ADVISES Project: Scenario-based Requirements Analysis for e-Science Applications

Alistair Sutcliffe, Sarah Thew, Colin Venters, Oscar de Bruijn
Business School, University of Manchester
+ 44 161 306 3315, ags@man.ac.uk
John McNaught
National Centre for Text Mining, University of Manchester
Rob Procter
National Centre for e-Social Science, University of Manchester
Iain Buchan
North West Institute for Bio-Health Informatics (NIHBI)
School of Medicine, University of Manchester

Abstract

We describe the requirements analysis method for e-Science which is being developed in the ADVISES project. The procedure and techniques of the method are illustrated by application to two cases studies in epidemiological bio-health informatics. Lessons learned in applying the scenario based analysis method which also embeds Human Computer Interaction principles for design and sub language analysis are discussed, in the context of data intensive e-Science research where exploratory data analysis complements more traditional hypotheses driven experimental research.

1. Introduction

Requirements analysis is acknowledged to be the most error-prone phase of systems development and a major source of failure in software projects (Standish group, 2007) and e-Science projects are no exception. In fact, arguably, e-Science poses even greater challenges for requirements analysis (Zimmerman and Nardi, 2006). For example, e-Science projects may involve significant innovations in scientific work practices, the precise nature of which may be uncertain at the beginning; and they are often distributed, large scale and multi-disciplinary which may make it difficult to establish a consensus among users. In spite of these apparent problems, little attention has been paid to requirements in the UK e-Science programme. Some exceptions are Perrone et al. (2006) which describes the application of a use case and goals-questions-results requirements analysis method, and Jirotka et al. (2004) which describes the application of ethnographic methods for understanding users' working practices in the eDiaMoND project.

The ADVISES project is producing visualisation and data analysis software for medical epidemiologists. The project has two research aims: to improve requirements analysis methodology in e-Science, and to apply Human Computer Interaction (HCI) knowledge to the design of usable and effective visualisations for bio-health informatics. As well as producing usable and useful applications for our users, we intend to produce generalisable methods for requirements analysis and usability engineering for e-Science; both acknowledged needs (Beckles, 2005). In this paper we describe the requirements analysis and usability engineering approaches we have been developing and applying in the project. The following section reviews requirements analysis approaches and describes the ADVISES method. Section three addresses the design phase and reports on the approach we have adopted to design flexible composable visualisation for epidemiology research, and illustrates how HCI knowledge has been applied in the context of an user-centred and iterative methodology in which an eclectic 'toolkit' of techniques developed by the requirements analysis community over the past fifteen years are flexibly and pragmatically applied. This is followed by lessons learned from our experience of requirements investigation. We conclude with a discussion of the approaches we have adopted and their limitations.

2. Requirements Analysis Methodology

The ADVISES approach to requirements analysis is based on scenario-based design (Carroll, 2000) and user-centred requirements engineering (Sommerville and Sawyer 1997; Sutcliffe, 2003), both of which advocate the use of scenarios, storyboards and prototypes in an iterative cycle of requirements elicitation, design exploration and user feedback. Scenario-based design (SBD) was chosen in view of the often volatile and complex requirements of e-Science applications. As research practices often change as an investigation evolves, requirements can become a moving target, which is particularly true in the rapidly developing field of bio-health informatics. SBD is well suited to such circumstances because of its iterative approach which facilitates user-developer dialogue. This rapid iterative approach to development is shared by rapid application development methods (e.g., DSDM, 1995) and, agile software development methods (Beck, 1999).

Another advantage of using scenarios is that they allow for an effective exploration of users' problems and goals, without focussing on the detail of the solution. These scenarios can then be used to integrate usability engineering within requirements analysis by applying HCI design principles in the form of psychological design rationale or *claims*. Claims (Carroll et al., 1998) record design experience as generalisable, hence reusable, HCI knowledge in the form of principles with upsides (advantages) and downsides (disadvantages), illustrated by an exemplar design and motivating scenario that captures the users' requirements. Claims thus give designers packets of reusable advice expressed in a context with trade-offs. In ADVISES we reuse existing claims as well as proposing new claims that can be reused for visualisation design in e-Science.

The focus of our scenarios is the research questions that scientists pose and investigate. Hence, our aim was to see *requirements as questions*. This builds on the GQR (goals-questions-results) method applied in data intensive medical research (Perrone et al., 2006). Analysing researchers' questions implies a language-based approach to requirements analysis to understand the nature of research questions and their expression. Our approach was to capture research questions in epidemiological bio-health informatics and transform these into requirements for services, analysis workflow and the design of the user interface operations to allow users to pose questions in a flexible manner. This approach was also motivated by the natural language query interface for a medical records database in the CLEF project (Hallett et al., 2005), which allows users to pose questions in restricted natural language using a form-filling dialogue.

The ADVISES requirements approach is summarised in figure 1. In one track we analyse our users' language by applying manual and automated information extraction tools to archive publications and recorded conversations of research meetings. In the other track, users' requirements are elicited by more conventional methods of interviewing and observing users' work practices with ethnographic techniques.

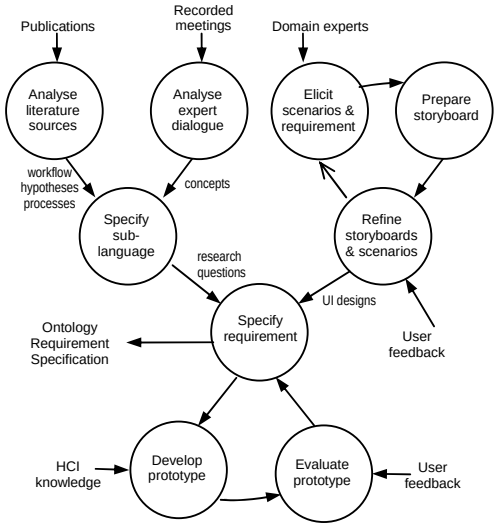

Figure 1. Requirements analysis process.

Throughout the requirements process we treat our chosen methods (observation, storyboarding, interviews, scenario creation, etc.) as a toolkit. That is to say, we do not follow a rigid process or methodology, but select an approach based on both project goals and the circumstances at that time. Early in the requirements process, storyboards representing initial views of screen designs and workflow processes are used to illustrate our ideas and gain feedback from our users. Prototypes are introduced when requirements have been refined by storyboard analysis. The two tracks are brought together in the prototyping phase when user research questions are specified and then translated into user interface operations and commands.

Experience in applying this method is reported in two domains: general epidemiology and genetic epidemiology.

3. Design

The ADVISES project is working with users from two domains: general epidemiology in public health and genetic epidemiology. Both share a common interest in developing software that can support hypothesis and data driven research.

With the burgeoning growth in scientific data, researchers in bio- and health informatics are exploring data-driven hypothesis generation as a complement to hypothesis-driven experimentation (Kell & Oliver, 2004). At present, there is limited adoption of this early 'hypothesis generating' data exploration stage. This is partly due to a lack of tools to support data exploration in their particular domain, and also because it is unfamiliar to scientific researchers whose accepted way of working is to develop a detailed hypothesis based on existing knowledge, and only then test it on data. Consequently, the tools we develop must support the users in extending their mental model of the ways in which they work. SBD facilitates development of vision scenarios and prototype design illustrations for new data driven e-Science research.

In this project, the focus on 'hypothesis generation' first arose during discussions between a user and analyst about the massive increase in data volume available to epidemiological researchers. This was followed by an interview with a second epidemiologist, discussing the software tools used to support research. A question about image creation software led to a discussion about epidemiologists' preference for numbers over images – they feel images can be ambiguous and open to interpretation, and prefer to see raw numbers. This led us to the conclusion that use of visualization in hypothesis testing was unlikely to be acceptable to our users. Instead, we looked at areas of their work that might benefit from the use of visualization software and proposed the idea of a system to allow visual exploration of large datasets, for use early in the research process. This could raise new and refine partly-formed hypotheses which could be followed by classical hypothesis testing. Further exploration of users' working practices revealed that, although they considered maps to be a powerful tool, these are rarely used as existing map software is difficult for the infrequent user to use. Thus, our proposal to develop visualization software for the easy generation of epidemiological maps was informed by dialogue with users, which led to an appreciation of users' values, as well as observation of their working practices.

3.1 General Epidemiology

Our users are researching the causes and effects of obesity, particularly childhood obesity. They use large datasets collected by local health authorities (Primary Care Trusts or PCTs) over a number of years that record the height, weight, age, gender and socio-demographic data of schoolchildren.

First the data are cleaned to eliminate records which are incomplete or have data entry errors. Then new variables, such as Body Mass Index (BMI), are derived from the basic data. These derived values are used for statistical analyses to determine the differences between cohorts by analysis of variance (e.g., gender, age, socio-demographic sub-populations) and causal analysis by multiple regressions, e.g., BMI as the dependent variable is tested against age, height, growth rates and many other variables. The results are plotted by geographic area, scaled from local PCTs to regional health authorities and compared over years to examine trends.

Requirements specification
Requirements are specified as scenarios, use cases and research questions. Use cases were documented for data cleaning, data smoothing, statistical analyses and results presentation.

Research Questions
Research questions ranged from pre-determined questions based on current literature and emerging findings to more opportunistic questions arising from exploration of data or compelling ideas. Examples of questions are:

- Is low-income a risk for obesity in women living in rural as well as urban areas?
- Is obesity varying throughout childhood less than it did 10 years ago?

Such questions were derived from a variety of sources, combining to provide insight into the different stages in the lifecycle of an epidemiological research question. First, observation of meetings when our epidemiologists discuss upcoming work and new areas of research provided an understanding of the process of shaping new research questions. This includes evaluating which questions are worth pursuing, identifying information sources, deciding which variables are significant and how to phrase questions. Second, we followed the progress of individual research questions, from the initial stages of

modelling and design, through data cleaning, statistical analysis, to write-up and presentation of results. This was achieved via a combination of direct interviewing and discussions, observation, and analysis of statistical software scripts. Finally, analysis of scripts and published epidemiological research provided 'polished' research questions. This use of a combination of information sources provides a broader insight into the types of questions asked, and the formulation of these questions, than any single technique could have done. Furthermore, supplementing interviews and observation of working practice with written material and meeting observations was a practical approach to solving the problem of limited user availability.

Sublanguage analysis

As can be seen from the above examples, epidemiological research questions are often complex queries which need to invoke statistical analysis scripts rather than simple database queries. Some questions require results displayed in graphs, whereas others need maps. In addition to the core questions, opportunistic questions evolve as new hypotheses are explored, e.g., investigating the distribution of obesity in areas with high and low population density, or examining the association between obesity and asthma in children. The initial sublanguage analysis categorises questions into types before proposing an interrogative grammar for the domain. In Epidemiology the core questions are causal, driven by the EOC (exposure-outcome-confounding) conceptual model articulated by our domain experts. Causal questions have a generic format:

- What is the effect of <exposure> on <outcome> in the defined study population?

This question can be specialised into temporal trend questions by additions of sub clauses for effect of exposure over time. A wide variety of question specialisations are generated by substituting variables as potential causes as illustrated by core questions illustrated earlier, e.g. effect of poverty/gender on consequence <obesity>. Comparative questions might compare groups over time, between places or between types of people – the classical 'time-place-person' sieve of epidemiology e.g.

- Are there differences between the increase in BMI in children from impoverished and normal backgrounds?
- Where are the differences in BMI at age 10 between areas in the North West?

The variety of possible specialisations in causal questions is vast since any combination of variables in the data set could be important (age, gender, socio economic status, etc), derived variables are often tested (e.g., growth rates, population densities) or linked to external data sets (e.g., healthy eating promotions in schools).

Design exploration

Since opportunistic questions can never be captured completely during requirements analysis, we have to adopt an approach that allows users to develop new questions themselves. In the current prototype, the solution has been to provide controls so users can pick any variables in the data to query, combined with user interface controls for specifying ranges, and attribute values. The users' requirements were translated as a set of user interface operations which allow various questions to be asked by selecting buttons and moving sliders to set ranges. Storyboards and simple paper prototypes were used to explore ideas for the interface in collaboration with users (see figure 2), and developed into a web based prototype (see figure 3).

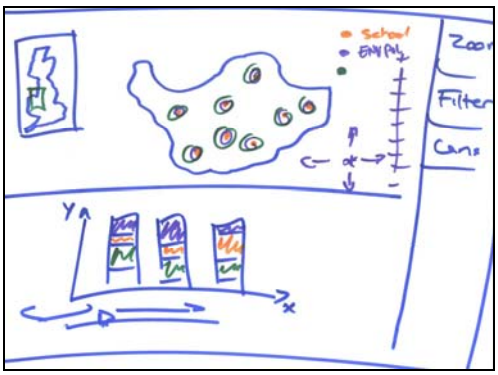

Figure 2. Early storyboard showing a simple map, graph and controls.

Questions can be posed in a flexible manner by selecting the appropriate variables to analyse. Other requirements, not illustrated in figure 3, support the application of filters, and smoothing algorithms whereby outliers in datasets can be interactively eliminated by moving quintile range dividers on the graphical display of a dataset distribution. Different data transformation and distribution smoothing algorithms can be applied; however, incorrectly applying these can lead to misleading results, so the system has an advisor module which warns less expert users when application of particular

algorithms may result in distributions that could produce doubtful results.

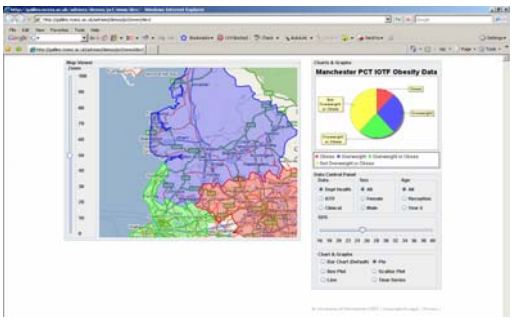

Figure 3. PCT prototype interface, showing a map and graph display with UI controls for selecting data variables and value ranges.

3.2 Genetic Epidemiology

Our users are rheumatoid arthritis researchers who are interested in the genetic determinants of the disease and its epidemiology. Recent changes in technology have enabled large scale genetic analyses — their current study investigates half a million loci of known genetic variation in 1500 arthritis sufferers vs 1500 controls. This application raises challenges of multi-level visualisation since the researchers wish to find genes associated with arthritis, link each gene to metabolic pathways associated with arthritis and then investigate the population for possible environmental influences, e.g., age, occupation, etc. The analytic workflow starts by comparing the DNA at sites of known genetic variation, SNPs (single nucleotide polymorphisms) for the control and disease populations. Many differences will be found between the populations, of which some will be false positives. The researcher uses statistical tests and contextual genetic and metabolic information to establish which SNPs are genuinely associated with the disease.

Requirements specification
The use case for this application and the user interface storyboard in figure 3 is as follows:

Goal: to identify a genetic region of interest that may be a cause of Rheumatoid Arthritis.
Actors: Researcher in Genetics
Assumptions: Data is supplied in the correct format, with Refseq identifiers, and p values calculated using 3 different algorithms (Allelic, Dominant, Trend models)
Steps
1. The user browses the file of interest on the desktop and selects the file for upload. 2. The application reads in the file and downloads information about each SNP (chromosome, location) and generates a chromosome overview display showing significant SNPs. 3. The user views a map of the SNPs plotted against the chromosome. 4. The users want to find SNPs which are in close proximity defined by within N thousand base pairs with p values lower than a set threshold. To home in on these SNPs, the user reduces the cut-off for both the p-value and the distance between pairs. 5. The user now spots a cluster of SNPs on chromosome 7 which all show a high degree of association. 6. The user zooms to the region of interest, using the 'Zoom' control – descending through several levels of magnification to isolate the area of interest. 7. The user is able to see 8 SNPs aligned along and around an unfamiliar gene. 2 of the SNPs map upstream to a putative promoter region, 6 are distributed along the length of the gene. 8. The user selects and copies all 8 SNP identifiers to the clipboard, and presses 'Map to Pathways' to try to learn more about the gene that has been isolated. 9. The application maps the 8 selected SNPs to their associated gene and then to the protein product. The KEGG diagram for the pathway is shown with the enzyme of interest highlighted.

The research questions in this application are linked to the analysis workflow, although more opportunistic questions might occur for population-level analysis of epidemiological influences. At the gene level, questions are driven by the calculated probability of a statistical association (i.e., P value):

- Which SNPs have P values < 0.05, 0.01, 0.001 for disease association?
- Where are significant SNPs located in a gene / gene region?
- Which metabolic pathway(s) component(s) does the SNP(s) map to?

Storyboard visualisations of the SNP to gene and metabolic pathway mappings are illustrated in figure 4.

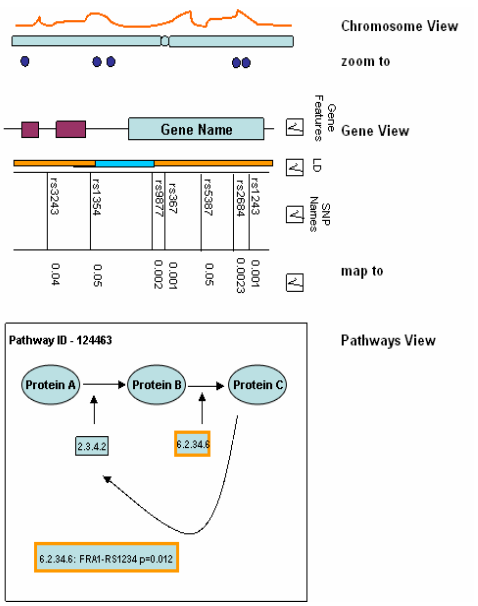

Figure 4. Storyboard visualisation for the Genetic Epidemiology application.

Three views are given for SNP analysis. First, the locus of significant mutations in the population is shown at the chromosome level. SNPs are coloured based on their P-value; by default, only SNPs with a P-value of 0.05 or less are shown. The chromosome view can be zoomed in to show regions of a particular gene with the identity of individual base pair mutations (SNPs). Colour coding denotes the P value of the SNP while the location structure of the SNP on the gene map enables interpretation of whether the mutation is in an active area of the gene or not. The graphs aligned to the gene map show the proportion of the population with the mutation. Software for the DNA database processing and statistical analysis already exists, so our requirements are to create interactive visualisations and controls to allow users to explore the databases based on the above research questions. This involves applying HCI principles to specify the user interface.

User Interface Specification

According to the use case, the user needs to interactively find suitable cut-off points for the distance between pairs and their p-value. This scenario points to the following Interactive Visualisation Claim:

Title: Closely coupled query controls and visual feedback of query results help exploration of datasets (Ahlberg & Shneiderman, 1994).
Upsides: results appear as soon as the query is submitted, linking the question directly to the answer. Values in queries can be changed iteratively to support sensitivity analysis.
Downsides: doesn't work well with discrete variables; needs user interface controls to be configured to change the query.
Scenario: The user moves sliders to specify the proximity range from 100 to 10,000K base pairs and the P value range from 0.05 to 0.01. The system displays a scatter plot showing each SNP plotted in the axes of base pair proximity 100-10,000 with 100K intervals and P values with 0.001 intervals. Each SNP is displayed with its ID colour coded for functional category. The researcher homes in on a cluster of SNPs in the range 500-600K proximity and P 0.02-0.01 where the colour coding shows a separation of introns and exons in the proximity range all associated with a suspected gene.
Example Design: Spotfire visualisation tool www.spotfire.com.

This claim has been adapted from its origins with a scenario that is more meaningful to our users and it illustrates how HCI design knowledge is imported into the requirements process. The design exploration considers using commercial software such as the Spotfire product, or developing more ambitious solutions linking scatterplot displays with gene and metabolic map displays, so as the user performs sensitivity analysis with slider controls, the display in all three views changes concurrently, such that the user can view SNP proximity distributions, loci on gene maps and implicated metabolic pathways. Research interactive, multi-view visualisations are part of the ADVISES research agenda that we incorporate in the requirements analysis process.

The visualisation claim can be reused in the design of results presentations for any query interface where continuous variables are present in questions. Interactive visualisation is suitable for dataset exploration because the slider enables sensitivity analysis; results 'pop out' in displays as soon as the slider is changed. However, some research questions do not have variables with continuous distributions so radio button controls are used instead. Other design trade-offs concern how many attributes to code on the results display and how to structure the display so the users can interpret the data easily. We reuse visualisation design principles (Ware, 2000; Tufte, 1997; Wilkinson, 2004) which recommend that position in a graph can be used to code two variables, with colour and shape coding representing another 2-3 variables. The claim in this case explains the trade-offs between more coding aiding quick visual

interpretation of the data, with the downside of the user having to learn the coding.

4. Lessons Learned

In requirements analysis, we found that interviews and requirements exploration sessions with storyboards and prototypes were effective. Even though one of the research team was co-located with the medical researchers, communication between our epidemiologists was difficult to observe as it mainly consisted of frequent short emails, instant messenger messages and corridor conversations. Consequently most information was gathered by interviewing and storyboard/scenario analysis with some additional understanding from observation of meetings and work practices. In particular, the scenario-based design approach we adopted enabled us to gather detailed requirements in a relatively short space of time.

The combination of requirements analysis with design exploration via storyboards and prototypes enables us to address the frequently encountered problem of 'I know what I want when I see it'. The approach also economised on the limited availability of our domain experts and helped elicit a considerable amount of tacit knowledge in critiques of the storyboards. Recording interviews and conversations between our domain experts provided a useful source of dialogue for extracting research questions as well as understanding the epidemiological workflow and research process.

The policy of iteratively refining requirements via storyboard analysis and prototyping helped stimulate the engagement of our users; however, we experienced some changes in focus as our users' ideas and priorities developed. As software developers, we have responded as flexibly as possible to user suggestions; this had a positive influence on collaboration during requirements analysis. Also, the ADVISES researchers had considerable prior knowledge of the genetic research domain, which helped in acquisition of the domain model and in engagement with users. Sublanguage analysis and development of a nascent ontology for epidemiological research is still in progress; however, initial results have helped to clarify domain concepts and research questions.

Our experience of reusing HCI knowledge is still in its early stages. The claims approach is useful, as it delivers knowledge in usable chunks with trade-off criteria; however, extracting new claims from the literature is difficult, but necessary, since few visualisation claims exist and design guidelines in the literature need considerable interpretation in the context of our application. For instance, there is copious advice on design of graphs and charts, but scant advice on visualisation design when a structural metaphor such as chromosome maps is needed to represent results.

5. Conclusions and Future Work

The ADVISES requirements analysis method is still under development in the sublanguage theme where we are analysing the range of research questions to create an interrogative grammar for epidemiological research and more generally for data-intensive investigation. The question-driven approach complements the GQR method of Perrone et al. (2006), as our method also employs use cases/scenarios and research questions; however, we focus more on early design realisation via storyboards. Our method also weaves HCI knowledge into the requirements process via claims, so requirements analysis is transformed from a dialogue to discover what the users want, to a conversation between developers and users that spans their goals and research questions, as well as the space of possible design to enhance their research practices.

One insight from the research questions approach is that complete requirements capture of all possible questions in data-intensive research is not possible; instead a flexible, adaptable query interface is necessary. While we anticipate that the sublanguage analysis will enable users to compose research questions by selecting and combining variables in their datasets, in many cases new multivariate questions arise. For example, if population density is not explicitly coded in the dataset, derived values have to be calculated. In our future work, we are extending requirements towards end-user development tools (Fischer et al., 2005) to enable users to extend the query interface and its sublanguage.

We found scenario and storyboard analysis helpful in provoking and facilitating discussions about users' thought processes. In epidemiology, there is considerable expertise required in separating out causal from non-causal (false positive) associations through a mixture of statistical analysis and reasoning about mechanisms, confounding and bias. We are only just beginning to gain insight into this expertise and there appears to be no substitute for interviews and analysis of expert conversations for gathering deeper insight into the domain.

The problem for all e-Science requirements analysis is where and when to stop. This can be illustrated by reference to a question which is fundamental to e-Science: which elements of the research process can be handed over to automated, software support, and which elements should be the preserve of human expertise. One solution to this dilemma may be found in the functional allocation literature which uses simple heuristics to discriminate between the kinds of tasks that machines are good at (e.g., error-free repetition) and those humans are good at (e.g., problem solving) (Dearden et al., 2000). However, we argue that the use of techniques such as function allocation must not be at the expense of a willingness to pursue an iterative and user-centred approach to requirements analysis and software development which enables researchers to define for themselves where the boundary between automation and the exercise of human expertise should be drawn.

In conclusion, the ADVISES experience has explored new avenues in e-Science requirements analysis; however, we have also raised more, as yet, unanswered questions.

References

Ahlberg, C. & Shneiderman, B. (1994). Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In B. Adelson, S. Dumais, & J. Olson, (Eds). In *Human Factors in Computing Systems: CHI 94 Conference Proceedings, Boston MA*, (pp. 313-317). New York: ACM Press.

Beck, K. (1999). Embracing change: Extreme programming explained. Addison Wesley

Beckles, B. (2005). User requirements for UK e-Science grid environments. In *Proceedings of UK All Hands E-science meetin*g, Nottingham, September.

Carroll, J. M. (2000). *Making use: Scenario-based design of human-computer interactions*. Cambridge MA: MIT Press.

Carroll, J. M., Rosson, M. B., Chin, G. & Koenemann, J. (1998). Requirements development in scenario-based design. *IEEE Transactions on Software Engineering*, *24*(12), 1156-1170.

Dearden, A., Harrison, M. & Wright, P. (2000). Allocation of function: Scenarios, context and the economies of effort. *International Journal of Human-Computer Studies*, *52*(2), 289-318.

DSDM. (1995). *DSDM Consortium: Dynamic Systems Development Method*. Farnham Surrey: Tesseract Publishers.

Fischer, G. & Giaccardi, E. (2005). Meta-design: A framework for the future of end user development. In H. Lieberman, F. Paternò, & V. Wulf (Eds.), *End user development* Dordrecht: Kluwer Academic Publishers.

Hallett, C., Scott, D. & Power, R. (2005). Intuitive querying of e-Health data repositories. In *Proceedings of the UK E-Science All-Hands Meeting*, Nottingham, UK.

Jirotka, M., Procter, R., Hartswood, M., Slack, R., Coopmans, C., Hinds, C. & Voss, A. (2004). Towards understanding Requirements for e-Science: the eDiaMoND case study. In *Proceedings of UK All Hands E-science meeting*, Nottingham, September.

Kell, D. & Oliver, S (2004). Here is the evidence, now what is the hypothesis? The complementary roles of inductive and hypothesis-driven science in the post-genomic era. *BioEssays26(1), 99-105.*

Perrone V., Finkelstein A., Goldin L., Krammer, J. Goldin L., Parkinson H. and Reddington F. (2006). Developing an Integrative Platform for cancer Research: A requirements engineering perspective. In *Proceedings of UK All Hands E-science meeting*, Nottingham, September.

Sommerville, I. & Sawyer, P. (1997). *Requirements engineering: A good practice guide*. Wiley.

Standish Group, Standish group CHAOS reports (2007). www.standishgroup.com/reports

Sutcliffe, A. G. (2003). Scenario-based requirements engineering. *Proceedings: IEEE Joint International Conference on Requirements Engineering,*. Los Alamitos CA: IEEE Computer Society Press.

Tufte, E. (1997). Visual explanations: images and quantities, evidence and narrative, Graphics Press.

Ware, C. (2000). *Information visualization: Perception for design*. San Francisco: Morgan Kaufmann.

Wilkinson, L. (2004). The Grammar of Graphics. Springer, 2nd Edition.

Zimmerman, A. and Nardi, B. (2006). Whither or whether HCI: requirements analysis for multi-sited, multi-user cyberinfrastructures. In Proceedings of the ACM Conference on Human Factors in Computing Systems, Montreal, ACM Press, New York, 1601-1606.