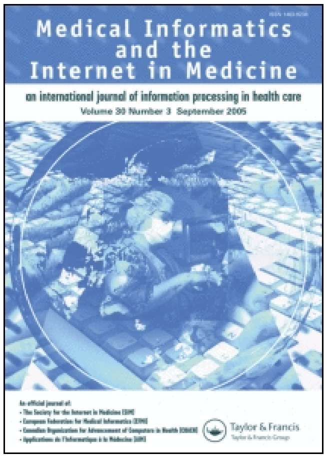

Medical Informatics & The Internet in Medicine
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713736879

Psychological cue use and implications for a clinical decision support system

Christopher Buckingham [a]
[a] Aston University, Aston Triangle, Birmingham, UK.

To cite this Article: Christopher Buckingham , 'Psychological cue use and implications for a clinical decision support system', Medical Informatics & The Internet in Medicine, 27:4, 237 - 251
To link to this article: DOI: 10.1080/1463923031000063342
URL: http://dx.doi.org/10.1080/1463923031000063342

MED. INFORM. (2002) VOL. 27, NO. 4, 237–251

Psychological cue use and implications for a clinical decision support system

CHRISTOPHER D. BUCKINGHAM
Aston University, Aston Triangle, Birmingham, UK

Abstract. Effective clinical decision making depends upon identifying possible outcomes for a patient, selecting relevant cues, and processing the cues to arrive at accurate judgements of each outcome's probability of occurrence. These activities can be considered as classification tasks. This paper describes a new model of psychological classification that explains how people use cues to determine class or outcome likelihoods. It proposes that clinicians respond to conditional probabilities of outcomes given cues and that these probabilities compete with each other for influence on classification. The model explains why people appear to respond to base rates inappropriately, thereby overestimating the occurrence of rare categories, and a clinical example is provided for predicting suicide risk. The model makes an effective representation for expert clinical judgements and its psychological validity enables it to generate explanations in a form that is comprehensible to clinicians. It is a strong candidate for incorporation within a decision support system for mental-health risk assessment, where it can link with statistical and pattern recognition tools applied to a database of patients. The symbiotic combination of empirical evidence and clinical expertise can provide an important web-based resource for risk assessment, including multi-disciplinary education and training.

Keywords: Clinical decision support systems; Classification; Cues; Base-rates; Mental health; Risk assessment.

1. Introduction

The life of a clinician would be greatly simplified if the relationship between data and clinical decisions was deterministic. The role of computer-based clinical decision support systems (DSSs) would be the simple one of recording factual associations between patient information and outcomes. Unfortunately for clinicians, most relationships between cues and patient prospects are probabilistic [1]. Compounding the problem is an incomplete understanding of these relationships, which means statistical models cannot guarantee optimal decisions. Instead, the best approach requires combining empirical data analysis with the less tangible decision-making processes of human clinical experts.

Clinical DSSs range from data organizers that inform decisions to those incorporating human expertise and proffering advice (see [2] for a review). This paper will discuss the latter approach by considering the way in which clinicians process patient cues and how this knowledge can influence the design of DSSs. It will begin by reviewing research evidence on the influences of cues on classification behaviour. Clinical significance will be illustrated using examples from ongoing research [3] into mental-health risk assessment. In particular, people's responses to the prior probabilities or base rates of potential outcome categories will be ex-

*Author for correspondence; e-mail: C.D.Buckingham@aston.ac.uk

Medical Informatics & The Internet in Medicine ISSN 1463-9238 print/ISSN 1464-5238 online
http://www.tandf.co.uk/journals
DOI: 10.1080/14639230310000633342

amined in detail because certain behaviour patterns have been exhibited that different theories (also referred to as 'models') of psychological classification have found difficult to explain. A model which does account for the results has been developed by Buckingham [4, 5]. Its principles will be described, along with its ability to explain base-rate bias and thus allow clinicians to adjust for it. The benefits of using the classification model within a computerized DSS will be proposed, both for reducing clinical errors and in providing an important tool for clinical education and training.

2. Cues and classification

Every patient manifests a set of descriptive attributes, such as eye colour, hair style, age, posture, etc. and more obvious clinical signs and symptoms such as blood pressure, pulse, weight, and temperature. This paper will refer to a specific value of an attribute as a cue. Clinical decisions can be regarded as classification tasks [6] where cues are used to assign patients to one of a number of potential categories. The fundamental importance of classification has been demonstrated in both medical [7–9] and nursing [10–13] domains, and reformulating clinical decision making as classification helps clarify potential causes of bias [14]. It will be discussed with respect to assessing the risk of suicide.

2.1. *Deciding whether a person is at high risk of committing suicide*

One reason mental-health clinicians assess people's risk of committing suicide is to determine whether they should be classified as those requiring intervention, or those who do not. The psychological term for the plausible classes is 'contrast concepts' [15], which equates with differential diagnoses in medicine. Knowledge elicitation activities carried out with a heterogeneous population of mental-health practitioners [3] identified many cues relevant to predicting suicide risk, two of which were the seriousness of intention and the realism of the plan to effect it. The next section considers how such cues influence risk.

2.1.1. Probabilities and cue use. Clinicians who make the most accurate estimates of outcome probabilities given the available evidence will have the best foundation for subsequent interventions. The drawback is that people are often poor at judging probabilities [16–18], one of the most important causes being their apparent failure to use base rates. This would mean that a clinician gives the same probability estimate that a patient commits suicide in a psychiatric hospital whether 99% of patients do or only 1%.

Tversky and Kahneman [19] invoked a 'representativeness' heuristic to explain the failure to use base rates whereby the probability that an object belongs to a class is determined by how representative or how good an example it is of the class. Suppose a psychiatrist wanted to predict the probability of suicide if a patient had a very serious intention (VSI) to commit it. The representativeness heuristic estimates the probability that the patient would come from the suicide class, which is given by $P(VSI\,|\,Suicide)$† and does not take into account the base rate because it assumes the suicide has already happened. The estimate of $P(VSI\,|\,Suicide)$ is

†Conditional probabilities are written with a straight line between two events and can be read as the probability that the left-hand event occurs given that the right-hand event is known to be true.

compared with the equivalent estimate of representativeness from the alternative class, $P(VSI \mid NoSuicide)$, which also ignores the base rate of patients not committing suicide. Hence sizes of the suicide and no-suicide classes do not enter calculations when the representativeness heuristic is used. Base-rate bias is the result because Bayes Theorem shows that $P(Suicide)$ is involved in calculating the desired probability of $P(Suicide \mid VSI)$, and it is dependent upon the relative class sizes:

$$P(Suicide|VSI) = \frac{P(VSI|Suicide)}{P(VSI)} \times P(Suicide) \tag{1}$$

Base-rate bias violates normative probability calculations and leads to over-estimation of rare occurrences. The consequences can be extremely serious: clinicians are known to make errors of judgement which can be interpreted as misuse of base rates [20] and these have been disastrous enough to lead to unnecessary mastectomies, for example [21]. The representativeness heuristic has been identified in clinicians [22] and its use with predicting suicide would be to overestimate the likelihood (suicide being the rarer outcome), with attendant detrimental effects such as unnecessary interventions. Alternatively, suicide risk might be ignored due to previous error rates ('crying wolf' too often) which could lead to patients committing suicide in the absence of an appropriate intervention.

Two psychological theories of how people represent classes have emerged: the prototype one where a class is represented by its single, most typical member [23, 24] and the exemplar one where people remember all previously experienced members of classes [25, 26]. Clinicians using the prototype approach would represent the class of patients at high suicide risk by a single, average or typical patient whereas those using the exemplar representation would remember all patients who were at high risk. The crucial difference is that class sizes are part of the exemplar model's representation, not the prototype one. The representativeness heuristic is really another way of stating that people use prototypes: the degree of representativeness of an unknown patient is determined by comparing its similarity to the most representative or prototypical member of the alternative classes.

It has been shown in general that people use a mixture of exemplar and prototype representations in various forms [27–29] and this is likely to be the case for clinicians [11, 13, 30, 31]. But the different implications for base-rate use between the two approaches has encouraged researchers to investigate base rates in an attempt to shed light on the exact nature of classification decisions. The next section will describe an experiment demonstrating base-rate bias.

2.2. *Base-rate bias*

A series of psychology experiments on base-rate use [24, 32–34] concluded that people display erratic use of base rates, sometimes neglecting them, sometimes responding correctly, and sometimes choosing the rarer category in defiance of base rates (known as the inverse base-rate effect). For clinical decision-making, it is extremely important to be able to predict whether base-rate bias might arise and how to counteract it by, for example, employing artificial decision aids such as computerized DSSs.

Kruschke's [24] first experiment illustrates the nature of the problem. Participants were required to learn associations between sets of hypothetical symptoms and diseases. Table 1 shows the experimental design where the differential diagnoses for symptoms S1, S2, and S3 were two diseases, one common and one rare, the common one occurring three times as often as the rare one. The symptom pattern S1 and S2 always occurred with the common disease and the pattern S1 and S3 always occurred with the rare disease.

After learning the symptom and disease associations, participants were asked to classify symptom combinations not previously encountered. Table 2 gives the proportions of participants choosing different diseases for various symptom combinations†. When all three symptoms were supplied, participants chose the common disease, but when the symptoms attached only to a single disease were put together (S2 and S3), participants favoured the rare disease. This is the inverse base-rate effect; S2 and S3 are equally predictive of their associated diseases and one would expect the common one to be chosen if base rates were being appropriately used. Instead, participants appear to respond to base rates under one circumstance and not another.

A new 'galatean' model of classification [4, 5] proposes a different explanation to existing theories for the inverse base-rate effect. The principles are briefly described in the next section with respect to suicide risk.

3. The galatean model of classification

The galatean model is a type of prototype model but instead of representing the average class member, its prototype encapsulates the hypothetical 'perfect' member—the one with the highest probability of membership (the name 'galatea' comes from Pygmalion's perfect woman). For example, the suicide galatea focuses on cues generating the highest risk of suicide, rather than cues most commonly found in suicide cases, which form traditional prototypes based on the central tendency; galatea cues maximize $P(Category \mid cue)$ instead of $P(cue \mid Category)$. The possible out-

Table 1. Abstract design of training symptoms for Kruschke's [24] Experiment 1: BR=base-rate.

BR	Symptoms	Disease
3	S1,S2	Common
1	S1,S3	Rare

Table 2. Response proportions for Kruschke's (1996) Experiment 1: RP = participants' response proportions (not including error categories)

	Disease RPs	
Symptoms	Common	Rare
S1	0.746	0.174
S2	0.933	0.031
S3	0.040	0.911
S2, S3	0.353	0.612
S1, S2, S3	0.580	0.402

†Proportions do not sum to one because the actual experiment had two isomorphic sets of contrast diseases and error categories are not shown.

come categories, such as suicide and no-suicide, are represented by their own galateas, where a galatea consists of components for each of the attributes relevant to classifying a person into the associated category.

3.1. *Selective attention to cues*

Not all cues have the same influence on classification and the galatean model incorporates different weightings by pairing the cue with the conditional probability of the category given the cue. This probability's role within the model is plausible because clinicians must already be responding to it when learning the perfect values that maximize it. It represents a direct and normative measure of each cue's independent selective attention (the relative weighting, and thus influence, a person gives the cue). The importance of selective attention in classification has been demonstrated both in general [35] and within clinical decision making [10, 36].

In Buckingham and Chan's [3] research, expert mental-health practitioners measured seriousness of intention to commit suicide from a maximum of 10 to a minimum of 0 and the same scale for the realism of a suicide plan. The cue maximizing membership in the suicide category for both attributes is 10 and the one maximizing membership in the non-suicide category is 0, the two categories being inverses of each other. The 'perfect' cue for one galatea will be the worst cue for the alternative galatea and each galatea component thus contains two cue values setting out the range of conditional probabilities, from highest to lowest. Any cues in between these limits are automatically assigned a conditional probability by linear interpolation: a cue half way between the limit and perfect cue would have an estimated conditional probability half way between the conditional probabilities attached to the perfect and limiting cues. The result is an approximate conditional probability distribution for galatea component values: all cues can be assigned selective attentions but the amount of information assimilated by people to do so is reduced to the salient values of the distribution (i.e. perfect and limiting cues in the range).

A second critical phenomenon modelled by galateas is cue competition, where each cue's conditional probability competes with the others within the galatea for expressing predictiveness. Its implementation within the galatean model will be explained in the next section.

3.2. *Cue competition*

Some form of cue competition has often been demonstrated in humans [37, 38] and is clearly an important factor in decision making. In the galatean model, competition occurs simply by limiting the total amount of predictive effect available within each category to 1 and sharing out the relative influences of cues in proportion to their individual predictiveness. A cue's selective attention, SA, becomes its normalized conditional probability:

$$SA_{cue} = \frac{P(Category|cue)}{\sum_{g=1}^{G} P(Category|cue_g)} \tag{2}$$

where G is the number of perfect and worst cues†† in the category galatea. These selective attentions are then used to classify objects.

††More complex probability distributions are possible, whereupon the average conditional probability is used.

3.3. *Galatean classification process*

The galatean model quantifies uncertainty in terms of a person's degree of membership or 'membership grade' [39], within a class. Zadeh [39] created membership grades to describe fuzzy sets containing objects with varying degrees of membership. Classical or 'crisp' sets are a specialization of fuzzy sets where their membership grades can only be 1 or 0 but fuzzy sets can have any membership grades between and including 1 and 0.

The notion of fuzzy membership is pertinent to clinical decision making because it is rare that one knows for certain either what disorder a patient has or what the outcome of that disorder will be. A patient can display evidence for more than one outcome, which can be represented as partial membership of each. The galatean model uses fuzzy memberships as the currency for clinical judgements such that the higher the membership grade for an outcome, the more likely it is considered to happen.

When an object is classified by the galatean model, each object cue is matched with the corresponding galatean cue. The associated selective attention represents the object cue's membership grade in the category. The membership grade for the whole object is the sum of the individual cue selective attentions (produced by equation 2) for those cues which match with the object components:

$$MGu = \frac{\sum_{g=1, g \in Obj}^{G} P(Category|cue_g)}{\sum_{g=1}^{G} P(Category|cue_g)} \tag{3}$$

where MGu is the membership grade before proportioning across the contrast categories, G is the number of perfect and worst cues in the category galatea, and Obj is the set of object cues. The model proposes that people learn the overall strength of evidence for a category, given by the total conditional probability associated with it (the denominator), and judge an object's membership in it by the proportion accumulated from matching components (the numerator). The denominator represents cue competition: the larger it is, the less influence of individual object cues. Because rare categories have a smaller denominator, their cue influences tend to be exaggerated, thereby causing the inverse base-rate effect. The final membership grade, MG, in outcome categories results from normalization so that it sums to unity within a contrast set:

$$MG = \frac{MGu}{\sum_{n=1}^{N} MGu_n} \tag{4}$$

where N is the total number of categories in the contrast set for the object.

Figure 1 shows how the galatean model classifies the object S1, S2, S3 from Kruschke's [24] Experiment 1. Each galatea has three components, equating to the three symptoms that occur with the two possible disease categories. These symptoms have two values (cues), one representing their occurrence (1) and one their absence (0). The cues are associated with the conditional probability of the disease given the cue and with the cue's selective attention after competition with the other cues in the category. For example, the galatean component, S1, for the common disease has a conditional probability of 0.75, which translates into a selective attention (cue weighting) of 0.273 after cue competition. When an object is classified, selective attentions from matching object and galatean cues are summed for each disease and

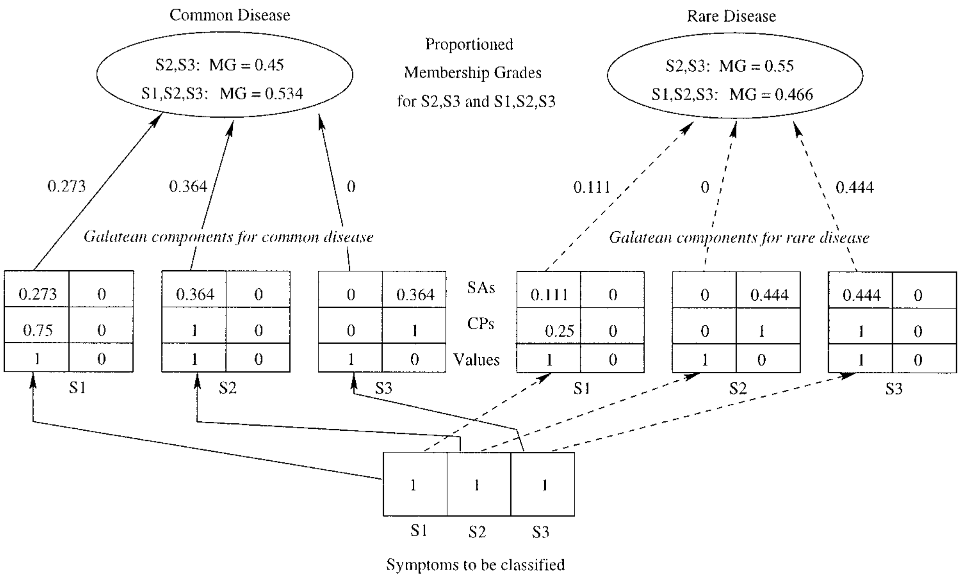

Figure 1. The basic galatean model's classification process: CPs=conditional probabilities, *P*(*Disease*|*value*), and SAs=selective attentions.

proportioned across the diseases to give the final object membership grades. For example, if the object has two symptoms, S2 and S3 (for this experiment, subjects were not explicitly told that symptoms were absent, which is why they are not part of the object), its matching selective attentions in the common disease will be 0.364 and 0 respectively, which, when proportioned with the selective attentions in the rare disease, gives a total membership grade in the common disease of 0.45.

Figure 1 illustrates the model prediction that when all three symptoms of the contrast set of diseases are presented, the common disease has a higher membership grade (0.534) and people will tend to choose it, as Kruschke [24] found (table 2). If the symptom pattern for classification is S2 and S3, the model predicts people will tend to choose the rare category (the membership grade is 0.55) in contravention of base rates, also found by Kruschke.

The core galatean principles of selective attention based on competing conditional probabilities have explained the inverse base-rate effect and have done so by encapsulating a set of quantitatively-defined psychological phenomena that are not collectively contained by any other classification model. In its parameterized form (all classification models are given parameters to generate a closer fit to experimental results), the galatean model performed as well as or better than competing theories on a series of different experimental designs (see [5] for a full description of the model and its fits to experiments). For example, the galatean model classification predictions for Kruschke's Experiment 1 differed from those of the participants with a root mean squared deviation of 0.0297 compared to 0.0308 for Kruschke's connectionist prototype model, ADIT [24].

The remainder of this paper will explore how the galatean model helps our understanding of clinical decision making and how it can be built into effective computer DSSs. First, the galatean model will explain how suicide risk could be exaggerated by clinicians due to the inverse base-rate effect.

4. The inverse base-rate and suicide risk

Figure 2 shows how a strictly hypothetical example of the galatean model might represent suicide risk predicted by seriousness of intention and realism of a plan. The conditional probabilities of the suicide cues show that suicide is far less probable than no suicide (all the cues strongly favour the no-suicide category) but suicide is the more probable outcome according to the model's final membership grades in the two categories. The reason for the bias towards suicide is that its cues are less predictive of it overall and so competition diminishes the conditional probabilities less than it does in the no-suicide category.

Figure 2 demonstrates the galatean model's ability to externalize biases, exposing them and laying them open to countermanding measures. By simulating human psychological processes, the galatean model can represent expert clinical assessment and diagnostic processes in a form comprehensible to clinicians. This facilitates knowledge elicitation, because experts are not being asked to provide knowledge in an alien form, and they are able to amend the model if it does not accurately represent their understanding of the domain. Biases can be removed, thereby improving classification judgements. In other words, the galatean model makes an ideal core of a clinical DSS, but this requires a more complex hierarchical version than the simple, one-level embodiment of figure 2.

5. The hierarchical galatean model

So far, the galatean model has been described as having a single layer of cues that provide membership grades directly to the outcome classes (figure 2). However, experts are aware of relationships between cues and these are incorporated by the full galatean model within a hierarchical knowledge struc-

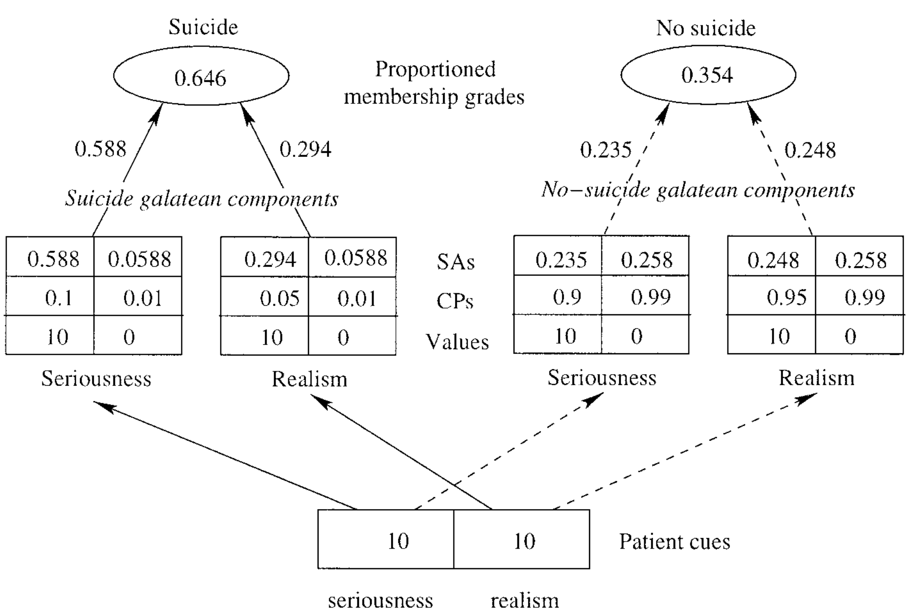

Figure 2. Hypothetical example of the galatean model's classification of a patient into suicide and no-suicide categories. CPs=conditional probabilities and SAs=selective attentions.

ture. Current research is investigating an application of the model to mental-health risk assessment [3] and figure 3 shows a portion of the galatea for people with high suicide risk, elicited from a group of multi-disciplinary mental-health practitioners. The group considered intention to commit suicide to be a concept contributing directly to risk of suicide but not something which is a single, directly measured piece of information such as a person's height. Instead, it is a component requiring further decomposition before reducing to patient cues. It consists of two subcomponents: seriousness of intention, which the clinicians agreed could be directly measured, and information about the plan or method of suicide, which was itself a concept composed of two datum components, one measuring the realism of the plan and the other the steps taken.

The mental-health research has elicited a full hierarchy of concepts and datum components from the expert practitioners and created the Galatean Risk-Screening Tool, GRiST, for gathering data [3]. Figure 4 shows part of GRiST relating to suicide, with the last three questions being those for the intention datum components of figure 3. The range of values (cues) each datum component can take have thus been established but the galatean DSS needs to quantify the degree of risk represented by a patient. This requires the experts to provide selective attentions to the cues and concepts, along the lines of those for the galatean model's application to psychology experiments.

5.1. *Quantifying the galatean model to represent uncertainty*

Equation 2 shows that datum-component weights (selective attentions) are normalized conditional probabilities of cues within a concept's subcomponents (e.g. the contribution of the presence of S1 to the common disease is 0.273, its selective attention). It was clear from previous research on building galateas [4] that experts find it extremely difficult to estimate accurate conditional probabilities directly. Instead, the expert was given two simplified tasks, to generate values that together led to the required cue weightings. First, the expert was required to identify a datum cue that provided the highest support for the corresponding category and give that

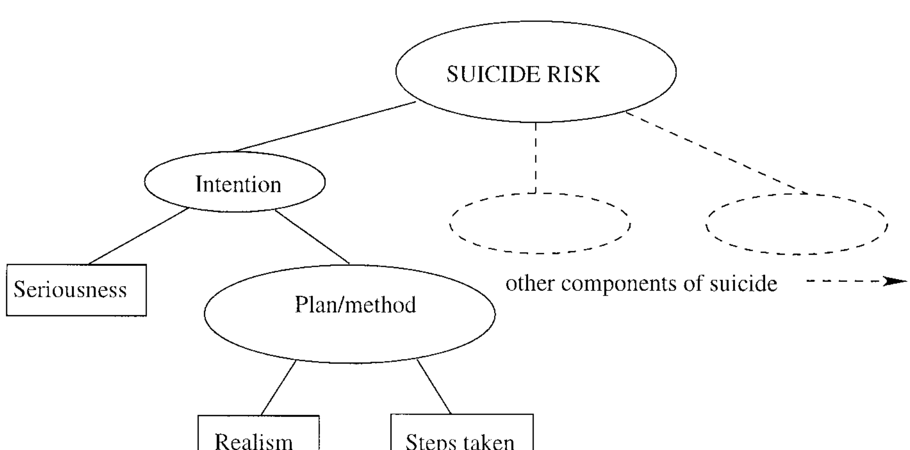

Figure 3. Part of the intention to commit suicide subconcept of the suicide-risk structure; Concepts are shown as ovals and datum components are rectangles.

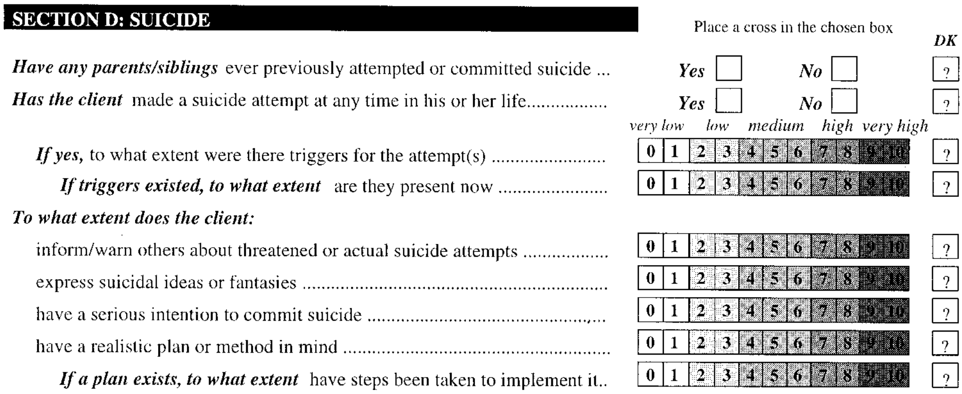

SECTION D: SUICIDE

Place a cross in the chosen box

			DK
Have any parents/siblings ever previously attempted or committed suicide ...	*Yes* ☐	*No* ☐	?
Has the client made a suicide attempt at any time in his or her life................	*Yes* ☐	*No* ☐	?
	very low low medium high very high		
If yes, to what extent were there triggers for the attempt(s)	0 1 2 3 4 5 6 7 8 9 10		?
If triggers existed, to what extent are they present now	0 1 2 3 4 5 6 7 8 9 10		?
To what extent does the client:			
inform/warn others about threatened or actual suicide attempts	0 1 2 3 4 5 6 7 8 9 10		?
express suicidal ideas or fantasies	0 1 2 3 4 5 6 7 8 9 10		?
have a serious intention to commit suicide	0 1 2 3 4 5 6 7 8 9 10		?
have a realistic plan or method in mind	0 1 2 3 4 5 6 7 8 9 10		?
If a plan exists, to what extent have steps been taken to implement it..	0 1 2 3 4 5 6 7 8 9 10		?

Figure 4. Part of GRiST that gathers information about suicide.

cue the maximum membership in the category of 1. For example, the seriousness datum component in figure 3 is evaluated for a patient by the GRiST question 'To what extent does the patient have a serious intention to commit suicide' (figure 4) and has values ranging from 0, meaning not serious at all, and 10, meaning completely serious. A value of 10 clearly provides the greatest risk and this is given the maximum membership grade of one. The expert then gives the worst value zero membership. The galatea for low suicide risk will have the inverse of these membership grades because a high membership in the high-risk category necessarily means a low membership in the alternative low-risk category. In all cases, the total membership grade associated with any datum component value will add up to one across the domain galateas.

Sometimes a datum component will have a perfect value in between 0 and 10, in which case the expert needs to determine which extreme value is the worst and then provide a suitable membership grade between 0 and 1 for the other extreme value. The result is that the datum component cues are given a distribution of membership grades where the membership grade of a value between the perfect and limiting values can be found by linear interpolation. However, membership grades require modification as a result of cue competition and the expert's second task is to encapsulate it by assigning relative influences on classification to each concept's subcomponents.

Figure 5 shows a hypothetical assignment of membership grades and relative influences to the intention concept and how these generate membership grades for a particular patient's values. The galatea subtree is for the high suicide-risk category. It will evaluate a patient's suicide risk with respect to intention by matching patient values with corresponding datum components, calculating resulting membership grades in the datum components, multiplying them by their components' relative influences, and passing the products up the tree until a membership grade is generated for the intention concept. For example, the realism of a patient's plan has values between 0 and 10, with 10 providing the maximum membership (one) and 0 the minimum membership grade (zero). If the patient being classified has a plan with a realism value judged to be 7, then it will generate a membership grade of 0.7 in the realism datum component (shown by the number in italics next to the component).

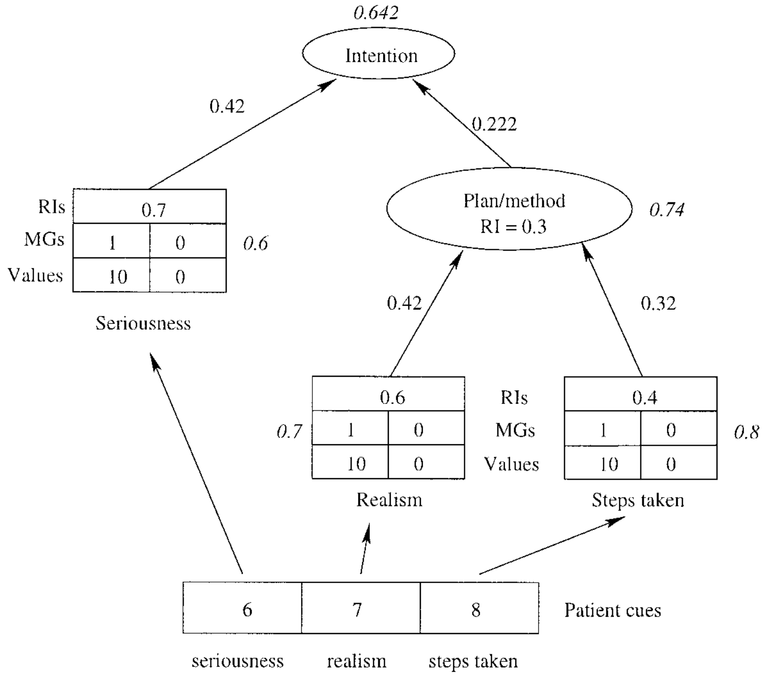

Figure 5. Propagation of membership grades when classifying a patient in the intention to commit suicide hypothetical galatea concept: RI = relative influence; MG = membership grade; numbers in italics beside galatean components represent the patient's membership grade in those components.

The hypothetical expert has assigned a relative influence of 0.6 to the realism datum component and 0.4 to the steps-taken component, indicating that the realism of a plan is more important than what steps have been taken to implement it. All subcomponents of concepts are similarly assigned relative influences, which must add up to one across the subcomponents (the constraints on relative influences and cue membership grades mean that membership grades sum to unity for each subcomponent and the final categories, unless data is missing from the object being classified). The patient's membership grade of 0.7 in the realism galatean component is thus multiplied by the relative influence of 0.6 to pass a membership grade of 0.42 up to the plan/method concept. The summed membership grade in the plan/method concept from its two subcomponents is 0.74 and its contribution to intention after multiplication by its relative influence of 0.3 is 0.222. In this way, the patient's memberships in all the components can be evaluated to give, for the patient shown, a risk membership grade of 0.642 with respect to the patient's intention to commit suicide.

From this description, it can be seen that the selective attentions of competing conditional probabilities given in equation (2) have been replaced by two values, a membership grade and a relative influence. These are multiplied to provide an estimate of the selective attention of equation (2), which can be regarded as a subjective probability [40]. By manipulating the membership grade distribution and relative influences, the experts can 'tune' the classification process to provide class membership grades that correspond to their own estimates of risk for different patient values.

The galatean model's effectiveness in representing expertise was demonstrated for a consultant psychotherapist's assessment of the suitability of patients for psychotherapy [4]. The therapist interviewed 24 potential patients. For each one, she completed a questionnaire that gathered the data identified as being relevant to assessments. The final question asked the therapist to rate the extent to which each patient was suitable, with 1 being none and 6 the maximum extent. Patient data (but not the final suitability rating) were then submitted to the galatean model of the therapist's assessment processes. If the model accurately captured the processes, then it would generate a high membership grade in the suitability category for patients given a high rating by the therapist. Such a positive correlation was strongly exhibited (r=0.825, $p<0.0001$).

Successfully representing expertise is the galatean model's objective but it might mean incorporating inaccuracies such as those caused by the base-rate bias. These could be identified through subsequent empirical analysis and the galatean values adjusted to compensate. Hence the model is able to represent and improve on expertise, both of which are central functions of a DSS. The model's potential implementation within a DSS will be described next.

6. The galatean model as a decision support system

There is considerable evidence suggesting that simple linear models of classification perform better than experts at many tasks, including clinical ones (e.g. [41, 42]). However, these experiments were constrained so that humans were given the same information in the same form as machines. In reality, humans have access to a huge pool of background knowledge which is not open to machines. In domains where experts can employ their wider knowledge, the best decisions are obtained when they work in conjunction with the machine rather than either party operating alone. Such enhancement is 'most likely to occur in dynamic decision environments where a knowledgeable decision maker is able to capitalize on information not captured by the mechanical model' [43, p. 326] which is a fair summary of most clinical situations.

The key to combining the strengths of clinicians with computerized decision support is to ensure that knowledge contained within the DSS is easily understood by clinicians. Because the galatean model aspires to being an accurate rendition of psychological classification, it meets this requirement. Figure 5 shows how a patient's risk assessment can be traced through the galatean hierarchy. It provides a graphical explanation of how expert clinicians generate their judgements, enabling the reasoning processes to be intuitively understood and enhancing clinicians' ability to identify errors and improve DSS performance. Clinicians are able to link their knowledge with that of the galatean DSS rather than being held hostage by incomprehensible algorithms. This transparency also enables the DSS knowledge to be scrutinized by independent clinicians, which opens up an important avenue of external validation.

The mental-health risk assessment project team has recently been awarded a substantial grant to develop a web-based DSS based on the galatean model. This paper has provided psychological and clinical arguments for such a system. The final section will discuss the functionality a galatean DSS could possess and how it would provide an integrated, coherent resource for clinicians, universally available over the world-wide web.

6.1. *Functionality of the proposed DSS*

The philosophy of a galatean DSS is to facilitate a symbiotic relationship between clinician and computer. The galatean model's ability to represent clinical expertise rests in its psychological validity; the computer's side of the symbiosis comes with its powers of data storage and analysis.

The galatean hierarchy will be instantiated with values for quantifying risk using a multi-disciplinary panel of mental-health practitioners. Multi-expert elicitation has become increasingly common (e.g. [44–46]), including elicitation of agreed values for continuously-distributed quantifications of uncertainty [47]. At the same time, patient information will accumulate through GRiST, the paper-based questionnaire, to become the database of a DSS and provide empirical evidence for the computer's data-analytical tools. The DSS will thus contain resources of three types:

1. A database of patient cues and associated risk judgements provided by practitioners as part of their clinical practice
2. A suite of statistical and pattern recognition tools for analysing the database and elucidating the association of cues and risk; some of these tools (e.g. bayesian belief networks and neural nets) will generate risk predictions that can be compared with those of the galateas
3. A model of expert risk assessment that analyses the risk attached to supplied patient information, with an intuitive analysis of how cues contribute to risk in different areas.

The statistical tools and expert risk model complement each other by providing different forms of knowledge. The former is based on empirical evidence and mathematical principles but makes risk predictions that may not be explicable in terms easily understood by people without the required academic background. The latter is based on a validated psychological model of risk assessment elicited from multi-disciplinary experts and explains the generation of risk in accessible terms for practitioners. Both perspectives may help improve the other's performance. The data-analytical tools can help expose inaccuracies in the galatean model's predictions and the galatean model's incorporation of expertise may inform the choice of tools and architectures for data analysis.

The key to the galatea's ability to provide comprehensible advice will be the graphical functions displaying the processing of risk. It needs facilities for tracing and explaining the generation of risk from patient cues, through a hierarchy of concepts underlying risks, to the high-level risks themselves. Risk may be displayed as numbers or even a form of analogue visualization such as the flow of liquid, where the 'deeper' the liquid, the greater the risk. The objective is to elucidate effects of changing galatea parameters on the accumulation of risk and inculcate greater confidence in the results.

The DSS will be made available over the internet, providing a widely-available, general resource to aid mental-health risk assessments. Analysis of the database will generate statistical predictions for integration with expert risk assessments. Together, the two approaches will both enhance risk assessment and provide facilities for education and training of practitioners.

7. Summary and conclusions

This paper began with a review of human reasoning and how it can lead to inaccurate conclusions. In particular, people are inconsistent in their use of base-rates when making probability predictions. The galatean model of psychological classification was shown to give an explanation of these inconsistencies and a hypothetical application of the model demonstrated how clinicians could overestimate rare events such as suicide. The paper then argued that the galatean model's psychological validity renders it suitable for representing expertise. By doing so, it will expose potential biases, enabling them to be countermanded. Hence the model makes an ideal heart of a DSS because it can link intuitive explanations of clinical expertise with empirical data analysis to enhance judgement accuracy. The two elements can also be combined to create an effective source of education and training, all of it potentially available over the internet, increasing dissemination and access.

An important benefit of the galatean model is that it facilitates multi-disciplinary consensus by formulating clinical judgement as a classification process. Classification can be applied to all forms of clinical decision making, thereby removing artificial boundaries created through incompatible theories and terminology. This is crucial in a world where clinical roles are becoming more fluid and information needs to be shared across a range of specialties. A DSS enabling such sharing, both through general comprehensibility of its information and through universal access, is the way forward for health providers and users.

References

1. Bursztajn, H. J., Feinbloom, R. I., Hamm, R. M. and Brodsky, A., 1990, *Medical choices, medical chances*, (London: Routledge).
2. Berner, E. S. (editor), 1999, *Clinical decision support systems: theory and practice*. Health Informatics., (New York: Springer).
3. Buckingham, C. D. and Chan, T., 2002, Developing a mental-health risk-screening tool. Technical report, (Surrey Hampshire Borders NHS Trust).
4. Buckingham, C. D. and Birtle, J., 1997, Representing the assessment process for psychodynamic psychotherapy within a computerized model of human classification. *British Journal of Medical Psychology*, **70**, 1–16.
5. Buckingham, C. D., unpublished, A prototype model of classification based on conditional probabilities.
6. Buckingham, C. D. and Adams, A. E., 2000, Classifying clinical decision making: A unifying approach. *Journal of Advanced Nursing*, **32**(4), 981–989.
7. Kassirer, J. P. and Gorry, G. A., 1978, Clinical problem solving: a behavioral analysis. *Annals of Internal Medicine*, **89**, 245–255.
8. Bordage, G. and Zacks, R., 1984, The structure of medical knowledge in the memories of medical students and general practitioners: Categories and prototypes. *Medical Education*, **18**(6), 406–416.
9. Brooks, L. R., Norman, R. G. and Allen, S. W., 1991, Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*, **120**(3), 278–287.
10. Jones, J. A., 1988, Clinical reasoning in nursing. *Journal of Advanced Nursing*, **13**, 185–192.
11. Radwin, L. E. Knowing the patient: a process model for individualized interventions. *Nursing Research*, **44**(6), 364–370.
12. Taylor, C., 1997, Problem solving in clinical nursing practice. *Journal of Advanced Nursing*, **26**, 329–336.
13. Radwin, L. E., 1998, Empirically generated attributes of experience in nursing. *Journal of Advanced Nursing*, **27**, 590–595.
14. Buckingham, C. D. and Adams, A. E., 2000, Classifying clinical decision making: Interpreting nursing intuition, heuristics, and medical diagnosis. *Journal of Advanced Nursing*, **32**(4), 990–998.
15. Rosch, E. and Mervis, C. B., 1975, Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, **7**, 573–605.
16. Kahneman, D., Slovic, P. and Tversky, A., 1982, *Judgement under uncertainty: Heuristics and biases*, (Cambridge: Cambridge University Press).
17. Evans, J St BT, 1989, *Bias in human reasoning*, (Hove: Lawrence Erlbaum Associates).

18. MCKENZIE, C. R. M., 1994, The accuracy of intuitive judgement strategies: Covariation assessment and Bayesian inference. *Cognitive Psychology*, **26**, 209–239.
19. TVERSKY, A. and KAHNEMAN, D., 1974, Judgment and uncertainty: Heuristics and biases. *Science*, **185**, 1124–1131.
20. EDDY, D. M., 1982, Probabilistic reasoning in clinical medicine: Problems and opportunities. In *Judgement under uncertainty: Heuristics and biases,* edited by D. Kahneman, P. Slovic, and A. Tversky, (Cambridge: Cambridge University Press)pp 153–160.
21. DAWES, R. M., 1983, Is irrationality systematic? *The Behavioural and Brain Sciences*, **3**, 491–492.
22. CIOFFI, J. and MARKHAM, R., 1997, Clinical decision-making by midwives: managing case complexity. *Journal of Advanced Nursing*, **25**, 265–272.
23. REED, S. K., 1972, Pattern recognition and categorization. *Cognitive Psychology*, **3**, 382–407.
24. KRUSCHKE, J. K., 1996, Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**(1), 3–26.
25. NOSOFSKY, R. M., 1986, Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, **115**(1), 39–57.
26. LAMBERTS, K. and FREEMAN, R. P. J., 1999, Building object representations from parts: test of a stochastic sampling model. *Journal of Experimental Psychology: Human Perception and Performance*, **25**(4), 904–926.
27. HOMA, D., STERLING, S. and TREPEL, L., 1981, Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory*, **7**(6), 418–439.
28. MALT, B. C., 1989, An on-line investigation of prototype and exemplar strategies in classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **15**(4), 539–555.
29. ERICKSON, M. A. and KRUSCHKE, J. K., 1998, Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, **127**(2), 107–140.
30. CIOFFI, J., 1997, Heuristics, servants to intuition, in clinical decision-making. *Journal of Advanced Nursing*, **26**, 203–208.
31. OFFREDY, M., 1998, The application of decision making concepts by nurse practitioners in general practice. *Journal of Advanced Nursing*, **28**(5), 988–1000.
32. MEDIN, D. L. and EDELSON, S. M., 1988, Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General*, **117**(1), 68–85.
33. KRUSCHKE, J. K. and JOHANSEN, M. K., 1999, A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **25**(5), 1083–1119.
34. KRUSCHKE, J. K., 2001, The inverse base rate effect is not explained by eliminative inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **27**, 1385–1400.
35. MCKINLEY, S. C. and NOSOFSKY, R. M., 1996, Selective attention and the formation of linear decisions boundaries. *Journal of Experimental Psychology: Human Perception and Performance*, **22**(2), 294–317.
36. HOLZWORTH, R. J. and WILLS, C. E., 1999, Nurses' judgments regarding seclusion and restraint of psychiatric patients: A social judgment analysis. *Research in Nursing and Health*, **22**(3), 189–201.
37. BAKER, A. G., MERCIER, P., VALLÉE-TOURANGEAU, F., FRANK, R. and PAN, M., 1993, Selective associations and causality judgments: Presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **19**, 414–432.
38. MATUTE, H., ARCEDIANO, F. and MILLER, R. R., 1996, Test question modulates cue competition between causes and between effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **22**(1), 182–196.
39. ZADEH, L. A., 1965, Fuzzy sets. *Information Control*, **8**, 338–353.
40. WRIGHT, G. and AYTON, P., 1994, *Subjective Probability*, (New York: Wiley).
41. MEEHL, P. E., 1959, A comparison of clinicians with five statistical methods of identifying psychotic MMPI profiles. *Journal of Counseling Psychology*, **6**(2), 102–109.
42. LOVIE, A. D. and LOVIE, P., 1989, Linear scoring models: lessons for experts and cognitive scientists. In *Expert knowledge and explanation,* edited by C. Ellis (Chichester: Ellis Horwood) pp. 58–75.
43. WHITECOTTON, S. M., SANDERS, D. E. and NORRIS, K. B., 1998, Improving predictive accuracy with a combination of human intuition and mechanical decision aids. *Organizational Behavior and Human Decision Processes*, **76**(3), 325–348.
44. WAGNER, W. P. and HOLSAPPLE, C. W., 1997, An analysis of knowledge acquisition roles and participants. *Expert Systems*, **14**(1), 3–14.
45. AMBROSIADOU, B. V. and GOULIS, D. G., 1999, The Delphi method as a consensus and knowledge acquisition tool for the evaluation of the diabetes system for insulin administration. *Medical Informatics and the Internet in Medicine*, **24**(4), 257–268.
46. VAN STEENKISTE, B. C., JACOBS, J. E., VERHEIJEN, N. M., LEVELINK, J. H. and BOTTEMA, B. J. A. M., 2002, A Delphi techniques as a method for selecting the content of an electronic patient record for asthma. *International Journal of Medical Informatics*, **65**, 7–16.
47. VAN DER FELS-KLERX, I. H. J., GOOSSENS, L. H. J., SAATKAMP, H. W. and HORST, S. H. S., 2002, Elicitation of quantitative data from a heterogeneous expert panel: formal process and application in animal health. *Risk Analysis*, **22**(1), 67–81.