Unthinkable Syndromes.
Paradoxa of Relevance and Constraints on Diagnostic Categories

Arthur Merin, Konstanz

1. Genesis of Diagnostic Categories

Bodies of collective knowledge evolve through individual action, like all products that have a use. They also can be evaluated from the engineer's optimizing design perspective. But can individual participants in their making recognize local optimality? Can they work to realize it? Are they unable to act seriosly in a way that would ensure acquisition of a certain suboptimal design feature?

One might hope for a simple answer: appeal to innate constraints on the form of categorization. But such constraints cannot wholly pre-empt the need for individual checks and thus for agency. The use of categorial schemes engages facts about the world that do not leap to the eye. They are unlike syntactic or phonological constraints of language, or edge-and orientation-detection algorithms of visual perception.

What we should thus be looking for will be modest: constraints, innate or otherwise, that rule out the choice of categorial schemes which are clearly suboptimal and that are, in principle, open to agents' instant observation. We look to a limited if important domain of systematic knowledge: medical diagnostic knowledge.

2. Basic Bayesian Reasonings

Signs or symptoms are not usually certain signs of diseases. A dry cough or hypertension could each be associated with many diverse diseases. Observing one or several symptoms will perhaps make a certain disease more likely, but it need not be a conclusive sign of it.

Bayesian reasoning with probability constraints is a well-explored framework for dealing with such inconclusive evidential relations. Its laws specify how the impact of evidence E is to transform prior probabilities P(H) of a hypothesis H or its odds P(H)/P(-H) into posterior probabilities P(H|E) or odds P(H|E)/P(-H|E). We obtain posterior odds by multiplying prior odds with the Bayes-Factor, P(E|H)/P(E|-H). Probability constraints represent a belief- or (possibly partial) knowledge-state. Bayesian networks are sets of probability constraints that represent probabilistic dependencies between variables, in our example between indicator variables X[K] whose values represent the truth of complementary propositions, K and –K, respectively. If there is some pair (H1,E1) of values of the variables X[H] and X[E], such that P(H1|E1) > P(H1), then the variables are said to be probabilistically dependent in a doxastic state P. A special case is deterministic dependence, when P(H1) < P(H1|E1) = 1.

To become computationally tractable, Bayesian nets rely on simplifying assumptions, particularly conditional independence assumptions. Propositions A and B are conditionally independent given proposition H under a probability function P iff P(AB|H) = P(A|H)P(B|H). (Unconditional independence is the special case where H is the tautology.) Independence conditional on a variable X[H] is independence conditional on each of the cells of the propositional bipartition {H,-H} of the space of possibilities. Call this doubly conditional independence (DCI). It is this condition which is crucial to Reichenbach's (1956) explication of a Common Cause (CC).

A necessary condition on H being a CC of A and B is that AB are DCI under H and are both positively relevant to H. (Temporal precedence of H is the remaining definitional requirement. Without this requirement we might simply speak of H as a 'common reason' (CR) for A and for B.)

DCI implies that A and B each positive to H are positively relevant also to one another. This suggests a common procedure. We search for common causes on observing a dependence (positive correlation) of A and B. We then infer as a CC or CR some condition H as yet unknown to us that we next go and search for.

Another and immediately obvious formal consequence of DCI is that the relevance of A and B to H, i.e. a measure of their evidential import for A, becomes compositional when assessed by the Bayes Factor: now P(AB|H)/P(AB|-H) = [P(A|H)/P(A|-H)]x[P(B|H)/P(B|-H)]. Taking logarithms we have the relevance of AB as the sum of relevances of A and of B.

Common Cause and Compositionality, both associated with DCI have given rise to a suggestive and, in its way, Kantian hypothesis. We postulate causes H (causal variables X[H])in order to reap the computational advantages of the condition, DCI, which goes into defining CC. The original context of this suggestion was medical diagnosis.

3. Computation, Causes and Syndromes

A syndrome, says the Shorter Oxford English Dictionary of 1983, is "a group of symptoms or pathological signs which consistently occur together, especially with an (originally) unknown cause; a condition characterized by such a set of associated symptoms." It is in this latter, rather liberal sense of "condition" that the epistemological literature on diagnostic categories appears to use the term "syndrome".

Charniak (1983) conjectured that medical diagnostic categories ("syndromes") are formed so as to afford independence of possible observations ("symptoms") conditional on presence of the syndrome [CI]. Pearl (1988:44) amplifies. Conditional independence on a syndrome variable i.e. DCI, allows probabilistic reasoning on dependencies to be decomposed into a set of independent subtasks in stages, using simple vector operations. Such simplicity, surmises Pearl, makes conditional independence a psychological necessity (presumably: wherever it can readily be assumed). "It may be to reap the computational advantages associated with such independence", concludes Pearl, "that we organize most of our knowledge in causal hierarchies."

Pearl's thesis on the structure and evolution of diagnostic categorial schemes is a philosophical hypothesis of the empirical kind. But, as stated and motivated, it is not readily testable. Numerical data are not generally available to verify whether the DCI constraint or at least compositionality of evidence holds in many particular cases. Thus

we could not expect a reliable answer if we asked a physician:

> If A and B are symptoms of H, will the indicator random variable X[H] fully account for the dependence between them?

Indeed, even if we made allowances for erroneous assumptions of conditional independence (in line with a ubiquitous tendency to assume independence where, in fact, it fails to hold), we should most likely draw a blank. This is not a question many physicians will be able to answer decisively. The same will be true for less fastidious and perhaps more intuitable, but reasonably unanswerable questions such as

> If A and B are symptoms of H, will this ensure that A is not a cause of B, nor B a cause of A, yet each is a sign that the other will occur?

A more careful variant, again, is still problematic:

> If A and B are symptoms of H, and if you either know for sure that H or know for sure that -H, will the presence of A or of B neither decrease nor increase your surprise at observing the other?

The condition where H is known might seem palatable as the basis for an exercise of the medical imagination. But the condition -H seems hard to fathom. (At least part of the problem seems to stem from psycholinguistic problems. But we have to contend with them.)

So what is wanted is a scheme of scanning diagnostic categories for possible counterexamples to Pearl's or even Charniak's thesis. Also wanted is a criterion that might guide individual agents of evolution. Syndrome categories are, after all, introduced or at least proposed by individual members of the medical community. And then we must ask: How can individual proponents of a taxonomic decision avoid suboptimal acts of category formation when they lack, as they must, the hindsight of the evolutionary long run which lies ahead of them?

4. Paradoxa of Relevance

The probability calculus admits well-known paradoxa of relevance (Carnap 1950: Ch.6 gives an exhaustive and exhausting discussion). Probability models can be constructed which satisfy one or the other of the following conditions:

> A and B are positive to H, but AB is negative to H.

> A and B are positive to H, but AvB is negative to H.

> AB is positive to H, but AvB is negative to H.

Paradox (i) has non-medical examples readily constructed to model it. Paradox (ii) is less readily imagined on the spot, but is well known as Simpson's Paradox. By De Morgan's equivalence and recalling that negation reverses relevance sign, we find that (i) is the condition -A and -B negative to H, -Av-B positive to H. So, formally, any instance of (i) and (ii) is also an instance of the other modulo sign. Predicament (iii) holds whenever (i) or (ii) holds and will be of no substantive further interest to us. (Its necessary and sufficient conditions allow one of A and B to have zero relevance to H.)

What makes (i) and (ii) counterintuitive? A first answer is that their analogues cannot arise in deductive, i.e conclusive reasoning. A proposition E is negative to a proposition H iff positive to -H. Now recall (i)-(iii) and replace in them "positive to" by "entails", and "negative to" by "entails the contradictory of". None of the conditions thus obtained is satisfiable. If both A and B entail H, then so do AB and AvB. If AB entails H, AvB cannot entail -H. The same will hold analogously when probability is used to explicate a notion of non-vacuous doxastic entailment, $P(H|E) = 1 > P(H)$, of extreme positive relevance of some E to H. Relevance configurations which cannot be retained smoothly when relevance goes to the extreme are presumably unnatural. But there are also prima facie different considerations involved.

5. Doubly Conditional Independence blocks Paradox

Neither of the paradoxa can arise under DCI.

> FACT 1: If A and B are DCI on {H,-H} and both are positive to H, then AB and AvB are also positive to H.

The proof is immediate. DCI ensures multiplicativity of Bayes factors. If both A and B are positive to H, their Bayes factors each exceed 1, and hence their product, equal to the Bayes factor of AB, must exceed 1. That Paradox (ii) is ruled out by DCI is less obvious to immediate intuition. But of course it is, as the structural equivalence to (i) already tells us and direct proof easily confirms. Probability models are also readily constructed to demonstrate

> FACT 2: CI is insufficient to rule out evidential paradox.

> FACT 3: Neither DCI nor even CI is necessary for compositionality (log-Bayes factor additivity) of evidential weight.

Fact 2 tells us that entailment by H, which makes for CI is not enough. Fact 3 would be a little worrisome if DCI were our prime target. (Only a little, though. DCI guarantees compositionality, and, since it can be specified elegantly, having DCI is better than hoping for haphazard compositionality). However, what we are after is additivity of evidence.

Since paradox (i) (hence paradox (ii)) is always a counterexample to additivity of evidence, category creators who failed to avoid it would be performing an act dominated with respect to computational optimality of the conceptual scheme to be extended. They would also ensure that their syndrome could not label a Reichenbach Common Cause of their observables.

6. Individual Intuitions and Natural Categories

Pearl's hypothesis is not readily verified, we noted. However, it is readily falsifiable. A diagnostic categorial scheme which admitted paradox of relevance would be a counterexample.

Now, Simpson's Paradox does affect real-life probabilistic reasoning. But its typical configurations (the Berkeley admissions paradox or inverse effects of medical treatment in sub-populations and their aggregate) are not instances of syndrome and symptoms. They do not involve what might be called 'natural' cultural categories.

Thus we should look at natural categories and doxastic agents in our chosen domain. I have informally checked physicians from a variety of specializations for incidence of paradoxical syndrome-symptom configurations. None could come up with a real medical configuration of

syndrome and symptoms that made for evidential paradox under either formulation (i), (ii) or (iii). None could even make up such a constellation for the nonce.

Perhaps the latter failure simply shows that the physicians queried lack imagination, the kind which philosophers are trained to acquire. Yet (the?) two philosophers who have treated of the paradoxa of relevance in a medical setting, Rudolf Carnap and Wesley Salmon, also exhibit a similar and very sensible blind spot of the imagination. Both authors refer to abstract constellations (i), (ii), Carnap also to (iii). Both make up medical (pseudo)-examples to illustrate paradoxical eviedential situations. But the examples constructed, involving three propositional variables, are never instances of two symptoms each speaking for a syndrome, while their conjunction or disjunction speaks against.

Carnap (1950:367), in ostensible illustration of paradox (ii), considers, with uncharacteristic opacity, two putative syndromes ("virus pneumonia" (H), "bacillus pneumonia" (G)) speaking for a hypothesis that does not itself have either symptom or syndrome status. He does not offer two symptoms speaking for a syndrome. Salmon (1975:27) constructs transparent scenarios where evidence E is positive to each of Carnap's H, G, but is negative to HG or HvG. Substantive plausibility apart, this does not, again, illustrate the syndrome/symptom relation to be tested.

Even philosophers have not so far then, it seems, made up examples violating the "Avoid paradox" constraint on possible syndrome-symptoms configurations. Such configurations look like being unimaginable, as soon as our imagination is informed by constraints of subjective but serious possibility for the specific domain in question.

A significantly sized class of possible instances of diagnostic categorial schemes will thus be ruled out by individual categorizer's inability to entertain or unwillingness to admit schemes which make for paradox. These schemes are bound to violate compositionality of symptomatic evidence for a syndrome, and in doing so they violate DCI, which is a sufficient condition for it. Hence we have a useful constraint on formation of possible categorial schemes. One must suspect that our penchant to search for causes -- common causes -- i.e. for an etiology that reveals an H which can be acted on remedially is the real driving force for this: that we are, as Lichtenberg puts it, causational animals ,"Ursachentiere". Yet whichever of CC or Compositionality we alight upon, paradox-freedom is its tangible signpost.

We can interpret the observed constraint in two ways. First, it might be seen as concrete evidence that human minds are simply built to work in the way that Pearl hypothesizes it does. Secondly, it can be seen as an intuitive check for individuals not so built, against proposals, their own or others', of a categorial scheme which is bound to violate a sufficient condition for evidential compositionality.

Literature

Carnap, R. 1950 *Logical foundations of probability*. Chicago: University of Chicago Press.

Charniak, E. 1983 `The bayesian basis of common sense medical diagnosis.´ AAAI-83 (Proc. National Conference on Artificial Intelligence Aug. 22-26, Washington D.C.), 70-73.

Pearl, J. 1988 *Probabilistic reasoning in intelligent systems*. San Mateo CA: Morgan Kaufman.

Reichenbach, H. 1956 *The direction of time*. Berkeley: University of California Press.

Salmon, W.C. 1975 `Confirmation and relevance´. In G. Maxwell and R.M. Anderson, Jr. (eds.) *Minnesota Studies in the Philosophy of Science, Vol. 6, Induction, Probability, and Confirmation*. Minneapolis MT: University of Minnesota, 3-36.