

More sense from audit trails: Exploratory sequential data analysis

Terry S. Judd and Gregor E. Kennedy

Biomedical Multimedia Unit

The University of Melbourne

This paper introduces the use of exploratory sequential data analysis (ESDA) to detect, quantify and correlate patterns within audit trail data. We describe four sequence analysis techniques and use them to analyse data from 34 students' attempts at an interactive drag and drop task. Using a model sequence of events based on the task's underlying educational design as reference, we employed these techniques to: (i) calculate an 'average' sequence of events based on individual user sequences, (ii) characterise individual sequences in terms of their similarity to the design model, (iii) identify common partial sequences within individual sequences, and (iv) characterise transitions between two disparate actions within the task. We then used the results of these analyses to explore why most students failed to complete all components of the task. We suggest that it was not because the task was too long or that it lacked challenge but that students intentionally and selectively ignored certain non-key steps in the task. It is our contention that ESDA techniques, in conjunction with judiciously collected audit trail data, represent a powerful and compelling tool for educational designers and researchers.

Keywords: audit trails, sequence analysis, multimedia, interactivity

Introduction

The utility of audit trail analysis continues to attract scepticism, much of it directed towards issues associated with data management and analysis and interpretation (Reeves & Hedberg 2003). With respect to data management, audit trail systems can undoubtedly capture huge amounts of data, the bulk of which may be irrelevant to the investigation at hand. However, given our ability to capture, store and manage and retrieve increasingly large amounts of data using software agents, network protocols and scalable databases, physically managing the sheer volume of data generated is not the problem it once was. In addition, not all data capture systems are as verbose and produce data as abstract as server logs or keystroke capture agents but can instead be tailored to collect only those events and information requested by the investigator. With such targeted systems the volume of data captured is substantially reduced and internal re-coding of the captured data can greatly improve its accessibility (Judd & Kennedy 2001). In essence, data management is no different than any other system design issue. Sound design leads to the development of robust systems and good data management practices.

With respect to the analysis and interpretation of audit trail data, Reeves and Hedberg (2003) state that

The analysis of audit trail data within complex multimedia or hypermedia programs is especially challenging. When learners can go wherever they want in any sequence, the possibility of detecting interpretable paths without the input of learners becomes almost impossible (p. 182).

While acknowledging that audit trail analysis is often complex, and sometimes challenging, discerning meaning is by no means impossible. Audit trail data, like many other types of data based on behavioural observations are open to analysis through a wide range of statistical and numerical techniques. Although undoubtedly helpful, the use of external measures to supplement or inform audit trail analysis is by no means mandatory and meaningful patterns of usage can be derived in their absence, particularly where the sample population is large. Moreover, a sound understanding of how users interact with a given multimedia or hypermedia environment is essential *before* we can begin to consider why they interact with it in various ways. By way of example, in a recent paper (Kennedy & Judd, 2004) we described an audit trail analysis of students' usage of a multimedia program designed to assist medical students develop sound interviewing techniques. Students use the program to construct a virtual interview between a doctor and a patient, represented to the user as a series of audio and video clips and internally (within the program's logic) as a decision tree. The audit trail data captured included a sequential record of which

'nodes' in the decision tree the user visited and the various actions they engaged in at each of these nodes. Simple descriptive statistics of counts and times provided valuable insights into how students were accessing the program and utilising its principal features. A numerical clustering technique was then employed to resolve differences between students' use of a key subset of these features. We subsequently identified four distinct categories of users that encapsulated different practical (and by inference conceptual) approaches to using the program. These various approaches were manifest as distinct navigational paths within (but not necessarily between) the various nodes of the decision tree. The robustness of these categories has subsequently been confirmed by independent analysis of additional audit trail data collected from the program over three successive years (Kennedy & Judd, unpublished data). We were able to detect these differing navigational paths using simple counts of user actions as input because of several key aspects of the program's underlying learning design. These include the fact that users visit many nodes during the course of an interview, that there are a limited number of actions available to users at each node and that the order in which these actions can be performed is fixed. However, analysing simple counts of user actions cannot assist us to recognise patterns within the sequences of nodes visited by individuals or groups of users or, in a wider context, to recognise and interpret navigational paths within more complex or less structured multimedia or hypermedia environments or tasks. In such cases, we need to employ specialised techniques to detect, compare and contrast sequences of user actions, events or paths we may have captured using audit trails.

Sequence analysis remains a largely untapped field of investigation by multimedia and hypermedia researchers. However, the rise of genetic analysis, in particular, has led to the development and adoption of powerful techniques (Kruskal 1983, Lange 2002) for analysing sequential data, some of which have filtered through into the domain of human computer interaction (HCI). Of particular relevance to us is the application of various sequence analysis techniques, often referred to as exploratory sequential data analysis or ESDA, within the field of usability testing (Sanderson and Fisher 1994, Hilbert and Redmiles 2000). Hilbert and Redmiles (2000) identify three main categories of ESDA techniques: (i) sequence detection – techniques for detecting occurrences of defined target sequences within source sequences, (ii) sequence comparison – techniques for measuring correspondence between source and target sequences, and (iii) sequence characterisation – techniques for constructing abstract models from source sequences. Techniques drawn from each of these categories are potentially useful to researchers working with audit trails captured from multimedia and hypermedia environments. They might, for example, be employed to characterise the order in which users visit the various sections of a modular tutorial and to then compare these with model paths based on simple presentation order, embedded user support or more abstract paths based on combinations of linked concepts. Alternatively, researchers may wish to focus their attention at a finer scale of actions such as the order in which users complete a drag and drop task or select options in a multiple choice question.

The purpose of this paper is threefold. The first is to briefly introduce a number of relatively simple ESDA techniques that we have either adapted or developed specifically for use with audit trail data. The second is to demonstrate the utility of these techniques by applying them to the analysis of actual audit trail data collected during students' attempts at a relatively complex interactive task. Finally, in a previous study (Kennedy & Judd 2000), we reported various analyses of this task based primarily on simple usage statistics. However, our interpretation of these results, particularly with a view to reconciling user actions and designer expectations, was limited by our inability to analyse key sequential components of the user data. This paper attempts to redress that limitation.

Methods

The software environment

For the current study, we collected audit trail data from a multimedia program used by Medical students at the University of Melbourne. Medical Genetix (Metcalf 2003) is a CD ROM based program dealing with biomedical and clinical aspects of various genetic disorders and is used as a non-compulsory, self directed learning resource within a problem based medical curriculum. Students have ready access to this resource through their faculty's main computer laboratory and tutorial rooms.

Medical Genetix includes of a number of modules and sections that users can freely navigate between via a series of interface tabs. We were especially interested in students' use of a single section of the program

(Family Histories and Pedigrees) within a single module (Cystic Fibrosis) in conjunction with its recommendation as a learning resource for a 'problem of the week' dealing with cystic fibrosis. This section includes a 'drag and drop' task requiring users to complete a genetic pedigree, based on a supplied family history, by dragging tiles that reflect family members' sex and genetic status to 20 empty positions in the pedigree. Incorrect drags are immediately rejected. We created a modified version of Medical Genetix that included the audit trail system and configured the system to collect detailed data relating to students' usage of the Cystic Fibrosis module in general and the pedigree task in particular. This version was installed in Medical Faculty's main computer lab, and data was collected over a one week period during 2000, coinciding with its recommendation as a learning resource.

In an earlier paper (Kennedy & Judd 2000) we revealed that a total of 78 students accessed the program, 49 entered the Family History and Pedigree section of the Cystic Fibrosis and of those 34 attempted the pedigree task. Of these 34 students, all placed a minimum of 10 tiles and correctly placed 15.6 tiles on average. Interestingly, only one student successfully placed all 20 tiles (Judd & Kennedy, 2000). Moreover, our analysis of users' final three tile placements revealed that students were not abandoning the task due to lack of success or frustration – 97% of users correctly placed their final tile and 77% correctly placed their final three tiles. These results led us to suggest that either the task was either too long or insufficiently challenging to maintain students' interest to completion (Kennedy & Judd 2000).

For the current study, we undertook an analysis of the sequence in which students correctly placed the various tiles when completing the pedigree. In doing so, we were particularly interested in assessing how closely students adhered to the presented history. More specifically, was a student's degree of adherence to the presented history related to any failure to complete the task and if so, how? Were students tackling the task in the same (or similar) ways, or were they adopting different strategies?

The audit trail system

The macro design of the system we used to collect the audit trail data is described in detail in Judd and Kennedy (2001). Briefly, it is an embedded system that is installed and configured during the target software's development. For our purposes, the advantages of this approach – targeted data collection and integrated data management – strongly outweigh the perceived benefits (chiefly portability) of software agents that rely on keystroke capture. The system employs both sequential and object based data capture and management, allowing targeted events to be logged to a simple sequential record and/or used to update user defined data objects each with its own set of properties and methods (e.g. timers, counters, strings, histories etc). Data logged by the system is converted to XML format and either stored locally on the user's computer for later retrieval or emailed to the developer/researcher. A dedicated parsing application is then used to process the retrieved XML files and prepare the data for further analysis.

Sequence analysis techniques

We have recently developed or adapted a series of analytical techniques for detecting, characterising and analysing sequences in audit trail data. Four of these techniques have been employed in the current study and are described below (detail of descriptions is limited by size constraints).

Model development (technique A)

Technique A employs a number of sequence detection and characterisation routines loosely based on Fisher's cycles (Hilbert & Redmiles 2000) to derive an objective model sequence of prescribed events from a number of source sequences. Briefly, it takes a sequential list of events (items), such as the names of visited screens, for any number of users and through a repeated series of transformations, calculates a ranking for each item, where rank indicates the position of the item in the model sequence. Identical rankings are resolved through a series of pair wise comparisons of the equivalent items using the percentage of occurrences of one of the items either before or after the other as input (higher percentage = higher ranking). The resultant objectively derived sequence represents an 'average' sequence of events for all users. While taking repeat instances of events into account in the calculation of ranks, these are not represented in the final model. Models derived via this technique can serve as a basis for comparison with actual sequences or model sequences derived by alternative means (e.g. presentation order, design rationale).

Sequence comparison (technique B)

Technique B employs sequence comparison routines in which source and target sequences are reconciled via three types of string transformations (deletion, insertion and translation) and is based on a technique termed 'process validation' developed by Cook and Wolf (1997). It takes as input one or more actual sequences of events (source sequence) and a corresponding model sequence of events (target sequence) derived through technique A or some other means. Deletions are used to remove unwanted multiple instances of events from source sequences while insertions replace missing events that are present in the target sequence. A series of translations are then applied to convert the modified source sequence to the target sequence. Each transformation is optimised so as to require the minimum number of operations. For example a missing block of four adjacent items is treated as a single insertion. Source sequences can then be compared to a target sequence based either on their values for individual transformation types or custom metrics combining the values of two or more transformation types (e.g. distance = a x insertions + b x translations).

Sequence detection (technique C)

Technique C takes as input a series of source sequences and a specified sequence length. Each source sequence is processed iteratively to extract all (consecutive) sequences of the specified length. For example, the simple sequence ABCDE would yield four sequences of length 2 (AB, BC, CD, DE) three sequences of length 3 (ABC, BCD, CDE), two sequences of length 4 (ABCD, BCDE) and a single sequence of length 5 (ABCDE). As each source sequence is processed, any unique sequence is added to a list of target sequences. At the end of this process, all target sequences are matched against all source sequences to determine their frequency of occurrence across the entire sample. Sequence 'chunks' identified in this way can be matched against model sequences created via technique A or some other method or can be used as 'building blocks' in the generation of new models.

State transitions (technique D)

Technique D employs a sequence characterisation approach based on Guzdial's (1993) adaptation of Markov chain analysis. Its purpose is to derive process models that describe the probabilities of transitions between events (states). Unlike the previous techniques, which can accommodate large numbers of unique events, this technique is best limited to a few key events. This limitation is practical rather than technical as although state transition data can be simultaneously calculated for many events, the presentation and interpretation of this data relies on graphical representations that become increasingly difficult to render as the number of states increases. To work within this limitation it is generally best to either chunk related events prior to analysis or restrict the analysis to a smaller number of higher level events (e.g. investigate user movements/transitions between major sections of a program rather than between screens across multiple sections). The resulting state diagrams efficiently summarise the probability that a user will move from an event to any other event, including reselection of the originating event.

Data analysis

Model development

For the first part of this analysis, a model sequence of tile placements in the interactive pedigree task of Medical Genetix was created based on the simple order of mention of named characters in the provided family history (see Appendix for a full transcript of the history). This model (I) accurately reflects the task's underlying educational design – the order in which student could reasonably be expected to complete the task. A second model (II) representing the 'average' sequence of tile placements was then calculated from the students' actual sequences using technique A. The two models are presented in Table 1. There was a high level of association between the two models. Reconciling the two models via technique B required only two single item translations, corresponding with the differential rankings of Mark (19th vs. 5th) and Tanya (13th vs. 16th) (Table 1).

Sequence comparison

All 34 user sequences were also reconciled with the model I sequence using technique B. With a single exception (the lone student who correctly placed all 20 tiles) all of these transformations required a combination of insertions and translations, the results of which are presented as a combined density plot in Figure 1. Although the technique is able to accommodate deletions, none were required as individual tiles can only be placed once in the pedigree.

Table 1: Model sequences of tile placements

Models are based on (I) order of mention in the family history and (II) actual user sequence using technique A.

Order	Model		Order	Model	
	I	II		I	II
1	Mathew	Christine	11	Mathew	Andrew
2	Paula	Andrew	12	Paula	Anne
3	Julia	Anne	13	Julia	Tanya
4	Damian	Colin	14	Damian	Colin
5	Margaret	Debbie	15	Margaret	Debbie
6	Mark	Tanya	16	Peter	Mitchell
7	Peter	Mitchell	17	Susan	Melanie
8	Susan	Melanie	18	Joseph	Alan
9	Joseph	Alan	19	Sally	Mark
10	Sally	Fiona	20	Christine	Fiona

The majority of user sequences required at least four insertions (mean = 3.9, median = 4) and three translations (mean = 3.1, median = 3) to be reconciled with model I. Given that the average user successfully completed 15.6 tile placements (Kennedy & Judd 2000), these data confirm that most insertions (i.e. omissions by the student) involved individual tiles (corresponding to a single mention of a name in the history) and not groups of tiles (corresponding to a phrase, sentence or paragraph in the history) as the latter would be treated as single insertions. The interpretation of the translation data is more complex and is addressed in part in the following section.




Figure 1: Sequence transformations

Values represent the frequencies of insertions and translations required to derive model sequence I from the 34 source sequences. Area of circles is proportional to the number of source sequences.

Sequence detection

Technique C was used to create partial sequences of between two and ten items in length using the 34 user sequences as input. Partial sequences of nine items or greater in length were not considered further as none were common to two or more user sequences. Large numbers of distinct partial sequences were generated from the user sequences despite the high level of agreement between the 'average' user sequence (model II) and the model I sequence (Table 2). For example, if all user sequences conformed completely to a specific model (e.g. model I), we would detect only 19 partial sequences given that the maximum user sequence is 20 – similarly we would detect 18 partial sequences of length three, 17 of length four and so on. However, 135 distinct partial sequences of length two out of a theoretical maximum of 380 (randomly generated) partial sequences were detected, with over half of these occurring in two or more of the user sequences. The high number of partial sequences of length two strongly influenced the number of partial sequences of all other lengths that were detected, with totals ranging up to 345 for sequences of length five (Table 2).

Table 2: Statistics for partial sequences derived via technique C

Total = the number of discrete partial sequences recognised across all user sequences. Common = the number of identified partial sequences occurring in two or more user sequences. Maximum frequency = number of instances of the most common partial sequence across all user sequences.

Sequence length	Total	Common	Maximum frequency
2	135	72	21
3	255	85	12
4	328	50	6
5	345	31	6
6	340	16	3
7	317	8	2
8	290	1	2

Table 3: Grouping of names within model I sequence based on sentence structure

Names in italics are duplicate mentions. Square brackets contain the character's genetic status (xx = affected, xo = carrier, oo = unaffected, - = unknown status) according to the history and the percentage of students who placed the corresponding tile in the pedigree

Paragraph	Names
1	Mathew[xx, 73.5]
2	Paula[oo, 88.2] Julia[xx, 76.5]
3 a	Damian[--, 91.2] Margaret[--, 76.5]
b	<i>Damian</i> , Mark[--, 67.6] Peter[xo, 73.5]
4 a	<i>Peter</i> , Susan[xo, 85.3]
b	<i>Peter</i> Joseph[oo, 73.5] Sally[oo, 79.4]
5 a	<i>Peter</i> , <i>Sally</i>
b	Christine[oo, 79.4] Andrew[oo, 82.4]
6	<i>Margaret</i> Anne[oo, 61.8] Colin[--, 73.5]
7 a	<i>Colin</i> Debbie[--, 73.5]
b	Tanya[xx, 64.7] Mitchell[oo, 79.4] Melanie[xo, 76.5]
8	<i>Margaret</i> Alan[xo, 79.4] Fiona[--, 67.6]

Several of the most common partial sequences of various lengths were then matched against model I. To aid in this comparison, model I was segmented on the basis of the underlying sentence structure of the family history (Table 3), with names grouped according to which paragraph, sentence or phrase they were mentioned in. Of the most commonly occurring partial sequences of length two, all corresponded to joint mentions within individual phrases (e.g. *Damian and Margaret*, *Peter and Susan*, *Alan and Fiona* – see Appendix). The most common partial sequences of length three followed a similar pattern of distribution within the model including examples such as *Mathew*, *Paula and Julia* and *Sally*, *Christine and Andrew*. Of the longer partial sequences, the most common were those occurring within and across the first four paragraphs although these did not necessarily follow model I exactly, with variations including simple deletions such as the omission of Mark (paragraph 3b) and translations such as the order reversal of Paula and Julia (paragraph 2a) (Table 3).

State transitions

Technique D was employed to investigate transitions by users between the history and the drag and drop task. The number of instances that users either (a) followed a drag with another drag, (b) followed a drag with a visit to the history, (c) followed a visit to the history with a drag, or (d) followed a visit to the history with another visit to the history, were recorded for all individuals. These values were then used to calculate average probabilities for each of the four possible transitions (Figure 2). Table 4 provides some supporting statistics. We also investigated whether these transitions were influenced by the success or failure of any given drag (unsuccessful drags accounted for 20% of all attempts). Interestingly, users were

less likely to visit the history following an unsuccessful than a successful drag ($p = 0.53$ vs. $p = 0.32$; see Figure 2b and 2c). A possible explanation of this behaviour is that users typically attempt to replace a tile, utilising retained knowledge, following an incorrect drag but seek additional information from the history before attempting to place the next tile in the sequence.

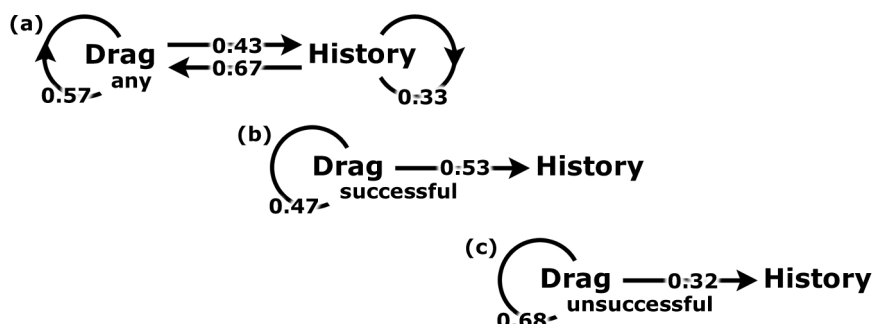


Figure 2: State transition diagram

Options represent the probability that for any action a user will either change actions (from drag to history or vice versa) or reselect the same action for any drag (a), following a successful drag (b) and following an unsuccessful drag (c). NB. History to history and history to drag probabilities are identical for (a), (b) and (c).

Table 4: Descriptive statistics for drags and history visits

Consecutive drags and history visits values represent the average number of consecutive drags or history visits (i.e. sequence length) across all users, not their incidence.

Variable	mean	stdev
unsuccessful drags	4.0	2.7
successful drags	15.6	2.0
total drags	19.6	3.4
unsuccessful/total drags	0.20	-
consecutive drags	2.2	-
history visits	17.9	6.2
consecutive history visits	1.6	-
drags per history visit	1.3	-

Discussion

We have introduced four ESDA techniques as suitable tools in the analysis of sequential audit trail data captured from multimedia or hypermedia environments. All are simple to apply provided the source data is in an appropriate format and we have provided examples of how the results of each can be presented and analysed further to facilitate interpretation. Our results also demonstrate the advantages of employing a combination of techniques within the same dataset. For example, our application of technique A (model development) led to the development of a model that was almost identical to one based on order of mention (Table 1), which viewed in isolation might suggest that individual users complete the task to a similar degree and in similar ways. However, non-sequential analysis of the data revealed considerable variation in both degree and success of completion of the task (Table 4, Kennedy and Judd 2000) while sequence analysis demonstrated a substantial level of deviation of individual sequences from the order of mention model (Tables 2 and 3; Figure 1).

In a previous paper (Judd and Kennedy 2000) we suggested that students typically failed to complete the pedigree task because it was either too long or insufficiently challenging. However, this conclusion – which was based only on counts of correct and incorrect tile placements – is not supported by our analysis of the sequence in which individual students placed tiles within the task. If the task was too long and students were failing to complete it as a result, we would have expected them not to have placed one to several of the last mentioned characters in the family history. This was not the case. At least half of the students completed at 16 or more (out of a possible 20) successful drags and the tiles they omitted from the pedigree were drawn from all parts of the history – even the first mentioned name was omitted by

more than 25% of students (Table 3). Of the four most frequently omitted tiles (indicated by asterisks in Figure 3), only the last mentioned character in the history (Fiona = 20) was linked to any subordinate branches in the pedigree and, even then, the tiles of each of the children of this character could have been correctly placed without reference to Fiona's genetic status. Furthermore, a character's declared genetic status appeared to be unrelated to whether that person's tile was or wasn't placed. Of the four most frequently omitted tiles, two had undeclared status, one was unaffected and one was affected (see Table 3). All of which suggests that it was a character's/tile's spatial and/or conceptual relationship to other characters in the history rather than their order of mention that was critical in determining whether they were successfully placed in the pedigree.

It is more difficult to ascertain with certainty whether the task lacked challenge. However, completing the task should have been a relatively straightforward process for any student with a basic understanding of the inheritance of recessive traits and the use of pedigrees (readers can judge the difficulty of the task for themselves by referring to Figure 3 and the appendix). Most students attempted the task with few errors (Table 4) and at least some of those can be accounted for in either ambiguities in the history or unfamiliarity with task protocols (eldest sibling at left; males at left in isolated couples – see Figure 3). As is the case if the task was too long, we might have expected to see students leaving the task sooner, if they lacked the motivation to complete it because it was too easy.

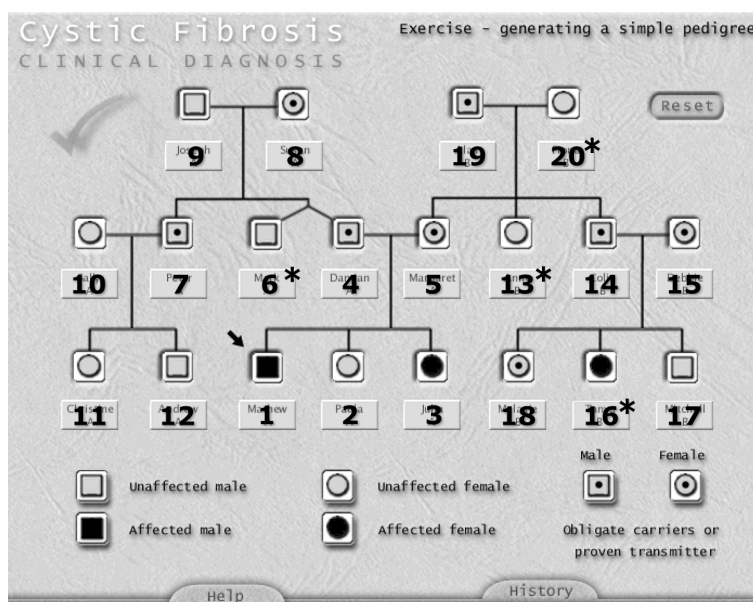


Figure 3: Completed pedigree task

Superimposed numbers represent tile placement order based on their order of mention in the family history (see Appendix). Asterisks indicate tiles with the lowest rate of placement (< 70%) across all users.

Why then did students fail to complete the task if, as we assumed, it was neither too long nor insufficiently challenging? A possible explanation is that students consciously chose to not to place certain characters in the pedigree as they worked their way through the history. As we describe above, the most frequently omitted characters tended to be either spatially or conceptually less important in the overall context of the task. That is, although students 'ignored' certain characters (by failing to place them in the pedigree) at a behavioural level, they may well have understood their status and relationship to other characters sufficiently to have, in a sense, 'placed' them at a cognitive level.

In conclusion, we believe that ESDA is a powerful and robust approach when judiciously applied to audit trail data. It promises the ability to investigate students' behavioural learning processes in great and varied detail for many types of interactive tasks. Such tasks are not limited to traditional multimedia applications and could be extended to the analysis of users' interactions in real time chat environments or discussion forums. We believe ESDA techniques will also prove to be of considerable use to educational designers seeking to verify whether interactive tasks are being used in accordance with their expectations. Good quality information of this sort is essential for informing both the effective design of new tasks and

the redesign of existing but ineffective tasks. The types of analysis we describe could also be used to provide real time interventions to users as they navigate tasks within educational technology environments. ESDA techniques open new avenues for exploring otherwise difficult to analyse data. They represent a valuable tool for those seeking to refine current ideas and generate new hypotheses in relation to user behaviour in multimedia and hypermedia environments.

References

- Cook, J.E. & Wolf, A.L. (1997). *Software process validation: Quantitatively measuring the correspondence of a process to a model*. Tech Report CU-CS-840-97, Department of Computer Science, University of Colorado at Boulder.
- Guzdial, M. (1993). *Deriving software usage patterns from log files*. Tech Report GIT-GVU-93-41.
- Hilbert, D.M. & Redmiles, D.F. (2000). Extracting usability information from user interface events. *ACM Computing Surveys*, 32(4), 384-421.
- Judd, T. & Kennedy, G. (2001). Extending the role of audit trails: A modular approach. *Journal of Educational Multimedia and Hypermedia*, 10(4), 377-395.
- Kennedy, G. & Judd, T. (2000). Pilot testing of a system of electronic evaluation. In R. Sims, M.O'Reilly & S. Sawkins (Eds), *Learning to Choose: Choosing to Learn (Short Papers and Works in Progress)* (pp.187-192). Lismore, NSW: Southern Cross University Press.
- Kennedy, G. E. & Judd, T.S. (2004). Making sense of audit trail data. *Australasian Journal of Educational Technology*, 20(1), 18-32. <http://www.ascilite.org.au/ajet/ajet20/kennedy.html>
- Kruskall, J. B. (1983). An overview of sequence comparison: time warps, string edits, and macromolecules. *SIAM Review*, 25(2), 201-237.
- Lange, K. (2002). *Mathematical and statistical methods for genetic analysis*. 2nd edition. Springer-Verlag, New York.
- Metcalfe S. (2003). *Medical genetix: Clinical and molecular aspects of human genetic disorders* (CD). Melbourne: The University of Melbourne.
- Misanchuk, E.R. & Schwier, R. (1992). Representing Interactive Multimedia and Hypermedia Audit Trails. *Journal of Educational Multimedia and Hypermedia*, 1(3), 355-372.
- Reeves, T.C. & Hedberg, J.G. (2003). *Evaluating interactive learning systems*. Athens GA: University of Georgia, College of Education.
- Sanderson, P.M. & Fisher, C. (1994). Exploratory sequential data analysis: foundations. *Human-Computer Interaction*, 9, 251-317.

Appendix: Family history from the Medical Genetix pedigree task

- Mathew A, the proband, is a 10-year-old boy diagnosed with cystic fibrosis (CF).
- He has two sisters, Paula (8) who is unaffected and Julia (7) who is also affected.
- His parents are Damian and Margaret. Damian has a fraternal twin Mark and an older brother Peter.
- Through cascading testing, Peter and his mother Susan were found to be carriers of CF, whereas his father Joseph and Peter's wife Sally were not carriers.
- Neither of Peter or Sally's two children, Christine or Andrew were found to be carriers.
- Margaret has a younger sister Anne, who is unaffected, and a younger brother Colin.
- Colin and this wife Debbie have three children: the middle child Tanya is affected with CF, Mitchell is unaffected and the eldest child Melanie was later found to be a carrier.
- Margaret's parents are Alan and Fiona and through cascade testing Alan was found to be a carrier.

Terry S Judd & Gregor E Kennedy, Biomedical Multimedia Unit, Faculty of Medicine, Dentistry and Health Sciences, The University of Melbourne. Email: tsj@unimelb.edu.au gek@unimelb.edu.au

Please cite as: Terry Judd, T. & Kennedy, G.E. (2004). More sense from audit trails: exploratory sequential data analysis. In R. Atkinson, C. McBeath, D. Jonas-Dwyer & R. Phillips (Eds), *Beyond the comfort zone: Proceedings of the 21st ASCILITE Conference* (pp. 476-484). Perth, 5-8 December. <http://www.ascilite.org.au/conferences/perth04/procs/judd.html>

Copyright © 2004 Terry Judd & Gregor Kennedy

The authors assign to ASCILITE and educational non-profit institutions a non-exclusive licence to use this document for personal use and in courses of instruction provided that the article is used in full and this copyright statement is reproduced. The authors also grant a non-exclusive licence to ASCILITE to publish this document on the ASCILITE web site (including any mirror or archival sites that may be developed) and in printed form within the ASCILITE 2004 Conference Proceedings. Any other usage is prohibited without the express permission of the authors.