Wellcome Film Project

The Logical Analysis of Clinical Medicine, Part Two

The Scientific Basis of Medicine

Presented by Professor Wilfred Card, University of Glasgow.

University of London Audio-Visual Centre, 1975.

Made for British Postgraduate Medical Federation.

Road marking photographs produced by permission of the Transport and Road Research Laboratory.

Computer animation by Robert Colvill, University of London Computer Centre.

Graphic design by Janice Minton and Hank Young.

Produced by David Sharp.

Black-and-white

Duration: 00:37:38:07

00:00:00:00

<**Opening titles**>

< **Card to camera**>

In the previous talk I discussed a simple model in this analysis of clinical medicine and the first thing we had to do was to try to specify a patient as a set of elements of evidence and we call them indicants.

<**Card narrates over animated mathematical model**>

We then have to try and define disease classes and allocate the patient to a disease class by some process of inference. For a given disease there is then a choice of treatments and each treatment has a choice of outcomes. We've got to attach some measure of worth or value to the outcome and we call this utility. The analysis I want to discuss today is the treatment and this part of the analysis has been much less worked on.

First we have to define a set of treatments and the outcome of treatments have also to be defined.

<**Card to camera**>

This can be very clear cut as in the two classes survival and death following mild cardio-infarction, or it can be less sharply defined.

<**Card narrates over table**>

Here is the set of outcomes used in a study of severe brain damage and they have to be exhaustive, they have to include everything and they have to be mutually exclusive – a patient can't be in more than one class.

<**Card narrates over mathematical diagrams**>

If we are to formalise our decisions, it is insufficient to say if one outcome is preferable to another, we have to assign this number, this utility, and we shall discuss later how this might be estimated. When we treat a given patient, we only have a certain chance or probability of achieving a given outcome. So that we have to add probabilities to these outcomes as you see here. We shall also have to add costs to the model, the costs of treatment, and these will have to include not only monetary costs but also any pain or danger incurred with the patient, so cost is a rather complicated concept.

How, now, do we measure the value of different treatments? So that we can choose between them. The expected value, Treatment 1, is p1U1 + p2U2 less c1. And expected value of Treatment 2 is (p3U3 +p4U4-c2) – we choose that treatment which has the greatest expected utility.

<**Card to camera**>

Now this is the essence of the problem of formalising treatment decisions. Obviously it can be made very much more complicated. We can have many more treatments, many more outcomes per treatment; we can also imagine a sequence of treatments. When we say treatment for a duodenal ulcer is either medical or surgical, we really mean medical now with the opportunity of surgery later.

<**Card narrates over animated mathematical diagram**>

On this illustration, S is surgical, M medical; we can have surgical or medical now and surgery later, or medical now, and then medical and then surgery later or medical, medical, medical.

00:05:05:23

<**Card to camera**>

In the model, the treatment is depicted as being based on a particular disease class, but we used to be taught to treat always the patient and never the disease.

<**Card narrates over animated mathematical table**>

So perhaps we ought to go back to the patient as described by a set of indicants and calculate our treatment on that. This is possible, as we shall see. But the essence of the whole problem is the estimation of the utilities of outcomes, the estimation of the probabilities, given the disease or given the state of the patient, and the estimation of costs.

Prognosis is a set of probabilities of outcomes, given the disease and given the treatment. But we have put the patient here into a disease class and it may be preferable to think of the prognosis given the state of the patient. In the symbolism here, the X in heavy type signifies the state of the patient – the vector X1, X2, all the way to Xn. How is this set of probabilities to be estimated?

<**Card to camera**>

Ideally the statisticians might like us to carry out an informative experiment. This would mean taking a series of patients, say 100, and giving each one the set of treatments at random and then observing the results in terms of defined outcomes. While such a formal therapeutic trial is sometimes possible, for example with a new drug, frequently it's impossible for ethical reasons. Where a doctor is concerned with the outcome of the patient given a particular treatment, for example in the case of myocardial infarction, he makes the prognosis in terms of a series of observations which he may weigh mentally with different degrees of seriousness. This kind of method can be imitated mathematically.

<**Card narrates over animated tables**>

First there are two classes in myocardial infarction of death and survival. He then chooses a set of indicants which he thinks may be important, and to each of these he attaches a weight. There is then the variable – the x1, x2, x3 as the value for a particular patient and this will total to a score X and X prime for the two patients and we imagine a population of patients whose outcomes we know and whose variants we know and we suppose the distribution of these scores is normal.

Here we see the distributions of the scores of those patients who died from myocardial infarction, and here we see the distributions of the scores of those patients who survived. There's an overlap between them and this will be the area of misclassification. What we should like to do is calculate the weights, the As, such that the distributions are separated as widely as possible.

These two distributions are also shown here. What we could do is to calculate the weights so that the means of the distributions of the two scores are separated as widely as possible. M and M-. Or we could else tighten the distributions by reducing their variants. Or else we could do both. In practice, we calculate the weights so as to maximise the ratio of the difference between the two means and the variants of the distribution.

00:10:24:00

<**Card narrates over table**>

Here's an example from an actual study. On the left you'll see the prediction, 299 were predicted to have survived, of whom, in fact, 9 died. Similarly, the prediction of death in which 118 people died but 28 people survived. The misclassification was that the crossover of the distributions is 37, that is 8.3%. In such a study, a large number of characters and variants are used, but it looks as though this number might be considerably reduced.

<**Card narrates over graph**>

In a study of the effects of severe head injury in which the authors were interested in the probability of epilepsy occurring a year later, the 8 possible classes of these 3 factors were plotted here and you will see how they relate to the probability of epilepsy.

<**Card to camera**>

Sometimes it won't be possible to carry out experiments such as therapeutic trials or analyse the natural course of the disease and we may be forced to make use of the subjective probabilities of the experts.

To attach a number or utility to a given state of health is a formidable problem which is, as yet, unsolved. First, whose to make an estimate of the state of health of the patient? At first sight it should be the patient himself seeing as he's the only one to experience the effects of his illness, but frequently he lacks much of the information and it looks as though the estimates will have to be made by a doctor who includes the attitudes of the patient and his family in his estimation. The only general principle to guide us is to look at the utilities as concepts which underlie the consistency of action of doctors. We can then only infer them by observation of the ways doctors behave either in real or simulated situations. We can get hold of a number or a utility by adopting a wagering technique which is derived from the axioms of decision theory. Here's an example.

Supposing you've gone blind and there's an operation which could save your sight. Unfortunately it's a dangerous operation and carries a considerable mortality. Your choice is between staying permanently blind and risking the operation, the outcomes of which are perfect vision or death. What is the highest mortality you would be prepared to accept? I asked this question of a group of neurosurgeons and the average highest mortality was 18%. We can now calculate utility of complete blindness on this assumption.

<**Card narrates over mathematical tables**>

UB is the utility of the state of complete blindness, UD the utility of the state of death – we call this 0, perfect vision / perfect health is 1 and P is the probability of death. In the utility of the present state of complete blindness is the probability of death and 1-P, the probability of perfect vision. That reduces the 1-P, since UD is 0, P in the neurosurgeons estimation was .8 and that makes the utility of the state of blindness .82 on this scale of 0-1.

Now, a single wager like this is insufficient, since we assume there is a certain attitude, technically utility function which runs through a set of decisions.

<**Card to camera**>

We have therefore experimented with a set of states of health in each of which is a loss of visual acuity. We imagine a concrete situation in a given patient.

00:15:17:00

<**Card narrates over tables**>

Here is a man age 38, married, he is an area manager, has to drive a car, salary £4500 a year, he's got a number of hobbies and he encounters a car accident. He completely loses the sight of one eye and there's a diffuse corneal opacity in the other eye and one can imagine various states of health corresponding to this reduced vision. Here is the set of states of health.

A, he has no vision at all; B, he can count fingers; C 1/60th, D 3/60th and so on up to H in which he's got perfect distant and near vision. From this set of states of health, we can now ask a subject to make a set of decisions and from these calculate his implicit utilities. We do this, 3 at a time, and here is an example: the patient is in state F and we imagine an operation which might improve him to state H with perfect vision. If it is unsuccessful, he will go down to state E with worse vision. We give him a probability of p, success, and ask him what his maximum q, that's the probability of failure, he's prepared to accept. The difference is the probability of remaining in the same state. We can get a whole series of wagers in this way, from them calculate his implicit utilities and plot these against a logarithm of visual acuity.

As you will see, we've got a linear utility function and we attain this in all our patients though at different slopes, which suggests that each subject was broadly consistent in his decisions. We can compare the utility functions of different observers.

<**Card to camera**>

Now these results are encouraging but the situations in real life are not uni-dimensional as they are with visual acuity, but multi-dimensional. And if you feel these are too complex to be soluble we should remind ourselves we solve such decision problems every day of our lives. The choice of a mode of travel, for instance, depends on considerations of speed, cost and comfort. We tried to get hold of numerical estimates of states of health by asking a number of gastroenterologists how they would compare in a given patient a perfectly working ileostomy say to the loss of an index finger, the loss of the thumb; a comparison of disabilities. These losses carry a fixed loss of disability under the industrial injuries act and in this way we could get a set of scores for each doctor and compare them.

<**Card narrates over mathematical table**>

And this shows the co-efficient of correlation on the right of each doctor with the average of the set of his colleagues: D6 – D10. And you will see that the measure of agreement is encouraging.

But if we plot the doctors' mean scores, you will see they show quite a shift and difference in their variants. But we can in this way get hold of a set of cardinal numbers but these are not true utilities and they will need conversion into a true utility scale using some wagering technique involving, as before, the fixed points of perfect health and death

<**Card to camera**>

Now I think our feeling is that while an estimation of the states of health remains a very formidable problem, it is perhaps not quite so impossible as it first appears.

00:20:11:02

It is indeed possible to construct a worthwhile decision tree simply using the utilities of 0 from 1, 0 for death and 1 for perfect health.

<**Card narrates over complex mathematical diagrams**>

In this tree here, the problem is whether to give heparin to a patient with myocardial infarction and this tree was constructed prospectively. If the patient's age is over 70, heparin is given. If the age is between 50-69 the next question is whether they've got varicose veins or had a previous pulmonary thrombotic episode. If the answer's yes, heparin is given, if it's no, you ask them if they're a smoker and so on. And such trees which are quite simple may prove to be quite useful and quite powerful.

If in a given decision system we had complete knowledge of all the utilities and probabilities, we could calculate the best treatment decision provided we ignored costs. By costs we mean to include not only monetary but also biological costs, that is pain, danger to life, etc. If we could include costs we could answer the question when faced with an expensive investigation – is it worth it? But this will mean finding an equivalent to utilities in terms of money. That is, we are forced to assign a monetary value to a state of health of human life.

<**Card to camera**>

At first sight this suggestion may repel you but I think this is due to a misunderstanding. All we are saying is that the resources of the health service are inadequate for all the demands made on it and have to be deployed most economically. This is fully recognised by the government.

<**Unnamed narrator quotes from Dr. David Owen, Minister of State for Health, The Times, 6 Feb. 1975**>

The Health Service is a rationed service, there will never be a government or country that has enough resources to meet all the demands any nation will make on a National Health Service.

<**Card to camera**>

If we are to deploy rationed resources most economically, no-one has suggested any better way than through the use of decision theory. The problem is not confined to the Health Service but is general. The Ministry of Environment is faced with this problem if it has a road junction at which a number of accidents take place and it has to decide whether to improve the safety of the junction.

<**Card narrates over still images of traffic**>

This photograph shows a method that has been employed of drawing lines across the road, exponentially decreasing at various intervals and this forces motorists to slow down. There is therefore a general problem of what is the monetary value of life at different ages.

<**Card to camera**>

The economists have calculated the average value of life, thinking of man as a consuming and producing unit, but also allowing evaluation for his existence as a human being.

<**Card narrates over flow chart**>

Vice Brod was one of the earliest to do this, this is before inflation, and here are two curves he calculated using different percentage rates of discount . The road research laboratory has also done some calculations on the net loss of output due to death, again, before inflation, and these two curves have the same function but they are both unimodal. One hump[?] where the maximum age is about 25-30 – this is the economic argument about the value of human life.

<**Card to camera**>

But to the health service we might see what it is prepared to spend to save a life. If a patient has renal failure, what does it cost to provide home dialysis?

<**Card narrates over table**>

We looked at the figures in the unit and the initial cost was over £5000 per patient, and the running costs were £4500. It was then possible, knowing the probability of living from year to year, taking certain discount rates to calculate the present cost of putting a patient on home dialysis. And the highest figure was over £50,000. While no-one would want to comment on this cost, it would be interesting to compare it with the costs of saving other lives to see how far there is any general coherency and in this kind of way we may one day be able to assign a notional valuation to a human life. And if this is possible, one can then calculate whether an expensive investigation is worthwhile or not.

00:26:05:00

<**Card narrates over mathematical diagram**>

This illustration shows the formalism of such a measurement. The top part shows the decision tree associated with treatment as we've already seen. At the lower part we can either do a further test and get E evidence or fail and get E bar, no evidence. If we get no evidence, the probabilities will of course remain the same p2, p3, p4 etc. If we get evidence, the probabilities containing the outcomes of health will have altered: q1, q2, q3, q4. And if we knew the values of all these parameters it would be possible to go back to the beginning and calculate whether the test should be done or not.

<**Card to camera**>

While we have focused in these talks on the logical analysis, that is the background to this thinking, it would not be possible to carry out much of this, in practise without a computer.

When a new machine is first introduced people say that it can never replace a human being, then the pendulum swings and they say the whole thing can be done by a machine, and finally the stage which I hope we have reached where we can consider the role of man and the machine and how we can use the most useful qualities of both.

What are the useful qualities of the computer? Given certain assumptions and certain data, it can calculate probabilities and human beings aren't always very good at this. Try this problem.

<**Card narrates over demonstration featuring two glass jars of beans**>

Here we have two glass jars, one containing predominantly black beans in the ratio of 70:30 and the other containing white beans in the ratio 70:30. We're going to put these beans into two bags. The beans will then be mixed up and one of the bags will be chosen, you won't know which. Now, we make a drawing of 12 beans from this bag, black, replacing it each time. Black, white, black, white, white, black, black, black, black, white, black.

<**Card to camera**>

8 black and 4 white – which was the bag which was chosen?

00:31:17:00

The black bag. Quite right and what were the odds that this was chosen? 26-1 on. And almost certainly, the great majority of you who don't know this probability distribution would very much have underestimated these odds. Human beings seem bad at making accurate calculations of this kind. They refuse to accept evidence which is staring them in the face. Human beings also get confused by too much evidence, in fact their performance may get worse, whereas it usually makes no difference to the performance of a machine. Human beings are very good at pattern recognition and we have seen something of this already and doctors will always be good at picking up a lot of little clues from watching a patient and coming to a judgment and a machine won't do this for a very long while and some people will say never.

But if we are going to use computers, how much computing power do we need and of what kind? This is something we don't know. We know that very massive computing power is necessary in developing a method and it's very easy to get into very big figures from quite simple problems.

<**Card narrates over table**>

If you're investigating some area of disease, say with 100 indicants, you might like to know which would be the 5 most powerful. You've got more than 75,000,000 combinations to test on your data. Supposing you only wanted combinations of 4 – nearly 4,000,000. Even 3 would give you over 161,000.

<**Card to camera**>

So there's little doubt that in developing the method or on the calculation of the tree, extensive computing power would be needed. But what is not clear is how much the ordinary doctor will need once a method is finally developed.

<**Card narrates over moving image of a calculator**>

And perhaps the programmable calculators contain most of the calculating power we need, though it may not be in the right form and they haven't got enough storage. But if we do use a formal approach to medicine and someone else will do the calculations – shall we need any computing power at all?

<**Card narrates over mathematical diagram**>

Here is a very simple decision tree which has been calculated from a great deal of data to decide whether a jaundiced patient should be classed as medical or surgical. Despite its simplicity it's extremely powerful so that it looks as if simple decision might be put in this kind of way. This simplified tree illustrates a point that comes out of this work – the possibility of what one may call test reduction.

<**Card to camera**>

If we develop this formalisation of medicine that I've outlined – what effect would it have on the practise of medicine? Formalisation looks at medicine as getting more and more information at increasing cost. How might this be organised?

<**Card narrates over diagrammatic chart**>

The patient is the first person who deals with himself and about 60% of human ills are dealt with by self-medication. The next person in the chain I think is going to be an auxiliary or paramedical staff, 3 will be the general practitioner, we shall probably still need a general consultant 4, and 5 will be the specialist.

<**Card to camera**>

It's very difficult to believe that the kind of methods I've outlined won't find a useful place in developing countries. And our general feeling perhaps is that we may return to the words of Leibniz that there will be no more need of dispute between two doctors than between two accountants. It will be sufficient for them to call for their slates, take their pencils in their hands and say, "Let us calculate."

<**End credits**>