# Designing multiple user perspectives and functionality for clinical decision support systems

Christopher D. Buckingham
and Abu Ahmed
School of Engineering and Applied Science
Aston University
Birmingham, United Kingdom
Email: c.d.buckingham@aston.ac.uk

Ann Adams
Warwick Medical School
University of Warwick
Coventry, United Kingdom
Email: a.e.adams@warwick.ac.uk

*Abstract*—**Clinical Decision Support Systems (CDSSs) need to disseminate expertise in formats that suit different end users and with functionality tuned to the context of assessment. This paper reports research into a method for designing and implementing knowledge structures that facilitate the required flexibility. A psychological model of expertise is represented using a series of formally specified and linked XML trees that capture increasing elements of the model, starting with hierarchical structuring, incorporating reasoning with uncertainty, and ending with delivering the final CDSS. The method was applied to the Galatean Risk and Safety Tool, GRiST, which is a web-based clinical decision support system (www.egrist.org) for assessing mental-health risks. Results of its clinical implementation demonstrate that the method can produce a system that is able to deliver expertise targetted and formatted for specific patient groups, different clinical disciplines, and alternative assessment settings. The approach may be useful for developing other real-world systems using human expertise and is currently being applied to a logistics domain.**

## I. INTRODUCTION

**M**ANY Clinical Decision Support Systems (CDSSs) with appropriate functionality have been successfully developed in academic institutions but never seen the light of day within healthcare practice. There are two fundamental reasons why these systems are not adopted. One is the failure to integrate with the way organisations and their individual employees work. The other is the inability to communicate information effectively beyond the immediate remit of the CDSS, which is often too narrow in the first place. This paper describes a research approach that attempts to circumvent both problems by developing a CDSS that has flexible requirements and data sharing protocols built into the design process from the very beginning. The CDSS is the Galatean Risk and Safety Tool, GRiST [1], [2], that helps assess and manage risks associated with mental-health problems.

The aim of the research was to design GRiST so that it could disseminate mental-health expertise using appropriate language for the particular type of recipient and in a format commensurate with the variable circumstances of assessment. This is no easy task because it would need to accommodate end users ranging from psychiatrists with years of specialist medical education to carers or charity workers who may have minimal training. In fact, GRiST was later adapted for self-assessments, by patients who do not have any predefined common ground apart from mental-health problems. Assessment contexts were also highly variable because GRiST was intended to be deployed for mental-health patients across the care pathway, from primary care, through secondary care and specialist services, and back to care in the community.

The complexity of health services in general and mental health in particular is one reason why the UK Government had so many problems with its National Programme for Information Technology [3] that was intended to revolutionise information systems and processes within the National Health Service (NHS). When GRiST was available for deployment in 2006, the oft-acknowledged "cinderella" mental-health services were still more paper-based than most in the NHS. GRiST set out to tackle barriers to information technology (IT) and its adoption by a design process dedicated to developing flexible interfaces and delivery formats for heterogeneous users and contexts. The research questions were: (i) how can the knowledge base be presented in the format and language most appropriate for each intended type of user? and (ii) how can the information technology generate flexible interfaces to the knowledge so that they fit with the different contexts of assessment?

The paper will first briefly review the clinical rationale for GRiST before describing the main functionality and underlying philosophy of the system. This will provide the context for the cognitive engineering approach that was used to develop knowledge structures providing risk assessments and advice. Their implementation as a sophisticated set of linked XML trees will be described, showing how they support the full GRiST CDSS and its deployment across health and community settings. Examples of the variety of interfaces and language used will be given along with an evaluation of the clinical implementation and adoption. The paper will end by considering the next steps for the research programme and how these have been facilitated by the knowledge structure design principles.

## II. BACKGROUND

To prevent serious untoward incidents (SUIs) amongst inpatients and in the community, clinicians need new and reliable

research evidence to help them detect high risk patients and to support risk management decisions. Despite patient safety being central to NHS policy [4], SUIs remain worryingly common [5]–[7]. Identified causes are lack of sufficient, accessible information about patients' risk profiles [5] and poor risk management or care planning [5], [6]. Risk assessment and management are core competencies for mental health clinicians [8], [9], but the two processes are often not properly connected [10].

There is a clear need to improve clinical practice, which was the motivation for GRiST. It is set apart from alternative risk-assessment and management tools by explicitly modelling human expertise within a generic model of psychological classification. This was a fundamental design principle; if GRiST is based on how humans in general organise their knowledge and reason with it, then its expertise will be in a universally accessible format. It enables GRiST to transcend disciplinary specialisms and opens its expertise up to people with no training at all.

GRiST is designed to assist the early detection of multiple risks amongst people with mental health problems, including suicide, self-harm, harm to others, self-neglect, and vulnerability. It records patient data (cues) to provide a precise information profile that supports the risk judgements given by clinicians.

Risk assessment can be formulated as a classification problem where each risk such as suicide or harm to others is a class and the support for each class determines the level of risk. The factors determining which risk gains most support will be the patient cues such as previous risk history, current intention, emotional and mental state, as deemed relevant by the assessor. The classification task is to formulate the support for each risk from the input data and activate appropriate interventions associated with the most supported class.

In GRiST, risk classes are represented by hierarchical knowledge structures or trees called galateas [11], which are used to represent mental-health expertise. The trunk or root node of the tree is the risk. It is deconstructed into subconcepts that are themselves trees until the leaf nodes are reached, representing the input data.

Figure 1 provides a hypothetical illustration of how the galateas represent classes and their support. The data used for input to the tree can be any type but it is then converted into a fuzzy-set membership grade, MG, from 0 to 1. Zero represents no support for the root decision class and 1 represents maximum support, but for this item of information alone; its MG at this point is independent of any other item. The main role of the MGs is in converting from real-world patient data to the model input. This is shown in Figure 1 by the MG row of the datum nodes, which defines a distribution of MGs matching the range of potential input data values. Values above or below the range take the MG associated with the maximum or minimum value respectively; values within the range are found by linear interpolation if they don't match a value specified in the distribution. For example, the "number of attempts" datum in Figure 1 has the value 3, which is 0.6
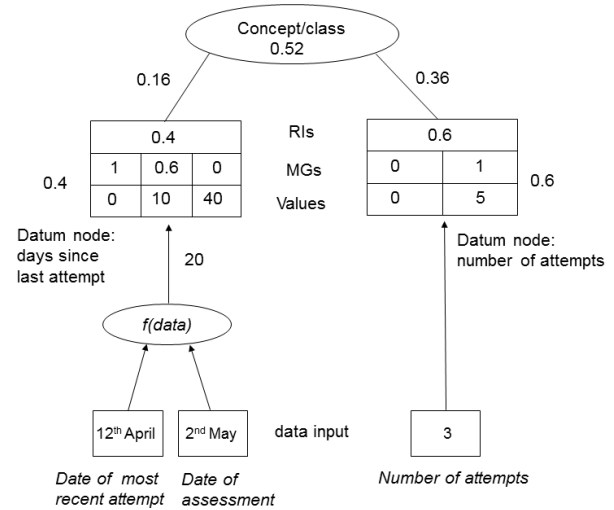


Fig. 1. Hypothetical example of how membership grades (MGs) are processed from patient input data to risk concept. RIs are relative influences, which represent the weights of data and concepts.

along the value range between 0 and 5 and so is assigned an MG that is 0.6 between 0 and 1; i.e. 0.6, as given by the MG outside the box for that datum. For others, such as "days since last attempt", the patient data is passed through a function, *f(data)*, before matching the value-mg distribution: two dates, in this case, which the function uses to generate the number of days between them, 20. Twenty is one-third between 10 and 40 and so is one third along the MG continuum between 0.6 and 0, producing an MG of 0.4, likewise shown outside the datum-node box. The MG is then multiplied by the RI associated with the datum, as shown in the RIs row, to give the MG contribution to the parent concept.

The parent concept MG is the sum of its children contributions, which is how 0.52 is assigned to the concept node in Figure 1. If this concept also had a parent, then the concept would have its own RI and its contribution to the parent would be the product of its RI and MG in the same way that it received its children MGs. The MGs percolate in this manner through to the root node to produce the overall class membership and thus the risk evaluation. Equation 1 formalises the process

$$MG_C = \sum_{i=1}^{n} MG_i RI_{pi} \qquad (1)$$

where $C$ is a concept, $MG$ is the membership grade generated at each datum node, $i$, of the concept, and $RI_{pi}$ is the product of all the RIs along the path, $p$, from the datum node to the concept.

The focus of this paper is how the hierarchical modelling of expertise translates into an ontology that can drive the GRiST CDSS. The added value of the hierarchy is that it represents the conceptual structure understood by human decision makers when relating influential factors to the decisions taken. There is plenty of evidence for the psychological validity of this

hierarchical knowledge structuring. For example, expert chess players "chunk" positions of chess pieces into hierarchical types of game states [12]. Similar strategies have been shown in other domains such as architecture [13], fault diagnosis [14], and medicine [15]. A review of the evidence [16] concluded that "on balance, it is difficult to dismiss hierarchical organisation as only a construct" (p31) and more recent research has begun to show its neural correlates [17], [18].

The psychological grounding of GRiST is not unique, of course, when it comes to intelligent knowledge-based systems (IKBSs) [19], [20]. However, it uses a generic classification model that represents expertise in a non-specialist format. It can be understood without requiring clinical training and makes it ideal for communicating to heterogeneous users (see [11] for more on the Galatean model rationale).

*A. Cognitive engineering and the GRiST ontology*

The GRiST approach to constructing decision support systems can be categorised as cognitive engineering because it is the application of cognitive science to IT systems that are intended to help solve real-world problems [21]. For cognitive engineering, models need to encapsulate expertise in a format that can be accessed by the experts and that is commensurate with the inputs and outputs those experts are familiar with in their problem-solving worlds [22]. IKBS Engineers were coming to this conclusion with the idea of situated cognition [23], which argues that thinking cannot be separated from the environment [24], [25]. These environments change and static IKBSs based on a single, giant elicitation exercise are doomed to fail because they will not be flexible enough to evolve or even be maintained easily [26, pg. 767].

There has been a change in tack from psychology to the data itself, with machine learning, data mining, and pattern recognition approaches coming to the fore. In a recent review of artificial intelligence in medicine [27] Peter Szolovits points out that in the early days, "we thought we knew a lot, but had little or no actual data. Today we are inundated with data, but have correspondingly devalued expertise" (pg. 12). The focus is on the machine, how knowledge can be structured for easy processing, and how useful outputs can be induced from the data. It is the same focus that stimulated the rise of ontologies for organising data into shared knowledge bases. Nevertheless, despite Musen's claim that cognitive models do not lead to scalable and maintainable IKBSs [28], "the symbiosis between cognitive science and cognitive engineering shows no sign of abating" [21, pg. 582] and continues to be the case in medicine [29].

The GRiST research tries to bring human and machine closer together using a form of ontology that has an intuitive connection with the knowledge used by mental-health experts. The interface between human and machine ontologies should be a primary focus for knowledge engineering [30], especially for CDSSs based on clinical expertise. The most basic form of ontology is a controlled and extensible vocabulary [31], [32], which means that dictionaries and thesauri would count. These are very familiar to people and emphasises the point that

ontologies are not strictly the preserve of machines. Indeed, all sensate beings create some kind of ontology for interacting with their environment [33].

Maintaining the intuitive representation of the GRiST knowledge base meant that the terms should reflect the natural language of human users [34], [35]. This is particularly important in mental-health risk screening because of the diversity of information that relates to risk and the lack of any all-encompassing coding schemes. Where schemes do exist, e.g., ICD-10 [36] and DSM-IV [37], they focus on diagnostic categories for mental disorders such as depression and schizophrenia and do not encompass the diversity of peoples' histories and current behaviour that impact on risk. Attempts to create ontologies within mental health have also focused on diagnoses [38], not risk, and have been aimed at data interoperability rather than formalising expertise and clinical decision making.

The GRiST ontology development was designed to ensure the end product met the needs of its users and organisational settings [39] by extensive iteration between clinicians and the evolving CDSS. The galatean psychological model kept the human-machine interface open and intuitive. It has a precisely specified semantics for hierarchical knowledge, incorporating parameters required for processing uncertainty, and the mathematical functions for propagating them through the hierarchy. This coupling of the ontology with its problem-solving method (classification) helps construct a system that solves real-world tasks [40], but does so by emphasising the fluid relationships of human intuition rather than machine formalisms and logic-based reasoners [41]. The next section explains the method in detail.

## III. METHOD

The goal of the methods reported in this paper was to create galatea knowledge structures that were able to evolve with expert consensus and support customised knowledge delivery for a variety of end users accessing it in different contexts. The first problem was how to develop and manage the hierarchical knowledge, which was solved using mind maps.

*A. Mind maps*

One of the most intuitive aids for note-taking, brainstorming, and generally organising ideas is the mind map [42]. Its layout reflects the goal of representing free-flowing, unconstrained associations of the mind at the same time as structuring knowledge hierarchically; it exactly accords with the knowledge-engineering requirements of GRiST.

There are many mind mapping software programs available. Freemind [43] was chosen because it is: (i) open-source; (ii) available across platforms; (iii) creates node structures that can be easily edited; (iv) enables icons to be incorporated into the nodes; (v) attaches notes to the node without obscuring the structure; and, most importantly, (vi) uses XML directly for representing the mind map rather than it being only an export choice. Its structure-editing role was integrated with the GRiST knowledge-engineering toolkit by creating an XSLT
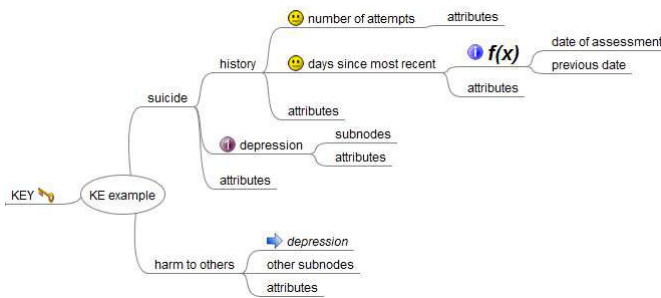
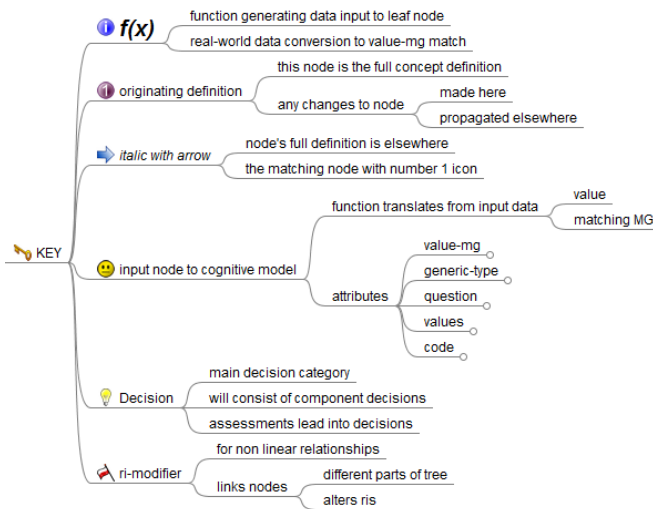Fig. 2.  Hypothetical and simplified mind map of risk assessment expertise



Fig. 3.  Expanded key node showing icons that help drive knowledge engineering

document that transformed the Freemind mind map XML into the GRiST structure tree.

A useful resource for helping users control structure changes and also to direct the style sheet is Freemind's icons. Figure 2 shows a simplified example of how Freemind defines the knowledge structure; Figure 3 expands the "Key" node that explains the icons helping control the translation between mind map specification of knowledge structures and the subsequent GRiST XML trees. Many concepts, such as depression, underlie all risks and so are repeated in the knowledge hierarchy. The blue arrow icon enables the mind map to define the full structure in one place. When the style sheet detects the blue arrow, it looks for a node with the same name that has the round number 1 icon associated with it (see Figure 2). The other icons are similarly used to specify aspects of the galatean structure, such as the face, which identifies leaf nodes of the galatea where value-mg distributions are defined. The $f(x)$ node indicates which patient data are required to generate the matching value to the value-mg distribution. This is the case for the time period between the assessment and the most recent suicide attempt, for example, as shown in Figure 1.

Every risk node, both leaf and concept, has a subnode called "attributes", which contains attributes required by the GRiST

XML trees. These enable the XSLT conversion document to translate between Freemind mind map format and GRiST XML nodes by making the attributes an explicit structure in Freemind. Otherwise, they would be unrecognised and ignored by Freemind when creating the mind map. The XSLT conversion document looks in the attributes subnode and creates them as well-formed attributes of the output XML tree that is at the root of the GRiST ontology. The next section introduces the GRiST XML trees and their attributes in more detail.

### B. GRiST XML tree functionalities, attributes, and relationships

Once the initial knowledge structure has been specified in Freemind, it is translated via the XSLT specification into the GRiST initial XML. The idea is to have a base tree that incorporates the requirements for all patient types and assessment circumstances. It is a kind of "universal" or Everyman tree incorporating every issue for every user. The knowledge-engineering task is to encapsulate the different subtypes with their particular perspectives and priorities within the GRiST XML trees and extract them for delivery within the CDSS.

For GRiST, the subtypes reflect the variety of patient being assessed and the contexts of assessment. Four patient types or *populations* were quickly distinguished as GRiST developed: children and adolescents; working-age adults; older adults; and learning disabilities. Delivering GRiST across assessment contexts also required functional variations that may apply to more than one population (i.e. are not unified with populations) and so needed to be treated separately. Customisations of this type are called services. GRiST is thus tailored along the axes of populations for different assessment trees and services, if it turns out that the functionality needs customising as well as the underlying classification tree; service functionality can be applied to combined subsets of more than one population.

Three objectives were pursued when designing the GRiST XML structures: (i) define the structures of all trees; (ii) instantiate the trees with parameters required for classifying patients according to their particular population; and (iii) generate the specific data structures required for delivering the variety of CDSS functionality for end users. The trees can be summarised as follows (Table I defines the main attributes):

*The Super Structure Tree (SST)*, which contains for all populations, all structural information about nodes and the questions attached to them, with associated values and membership grades. The SST also enumerates all the services that it may be used in with the accompanying more modest functional customisations required across different end users. The SST is the base "Everyman" tree that holds common information across them all as well as information about how to generate the distinctive sub-trees.

*The Structure Tree (ST)*, which contains structural, question, value, and value-mg information for an individual population. It is generated from the SST, and can be conceived of as an SST tailored for one and only one population. Service

TABLE I
EXAMPLES OF ATTRIBUTES USED BY THE GRiST XML SPECIFICATIONS TO DRIVE KNOWLEDGE ENGINEERING TOOLS AND DELIVERY FUNCTIONALITY
OF THE CDSS. THE TREES COLUMN IDENTIFIES THE TREES THAT CONTAIN THE ATTRIBUTE.

| Attribute | Semantics | Trees |
|---|---|---|
| `label` | name of tree node that can vary for populations | all trees |
| `code` | code for tree node, which is invariant | all trees |
| `populations="(population-name)"` | different populations of users defined by the tree | SST |
| `services` | defines services with particular configurations | SST, ST, RIT, |
| `help="(help text)"` | | SST, ST, RIT, QT |
| `generic="[path to generic node]"` | locates full definition of node | SST, ST, RIT |
| `generic-type="g"` | repeating nodes with invariant uncertainty parameters | SST, ST, RIT |
| `generic-type="gd"` | repeating nodes with varying uncertainty parameters | SST, ST, RIT, CAT |
| `generic-datum="[path to definition of datum]"` | locates full definition of node | SST, ST, RIT |
| `value-mg="((0 0)(7 1)(10 0.5))"` | association list of values and membership grades | SST, ST, RIT |
| `level` | prunes tree at different levels of assessor expertise | SST, ST, RIT |
| `question="question"` | question for collecting item of information | SST, ST, RIT |
| `values="values"` | defines the type of the item of information | SST, ST, RIT, QT |
| `layer="n"` | specifies order of initial data collection | SST, ST, RIT, CAT |
| `filter-q="question"` | question indicating whether subtree is applicable or not | SST, ST, RIT, CAT |
| `persistent="hard/soft/value"` | carries data forward from previous assessment | SST, ST, RIT, QT |
| `service` | configures node with given services customisation | SST, ST, RIT, CAT |
| `prune-for` | removes branch/node for a population | SST |
| `other attributes ...` | emerging out of knowledge engineering | any |

definitions contained in the SST will be carried over to the ST, meaning that customisations defined for a given service type will apply across all populations. The ST is used to generate the RIT corresponding to a population.

*The Relative Influence Tree (RIT)*, which holds the RIs for all nodes. Nodes with generic-type attribute of "gd" are expanded in all locations that point to them because these nodes may have different internal RIs (for concepts) or value-mgs (for datum nodes) in the locations. The RIT structure is generated from the ST and used to elicit and store the RIs.

Together, the ST and RIT are sufficient to specify all the information required for the Galatean psychological model of classification to be instantiated. They provide the complete ontology and problem-solving package and are the end products of the knowledge elicitation stage. However, four more trees are required by the end-user decision support tools. Three are derived from the ST and/or RIT and one is generated by the decision tool during the assessment of decisions, as follows:

*The Class Assessment Tree (CAT)* is generated from the RIT because it needs the RI attribute. It produces the full galatean tree for classifying objects and so has all nodes fully expanded in all locations, with no paths to separate generic nodes.

*The Question Tree (QT)* is generated from the RIT and has all the information required to display questions, obtain associated answers, and generate membership grades for the answers.

*The Answer Tree (AT)* is generated during an assessment by the data-gathering tool, and stores all user-supplied data.

*The Landmark Tree (LT)* is a tree used for helping assessors to navigate the CAT during assessments. It is a reduced version of the CAT, and will highlight nodes that may be of particular interest or relevance to the current assessment. It is envisaged that this will be derived from the Freemind mind map defining the base structure.

Table I provides examples of attributes that represent data defining the psychological model and its specifications of variations for populations and services. For example, `prune-for="(older working-age)"` means remove this branch of Everyman for older-age and working-age adults but not for any of the other listed populations. Many of the node attributes have an "enhanced" form where they can have different values for the populations of classification trees. This was required to encompass the variety of end-user perspectives that went beyond simply providing different views of the Everyman tree structure but also different representations of the data. For example, the tree node labels for the service-user self-assessment population are different to those seen by the mental-health practitioners assessing services users, as shown for the suicide node label:

```
label="(((service-user) "ending your own
life") ((iapt learning-disabilities older
child-adolescent working-age) "suicide"))"
```

Lisp-like association lists are used to pair the population or service with its customised value. They provide great flexibility for dynamically creating and updating customisations within the SST. They are integral to giving the correct values to the tree transformation procedures that generate the correct trees for each population and service.

Customisation of the same population tree across different service provider contexts is effected by the `services` attribute in the top-level root node of the ST. Each service

Fig. 4. Screenshot of the admin interface for managing population trees

provider represents a particular set of (minor) customisations/configurations that will be applied to GRiST's question set when conducting an assessment for that service. The top-level root node of the ST will have a:

```
services=
  "((structure
      ((service1 (association list of mods))
       (service2 (association list of mods))
       (...)))
    (rendition
      ((service1 (association list of mods))
       (service2 (association list of mods))
       (...)))"
```

attribute, defining all the services for which customisation/configuration data exists in the ST. Within the `services` attribute, these modifications (abbreviated to mods in the example) are organised as:

*Structural modifications:* those that involve dynamic (just-in-time) manipulations of the trees in some way prior to their being used to drive a GRiST assessment. Assessment tools will be agnostic of the structural changes, and will therefore not need to perform any additional processing.

*Rendition modifications:* those that involve dynamic (just-in-time) manipulations of the rendition of the GRiST assessment. Additional coding effort will be required in each assessment tool to realise the rendition manipulations.

The normal behaviour is that in order for a tree node to be amenable to the application of the defined modifications, it needs to "subscribe" to the modifications via a service attribute. For example, `service="rendition"` placed in a node will cause all the service-specified renditions to be applied to that node, such as providing it with a prefix or changing the font to bold, both of which were required for the IAPT service described in the results.

## IV. RESULTS

The principle behind managing the different XML trees is to have one single master tree (i.e., the SST) that is used to generate all the other ones. The conversions will be carried out using XSLT in the main and the resulting trees will be labelled so that they can be linked to the particular master tree from which they are derived.

An administrator's interface was provided for uploading, viewing and manipulating the trees (Figure 4). The website automates the derivation of all trees from each SST: namely each population's STs, RITs, and CATs and QTs at various levels (higher level trees pruned at concepts can be used for assessors with greater expertise who can use judgements in place of low-level data). Active trees (i.e. those delivered to end users) need to be marked so that experimental or legacy trees can be made or kept available alongside the current "live" ones for testing before deployment. When the GRiST CDSS is accessed from a patient record system, parameters are passed to the clinical server to indicate which patient assessent is being conducted and what population to invoke.

The success of the methodological approach has been clearly demonstrated over the years by the ability to create new trees suited to particular patient types when requested by mental-health practitioners. The first derivations enable GRiST to cover all age ranges, not just working age, and an additional specialist population was recently added for learning disabilities. Most conclusive was the need to produce a tool for self-assessments that patients could use in the community. This led to the development of myGRiST that exploited the enhanced attribute values to produce variations of nearly all the tree nodes. Not only were different data-collection questions and accompanying help text boxes required but also most tree branch names were changed to ones more suitable for non-clinicians. These substantial variations were all linked to the same common Everyman tree, which meant that clinical and patient answers were always directly comparable: the trees provide a common language and knowledge base despite their multiple manifestations.

The method facilitates a single genotype with a variety of phenotypes. But it also provides different drivers of the end-user CDSS tools. Clinicians tend to be under serious time pressure and want to collect data as efficiently and concisely as possible. This meant their tool was driven by the ST, which keeps class-specific questions separate from generic concepts that are applicable to all classes, and only expands generic concepts once. In other words, it reflects the original mind map with no structural redundancy. On the other hand, the service users conducting self assessments have more time to explore risks dynamically. They wanted less control over the order of access and they also wished to answer generic concepts such as their relationships, current behaviour, living conditions, etc. in the context of whichever risk they were considering. For them, their tool was driven by the CAT, which expands all concepts in all locations to provide full trees for every risk.

The first service customisations were also motivated by time pressures, especially for primary-care practices. Here, the visibility and rendition of risk tree data was needed across all populations (i.e. not limited to only working-age adults). At present, the main service customisation currently in use is for a primary-care service provided to general practitioners called Improving Access to Psychological Therapies (IAPT).

Since mental-health organisations started using the electronic (web service) version of GRiST in 2010, more than 2,000 clinicians have conducted the following number of completed assessments for the different populations and the IAPT service.

*working age population:* 50,193
*children and adolescents population:* 4,008
*older adults population:* 28,188
*iapt service:* 696

These results are testimony to the knowledge engineering method that facilitated accommodation of different end-user requirements. They also demonstrate the robustness of their implementation, with different populations and services being requested dynamically in real-time.

## V. CONCLUSIONS

This paper has described a methodology for eliciting, evolving, and delivering complex mental-health expertise using a psychological model of classification. The new research reported here develops GRiST from its initial construction [44], [45] into a fully-fledged knowledge engineering environment that can manage both structure changes and subtle variations in knowledge parameters within an integrated system. The sophisticated application of attributes and XSLT to deliver customised services has been proven by continuous delivery to mental-health practitioners, every hour of every day of every week.

GRiST is generating an accumulating database of patient information that will help parameterise the galatean model underlying the CDSS. The RIs will be learned from the data to provide models of expert consensus that can detect risks more accurately and target appropriate advice. Efficiency of data collection will exploit the latest research on fast and frugal classification [46], [47] where the most important information is identified first and processed very rapidly. More in-depth analysis only needs to take place if the decision-maker remains ambivalent.

New GRiST versions have been requested for forensic services, accident and emergency, and ante-natal clinics. The knowledge engineering methodology means they can be delivered in short timespans because few or no changes are required in end-user applications. Structural information in the GRiST ontology is given by the node nesting but information about the role of the node, both for the psychological classification model and its delivery within a CDSS, is held by attributes. There is no limit to the number of attributes that can be added, which provides great flexibility for amending node behaviours as knowledge engineering progresses.

The disadvantage is that the meaning of attributes has to be recorded outside the XML. Ontology specification languages such as OWL [48] have more power to include semantics and logical relationships but the galatean approach traded flexibility of knowledge structure with the need for a detailed specification document to accompany it. Any tools operating on the XML would need to refer to the specification document

and ensure their operations were in full accordance with the definitions. Once this has been achieved, then any changes in the ontology using existing attributes require no further software development.

The GRiST ontology was able to integrate natural language for people with codes more convenient to machines, helping to keep the human-machine interface closely aligned [34], [35]. Now that it has stabilised, it makes more sense to consider translating it into a formal specification language [49]. This would make it easier to validate the tree transformations and embed the rules more formally within the machine specification. An ontology language would also help with the problems of interoperability that are endemic within mental-health services, possibly to a greater extent than in other health areas. There is an inherent difficulty with categorising intangible mental-health constructs but data interchange on more straightforward health and social-care patient data would be beneficial. GRiST is explicit about how this generic information links to risks but mental-health organisations often collect it in other documentation not related to risk. Sharing it has proved problematic and its absence from risk documentation could present a danger to proper care [10].

Whether or not the current GRiST ontology is converted into a language such as OWL, instantiation of the galatean model as a formal specification of mental-health risk expertise has led to assessment tools that have generated considerable interest both in the UK and abroad. The method has applicability to any domain where human expertise can be disseminated within a DSS and where the DSS will be used by people with varying needs, characteristics, and work contexts. Evidence for this is emerging from the logistics domain where the galatean approach is being applied. The expertise and decision context is very different, with the goal being to optimise the use of vehicles for deliveries and collections based on predicting the number of orders. The knowledge trees have already successfully captured expertise that is able to represent alternative perspectives [50] and implementation of the CDSS is currently underway. Irrespective of the particular application domain, whether a CDSS is actually used in the real world depends on how flexible it is in meeting varying user requirements and *modus operandi*; this paper reports a method that should help improve its chances of adoption.

## REFERENCES

[1] GRiST, "Galatean risk and safety tool," www.egrist.org, 2013, accessed May 15th.
[2] C. D. Buckingham and A. Adams, "The grist web-based decision support system for mental-health risk assessment and management," in *Proceedings of the First BCS Health in Wales/ehi₂ joint Workshop*, 2011, pp. 37–40.

[3] R. M. Tackley, "The National Programme for Information Technology in the NHS," *Anaesthesia & intensive care medicine*, vol. 5, no. 12, pp. 400 – 401, 2004. [Online]. Available: http://www.sciencedirect.com/science/article/B8CXB-4MBDBCS-3/2/343b3d085e232189d0974d9eba5fe0ef

[4] Department of Health, *Equity and excellence: liberating the NHS*. London: DH Publications, 2010.

[5] Health Care Commisssion, *Learning from investigations*. London: Commission for Healthcare Audit and Inspection, 2008.

[6] Centre for Suicide Prevention, "Independent homicide investigations: the national confidential inquiry into suicide and homicide by people with mental illness," University of Manchester, Tech. Rep., 2008.

[7] N. Rose, "Oxford serious incident review: 7 years on," *The Psychiatrist*, vol. 32, pp. 307–309, 2008.

[8] Department of Health, *The ten essential shared capabilities - a framework for the whole of the mental health workforce*. London: DH Publications, 2004.

[9] ——, *Best practice competencies and capabilities for pre-registration mental health nurses in England: the chief nursing officer's review of mental health nursing*. London: DH Publications, 2006.

[10] E. Gilbert, A. Adams, and C. D. Buckingham, "Examining the relationship between risk assessment and risk management in mental health," *Psychiatric and Mental Health Nursing*, no. 10, pp. 862–868, 2011.

[11] C. D. Buckingham, "Psychological cue use and implications for a clinical decision support system," *Medical Informatics and the Internet in Medicine*, vol. 27, no. 4, pp. 237–251, 2002.

[12] H. Freyhoff, H. Gruber, and A. Ziegler, "Expertise and hierarchical knowledge representation in chess," *Psychological Research*, vol. 54, pp. 32–37, 1992.

[13] O. Akin, *Models of Architectural Knowledge*. London: Pion, 1980.

[14] D. E. Egan and B. J. Schwartz, "Chunking in recall of symbolic drawings," *Memory & Cognition*, vol. 7, pp. 149–158, 1979.

[15] V. L. Patel and J. F. Arocha, "The nature of constraints on collaborative decision-making in health care settings," in *Linking expertise and naturalistic decision making*, E. Salas and G. A. Klein, Eds. Mahwah, NJ: Erlbaum, 2001, pp. 383–405.

[16] G. Cohen, "Hierarchical models in cognition: Do they have psychological reality?" *European Journal of Cognitive Psychology*, vol. 12, no. 1, pp. 1 – 36, 2000.

[17] J.Z.Tsien, "The memory code," *Scientific American*, pp. 52–59, June 2007.

[18] M. Declercq and J. De Houwer, "Evidence for a hierarchical structure underlying avoidance behavior," *Journal of Experimental Psychology: Animal Behavior Processes*, vol. 35, pp. 123–128, 2009.

[19] A.-M. Chang, C. W. Holsapple, and A. B. Whinston, "A hyperknowledge framework of decision support systems," *Information Processing & Management*, vol. 30, no. 4, pp. 473 – 498, 1994.

[20] G. Lindgaard, C. Pyper, M. Frize, and R. Walker, "Does bayes have it? decision support systems in diagnostic medicine," *International Journal of Industrial Ergonomics*, vol. 39, no. 3, pp. 524 – 532, 2009. [Online]. Available: http://www.sciencedirect.com/science/article/B6V31-4VJJW5J-2/2/1b62c3ff120a78426387049fe33c9de4

[21] W. D. Gray, "Cognitive modeling for cognitive engineering," in *The Cambridge handbook of computational psychology*, R. Sun, Ed. Cambridge: Cambridge University Press, 2008, pp. 565–588.

[22] G. Gigerenzer, *Adaptive thinking: rationality in the real world*, ser. Evolution and Cognition. Oxford University Press, 2000.

[23] W. J. Clancy, *Situated cognition*. Cambridge: Cambridge University Press, 1997.

[24] J. J. Gibson, *The ecological approach to visual perception*. Hillsdale, NJ: Lawrence Erlbaum, 1979/1986.

[25] H. Heft, *Ecological psychology in context: James Gibson, Roger Barker, and the legacy of William James's radical empiricism*. Lawrence Erlbaum, 2005.

[26] T. Menzies and W. Clancey, "Editorial: the challenge of situated cognition for symbolic knowledge-based systems," *International Journal of Human-Computer Studies*, vol. 49, no. 6, pp. 767–769, 1998.

[27] V. Patel, E. Shortliffe, M. Stefanelli, P. Szolovits, M. Berthold, R. Bellazzi, and A. Abu-Hanna, "The coming of age of artificial intelligence in medicine," *Artificial Intelligence in Medicine*, vol. 46, no. 1, pp. 5–17, 2009.

[28] M. A. Musen, "Technology for building intelligent systems: From psychology to engineering," in *Nebraska Symposium on Motivation*, B. Shuart, W. Spaulding, and J. Poland, Eds. Lincoln, Nebraska: U Nebraska P, 2009, vol. 52, pp. 145–184.

[29] A. Jalote-Parmar, P. Badke-Schaub, W. Ali, and E. Samset, "Cognitive processes as integrative component for developing expert decision-making systems: A workflow centered framework," *Journal of Biomedical Informatics*, vol. 43, no. 1, pp. 60 – 74, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/B6WHD-4WS2HXV-1/2/abbdd78b426b1f0c91d8684316a8eceb

[30] C. Brewster and K. O'Hara, "Knowledge representation with ontologies: Present challenges–Future possibilities," *International Journal of Human-Computer Studies*, vol. 65, no. 7, pp. 563–568, 2007.

[31] D. L. McGuinness and F. van Harmelen, "OWL web ontology language overview," W3C, W3C Recommendation, Feb. 2004, http://www.w3.org/TR/owl-features/.

[32] F. Villa, I. Athanasiadis, and A. Rizzoli, "Modelling with knowledge: A review of emerging semantic approaches to environmental modelling," *Environmental Modelling & Software*, vol. 24, no. 5, pp. 577–587, 2009.

[33] A. Sloman, "Putting the Pieces Together Again," in *The Cambridge handbook of computational psychology*, R. Sun, Ed. Cambridge: Cambridge University Press, 2008, pp. 684–709.

[34] B. Hu, S. Dasmahapatra, D. Dupplaw, P. Lewis, and N. Shadbolt, "Reflections on a medical ontology," *International journal of human-computer studies*, vol. 65, no. 7, pp. 569–582, 2007.

[35] Y. Wilks, "The Semantic Web: Apotheosis of annotation, but what are its semantics?" *IEEE Intelligent Systems*, vol. 23, no. 3, pp. 41–49, 2008.

[36] World Health Organization, *The ICD-10 classification of mental and behavioural disorders: Clinical descriptions and diagnostic guidelines*. Geneva, Switzerland: World Health Organisation, 1992.

[37] American Psychiatric Association, *Diagnostic and statistical manual of mental disorders*, 4th ed. American Psychiatric Association, 2000.

[38] J. S. Kola, J. Harris, S. Lawrie, A. Rector, C. Goble, and M. Martone, "Towards an ontology for psychosis," *Cognitive Systems Research*, vol. 11, no. 1, pp. 42 – 52, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/B6W6C-4TBXGJ7-1/2/53faa35b5d32dc03e1cd54cb9c72d349

[39] B. Kaplan, "Evaluating informatics applications - clinical decision support systems literature review," *International Journal of Medical Informatics*, vol. 64, pp. 15–37, 2001.

[40] M. Musen, "Ontologies: Necessary–Indeed Essential–but Not Sufficient," *IEEE Intelligent Systems*, vol. 19, no. 1, pp. 77–79, 2004, january–February issue.

[41] M. K. Smith, C. Welty, and D. L. McGuinness, "OWL web ontology language guide," W3C, W3C Recommendation, Feb 2004, http://www.w3.org/TR/owl-guide/.

[42] T. Buzan, *The Mind Map Book*. BBC Consumer Publishing: London, 2003.

[43] Freemind, http://freemind.sourceforge.net/wiki/index.php/Main_Page, 2013, url accessed in May, 2013.

[44] C. D. Buckingham, A. Ahmed, and A. E. Adams, "Using XML and XSLT for flexible elicitation of mental-health risk knowledge," *Medical Informatics and the Internet in Medicine*, vol. 32, no. 1, pp. 65–81, 2007.

[45] C. D. Buckingham, A. E. Adams, and C. Mace, "Cues and knowledge structures used by mental-health professionals when making risk assessments," *Journal of Mental Health*, vol. 17, no. 3, pp. 299–314, 2008.

[46] G. Gigerenzer, U. Hoffrage, and D. Goldstein, "Fast and frugal heuristics are plausible models of cognition: Reply to Dougher ty, Franco-Watkins, and Thomas (2008)." *Psychological Review*, vol. 115, no. 1, pp. 230–237, 2008.

[47] R. Hertwig and U. Hoffrage, *Simple heuristics in a social world*. New York: Oxford University Press, 2013.

[48] W3C, "Owl," www.w3.org/TR/rif-rdf-owl, accessed 21-5-2013.

[49] S. Buckingham-Shum, "Contentious, dynamic, multimodal domains ... and ontologies?" *IEEE Intelligent Systems*, vol. 19, no. 1, pp. 80–81, 2004, january–February issue.

[50] C. D. Buckingham, P. Buijs, P. G. Welch, A. Kumar, and A. Ahmed, "Developing a cognitive model of decision-making to support members of hub-and-spoke logistics networks," in *Proceedings of the 14th International Conference on Modern Information Technology in the Innovation Processes of the Industrial Enterprises*, L. M. Elisabeth Ilie-Zudor, Zsolt Kemény, Ed. Hungarian Academy of Sciences, Computer and Automation Research Institute, 2012, pp. 14–30. [Online]. Available: igor.xen.emi.sztaki.hu/mitip/media/MITIP2012_proceedings.pdf