lot of ground and providing a readable introduction to a lot of material, this book is worth the effort. If you do not know what all of the above topics are (and I sure did not), this book provides a quick overview. For deep understanding, however, see the original papers. It is clear that Poundstone does not really understand several of the topics he is discussing. In particular, the paradox of the unexpected hanging is resolved much better by O'Beirne or Quine, and Poundstone really misses Newcomb's problem. I think he suffers from reading the philosophy journals he cites at the top of the bibliography. Like all publish-or-perish rags, these journals are 99% gibberish. Sustained reading is guaranteed to addle anyone's brains.

I will end the review with a nice puzzle from Poundstone's Holmes pastiche:

A man receives an anonymous note requesting a meeting at midnight in the local graveyard. The man is terrified but his curiosity gets the better of him.

It is a still night with no wind and dark. The only light is from a thin crescent moon. The man waits in front of the family crypt. In the morning, the caretaker finds him dead with a hideous grin on his face.

Did the man vote for Teddy Roosevelt in the 1904 U.S. presidential election?

Solutions

Physics

The dog's position and direction are indeterminate, other than that the dog must be between the boy and girl (endpoints included). To see this, simply time reverse the problem. No matter where the dog starts out, the three of them wind up together in one hour.

This argument is not quite adequate. It is possible to construct problems where the orientation changes an infinite number of times initially, but for which there can be a definite result. This would be the case if the positions at time t are uniformly continuous in the positions at time s, s small.

But suppose that at time a the dog is with the girl. Then the boy is at 4a, and the time it takes the dog to reach the boy is a/6, because the relative speed is 6 mph. So the time b at which the dog reaches the boy is proportional to a. A similar argument shows that the time the dog next reaches the girl is b + b/13, and is hence proportional to b. This makes the position of the dog at time (t > a) a periodic function of the logarithm of a, and thus does not approach a limit as a -> 0.

Logic

B must have one red and one green stamp. The person who answered YES must have a solution that is symmetrical between the two colors, since the setup has this symmetry.

Mathematics

Take a person X. Of the five other people, there must either be at least three acquaintances of X or at least three strangers of X. Assume wlog that X has three strangers A,B,C. Unless A,B,C is the required triad of acquaintances, they must include a pair of strangers, wlog A,B. Then X,A,B is the required triad of strangers, QED.

English

The following is the list of acronyms which have become common words. An acronym is "a word formed from the initial letter or letters of each of the successive parts or major parts of a compound term" (Webster's Ninth). A common word will occur uncapitalized in Webster's Ninth. Let me know if you find others.

fido	Freaks + Irregulars + Defects + Oddities (1966)
gox	Gaseous OXygen (1959)
laser	Light Amplification by Stimulated Emission of Radiation (1957)
lox	Liquid OXygen (1923)
maser	Microwave Amplification by Stimulated Emission of Radiation (1955)
modem	MOdulator + DEModulator (ca. 1952)
napalm	NAphthenic and PALMitic acids (1942)
radar	RAdio Detection And Ranging (ca. 1941)
scuba	Self-Contained Underwater Breathing Apparatus (1952)
snafu	Situation Normal -- All Fouled Up (ca. 1940)
sonar	SOund NAvigation Ranging (1945)
zip	Zone Improvement Plan (1963)

Acronyms from other languages:

flak	FLiegerAbwehrKanonen (German, 1938)
gestapo	GEheime STAatsPOlizei (German, 1934)
gulag	Glavnoe Upravlenie ispravitel'notrudovykh LAGerei (Russian, 1974)
tokamak	toroidal magnetic chamber (Russian, 1965)

Hyphenated acronyms:

hi-fi	HIgh FIdelity (1948)

Trivia

One is in Eastern Oregon (in Mountain time), the other in Western Florida (in Central time), and it's daylight-savings changeover day at 1:30 AM.