Interaction Modelling for Digital Libraries

Nick Bryan-Kinns, Ann Blandford & Harold Thimbleby

School of Computing Science, Middlesex University,
Bounds Green Road, London N11 2NQ, U.K.

Abstract

This paper reports on the development of a framework to support interaction modelling in order to inform design and evaluation of interactive systems. The utility and nature of this Interaction Framework (IF) are illustrated by considering user interaction within digital libraries. In particular, we are interested in how requirements on user interfaces to information repositories, such as supporting serendipity and ease of familiarisation with the structure of the collection, can be described and understood. IF provides a means of mapping such abstract requirements onto concrete descriptions of the interaction as illustrated by examples in this paper.

1 Introduction

One of the major challenges for Human-Computer Interaction (HCI) is to deliver tools and techniques that can be used to improve the usability of particular classes of devices. HCI techniques lag behind technology, so that we currently have a reasonable battery of tools that can be applied to evaluate word processors or devices that are designed to support well-defined tasks, but our understanding of the use and usability of novel technologies is comparatively weak. Without the development of relevant theory, the design of novel technologies will remain a largely craft-based skill.

The user's experience of working with any technology – whether or not the user achieves interesting goals in a satisfying way – depends on a variety of design factors that need to be taken into account together. An integrating framework that accommodates user, device, domain and more general context-of-use concerns is needed. Interaction Framework [4] has the potential to fill this role. The focus of the project reported here is on usability of Digital Libraries, which present a range of usability challenges, and are likely to become increasingly important over the foreseeable future. When compared to the overall provision of the World Wide Web, they support a relatively well defined set of user tasks, and are typically designed with a particular user population in mind. In particular, as structured repositories of digital documents that are accessed both locally and over networks, factors that affect usability include:

- the information structure,
- network response time and the patterns of communication between users and computer systems,
- the content types of documents and the ways individual documents can be accessed, and
- the context that the user is working in.

The focus of the project is on working with all users of a library, including developers and librarians as well as end-users. Various libraries are participating in this work which provides a wide range of libraries and forms of interaction to study.

In particular, the Faculty of Health Studies at Middlesex University, which has sites at several hospitals in the North Thames region, is a member of a consortium that has obtained funding to set up an Electronic Library in Medicine. This will make publications available to staff, students and practitioners both at a central location (a new library on the Whittington Hospital site that provides both traditional and electronic documents) and via client computers located in wards and health centres. Important questions here clearly include how practitioners integrate library access with other aspects of the work and professional updating, and how relevant information can be made available in a timely manner.

Across Middlesex University, the Hybrid Library Project is taking the approach of complementing existing library provision with appropriate electronic resources for students and staff across all disciplines. This project is concerned with developing appropriate infrastructure to deliver both traditional and electronic library services to all staff and students across the University.

While these libraries will be the main focus for user studies – in order to understand better how people work with library-based material in a variety of situations – to support design activity, we are also working with the developers of the New Zealand Digital Library (NZDL; [13]). The source code for this library has been made available to us so that we can develop and test alternative interface designs.

There are two important aspects to the background of the project presented in this paper which are discussed in the following sections. The first aspect is past work on usability of digital libraries, the second is past work on the Interaction Framework that will be developed further within the project.

2 Digital libraries

There is high investment in digital libraries, much of it focused on the structuring and management of collections, on technical concerns and on issues raised by internationalisation (see [8], or [1], for a broad overview of research areas that have an impact on the development of Digital Libraries). Digital libraries provide a challenging domain of application for any usability-oriented modelling approach, since little work has been done on use and usability. In particular the questions of how the information structure, network infrastructure, interaction design and context of use interrelate, and how HCI modelling can be used to guide design and redesign, have received little attention.

To illustrate the problem, we consider here one example of a non-functional requirement that an ideal digital library would satisfy: that users should be able to find interesting documents through 'serendipity', and be able to quickly assess how interesting they are. Within traditional libraries, 'serendipity' often occurs when a user is looking for a particular book and happens to notice another interesting title on a nearby shelf. They might pick it up and check the contents page or flick through it to see if it still looks interesting. In a networked digital environment, the provision of such a capability would impose requirements on, for example:

- the information structure (the way documents are categorised and displayed),
- network response time (which determines how easy is it for the user to quickly find out more about a document),
- the way individual documents can be accessed (does the user have to download the entire document to view it? what alternative viewing mechanisms are provided?).

A device might be designed to support serendipity through careful design of these features. Alternatively, the use of a 'recommender' system, such as those used by some internet book sellers, or a 'personal librarian agent' might enable the computer system to actively support a kind of serendipity, based on document selections made by previous users. The same abstract interaction requirement can be implemented in alternative ways. The aim of the Interaction Framework is to map the abstract interaction requirements on to more concrete interactional properties discussed later; currently this mapping is under development.

Other interaction requirements can be identified; for example, that the library should support *familiarisation* (that users can rapidly gain an overall impression of what kinds of material are available and of how those materials are structured); that it should support what Hall [6] calls the *'IKIWISI principle'*: that users may not know exactly what they want, but that 'I'll Know It When I See It'; that the interaction should be *lucid* (interaction proceeds in a sane, rational, and easy to understand way with respect to agents' objectives and constraints, and their understanding of the interactional possibilities of the system); and that interaction trajectories should be *'canonical'* [4] – that is, as efficient as possible. With canonical interaction trajectories we start to see the mapping of abstract interaction requirements to interactional properties as discussed in the next section.

Certain interaction requirements may arise because of the context of use. For example, a student in a library may have an undetermined length of time to browse, and may be working with minimal interruptions, so that the quality of feedback from the system (whether that be measured in terms of the resolution of digital images downloaded, or the number of documents on a particular topic that are retrieved) may be more important than the pace of the interaction. Conversely, a nurse on a ward who is working under time pressure will need more support to deal with interruptions (*e.g.* being given reminders about the current state of the interaction), or may need to be able to enter a document request quickly, possibly returning at a later time to retrieve the document that has been downloaded and printed automatically. These requirements should guide the designer of a system to configure that system appropriately – for example, allocating functions to local client or to the host server to match the users' needs.

3 Interaction Framework for Digital Libraries

Clearly, any library that does not offer the essential core functionality is not usable. However, the factors that really determine whether or not a particular library is usable and useful go beyond basic functionality, and relate to more abstract usability properties such as those listed previously (serendipity, IKIWISI, *etc*.). There is existing work (*e.g.* Smith *et al.* [12]) that aims to relate the context of work to the functionality and task structures of the design, but we are not aware of work that specifically addresses non-functional usability requirements. The focus of the work described here is on these requirements, and the Interaction Framework offers an appropriate style of analysis for dealing with them.

Interaction Framework (IF) aims to be neutral between users and devices, and to deal with multi-agent systems. Also, it relates domain concerns, and properties of the context of use that influence interaction patterns, to device ones. Therefore, IF has the potential to offer useful insights into the usability of digital libraries, and should provide support for designers who are aiming to design and configure more usable libraries.

IF was originally proposed as an integrating approach that focuses on the interaction between users and devices, treating both as having equally important (but different) roles within the interaction. By rigorously describing abstract properties of interactions, it becomes possible to reason about how user properties and particular design decisions influence each other to yield more or less satisfactory interactions. For example, Harrison *et al.* [7] demonstrate how a notion of ‘user freedom’ can be used to generate requirements on a device design, while Barnard *et al.* [3] give an example of breakdowns in the interaction that can be accounted for by relating user properties and device design. IF is used in conjunction with other theory-based or empirical techniques; that is, it provides an integrating framework that accommodates insights from other approaches.

The framework provides a way of describing interactions at various levels of abstraction as illustrated in table 1 and discussed in following sections. It can be used both descriptively – to describe an interaction that has already happened and reason about properties of that interaction – and prescriptively – to specify requirements on the design of the interaction as illustrated by Blandford *et al* [4]. At the *fundamental level* interaction is described in terms of the agents in the interactional system, their grouping, communication channels between them, and events that are issued along these channels. Agents in the system have objectives. As illustrated in table 1, agents' objectives permeate descriptions of interaction from the abstract level of interactional properties down to the fundamental level where they are described in terms of pre and post conditions on the state of the system (the set of agents).

As mentioned previously, each level of description in the framework builds on a lower level (except for the fundamental level). As such the *trace level* builds on the fundamental level to describe *interaction trajectories* which are partially ordered events constituting an episode of communication between agents attempting to meet an objective. For example, in searching for a text in a collection (an objective) there is a sequence of events for a user attempting to find the text. We are interested in trajectories which do not meet objectives, and the reasons behind this.

Building on the trace level, the *symptoms level* categorizes trajectories in terms of whether the trajectory was *canonical* (most efficient for achieving the objective), and if not, what the symptoms of interactional *trouble* were. Troubled trajectories do not meet objectives, or do not meet the objectives in the least convoluted manner. For example, mistakes may be made by the user in issuing an event which leads to the interaction involving a detour to return to a point from which the objective can be met.

Once symptoms of interactional trouble have been identified the *causes level* is used to suggest causes for the trouble. At this point IF can use hooks to other theories and modeling techniques to provide explanations for trouble occurring. For example, models of human cognition such as ICS [2] could be used to suggest the relationship between users' goals (related to their interactional objectives) and the structure of the information conveyed by the computer system. IF also provides some suggestions for causes of interactional trouble *e.g.* the poor *discriminability* of events (how easy it is to discriminate between events) can be a reason for interactional trouble.

Finally IF considers the interactional *properties level*. At this level IF describes abstract properties of the interaction such as serendipity - happening upon relevant information. Such properties can be related to lower levels of the framework in order to inform both analysis of current systems and design of new systems. For example, to support serendipity the system needs to provide high event potential in viewing texts (at the fundamental level the relates to the number of events that could be issued to reach an objective), but events must be discriminable so that the user can see how the texts are related, yet different. The following sections work from the fundamental level up to the properties level to illustrate IF's utility.

<table>
<tr><td rowspan="6">↓ objectives feed down</td><td rowspan="6">↑ build on lower levels</td><td>Framework Level</td><td colspan="3">Concerned with</td></tr>
<tr><td>Properties</td><td colspan="3">Abstract interactional properties e.g. familiarisation, serendipity, lucidity.</td></tr>
<tr><td>Causes</td><td>What causes interactional trouble e.g. discriminability of events.</td><td>↔</td><td>Hooks to other modelling techniques or theories</td></tr>
<tr><td>Symptoms</td><td colspan="3">Whether trajectory is canonical, and if not, what interactional trouble exists e.g. interactional detours.</td></tr>
<tr><td>Traces</td><td colspan="3">Interaction trajectories - partially ordered events to attempt to achieve objectives. Whether objectives are achieved or not.</td></tr>
<tr><td>Fundamentals</td><td colspan="3">Agents, agent grouping, communication channels between agents, events between agents, system state, objectives described in terms of system state.</td></tr>
</table>

Table 1: Interaction Framework levels of description

3.1 Fundamentals: Understanding The Basics Of Interaction

To recap, the following properties are considered at the fundamental level of interaction analysis:

- the agents involved in the interaction (users, computer systems *etc.*)
- the grouping of the agents (groups of agents can themselves be regarded as agents)
- communication channels between agents, and the constraints on these channels
- the state of the system (the set of agents)
- the objectives of the agents which are defined in terms of pre and post conditions on the state of the system
- the communicative events between agents which cause a change in the state of the system

The New Zealand Digital Library provides a pertinent example of a current digital library involving various forms of interaction. IF's use is illustrated by considering possible interactions with the library. As a starting point we consider what agents are involved in the interaction when a digital library is used, for example when searching for a particular text. The system at least consists of the user and the computer with which they interact (figure 1a). However, it is more illuminating in terms of the factors outlined in the introduction (information structure, network constraints, content type, and context) to consider the computer as a multi-agent grouping (figure 1b). The user then interacts with relevant parts of the group as suits their objectives. Of course, the question then arises of what level of granularity is appropriate for given analyses. Our current *intuitive* view is that each agent should represent something that we are interested in making some assertions, or raising some questions, about. In this example the simplistic view of the agent structure (figure 1a) is not appropriate if we wish to consider different clients which provide different forms of access to the library server, or network constraints on access to the server. Here the browsing and searching clients may require different network resources which will in turn affect their interactions with the user. Similarly, it would be possible to consider the agents at too low a level of detail for the current purpose. For example, the cognitive processes involved in users' interactions with the computer system, such as those described by [2] could be represented as agents which communicate with each other. However, for the purposes of this analysis this would be too fine a grain of detail and would make such analyses unwieldy.

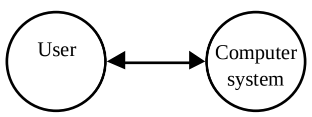

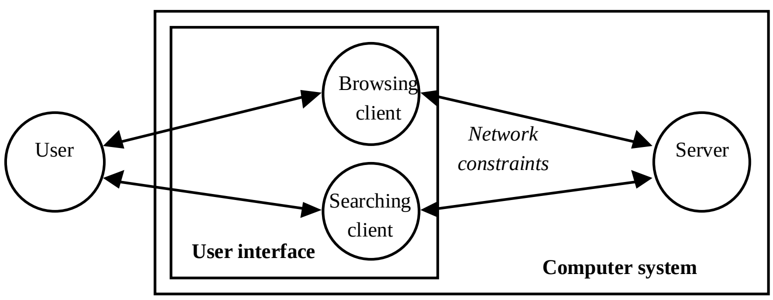

Figure 1a: Simplified agent view Figure 1b: Possible grouping within computer system

In introducing the grouping within the computer system (figure 1b) we have already alluded to two of the objectives that users have when using digital libraries: searching for and browsing for a text. Other *idealised* objectives (use of digital libraries as with other real world activities is never so clear cut) include scanning, reading, and writing [9]. To illustrate the use of IF we shall be considering just one of these objectives in this paper: searching for a particular text. Moreover, as mentioned previously, we shall be considering just one digital library here: the New Zealand Digital Library.

At the time of writing, the NZDL can be accessed *via* a world wide web client such as Netscape Navigator. The library itself consists of a set of 26 online collections of materials ranging from the BBC online archive, to the HCI bibliography, and a music video archive. A subset of these are illustrated in the screenshot of figure 2. Each of these collections provides different content and often different means to access the content. The music library, for instance, allows searching by melody. The particular example considered in this paper involves a user whose objective is to find the text of Jules Verne's *Survivors of the Chancellor*. In the set of collections provided by NZDL this text is to be found in the Gutenberg Collection, a collection containing typically out of copyright material, and some historically important documents whose texts are submitted and reviewed by volunteers. In short, the user's objective in the following example is:

- to find the text of Jules Verne's *Survivors of the Chancellor* in the Gutenberg Collection *via* the NZDL

Figure 2: Screenshot showing some of the set of collections provided by NZDL

The NZDL user interface which provides access to the Gutenberg collection is illustrated in the screenshot of figure 3. It allows users to view a list of titles or authors, or to search attributes based on free text terms. With respect to the user's objective this means that there are potentially four ways to find the text *Survivors of the Chancellor*; the first two relate to the browsing client of figure 1b, whereas the second two relate to the searching client:

- view the list of titles beginning with S
- view the list of authors whose surname begins with V
- search for texts entitled *Survivors of the Chancellor*
- search for texts by the author Jules Verne

Figure 3: Screenshot of NZDL's support for access to the Gutenberg collection

IF refers to the number of possible ways of achieving an objective as the *event potential* of the system. An interactional requirement on digital libraries derived from a desire to support serendipity may be that they should have high event potential *i.e.* many ways of achieving objectives such as locating texts, which may contribute to the potential for a user happening upon other relevant texts.

3.2 Trace Level: Event Sequences

As mentioned previously, the trace level of the interaction framework is concerned with descriptions of interaction trajectories which are episodes of communication between agents described by a partial ordering of communication events. At the trace level IF can be used to assess interaction in terms of whether interaction trajectories meet objectives, which leads to the identification of:

- **Abandoned objectives**: the agent cannot achieve their objective within the system, and the objective is unachievable. For example, a user may try to search for a text which is not stored in the collection. They would end up abandoning their objective when it became clear that the text is not present in the collection.
- **Incorrectly abandoned objectives**: the agent does not achieve their objective within the system and believes they cannot achieve their objective, when their objective is, in fact, achievable. For example, a user may try to find a particular text and reach a point in the interaction where they believe that the text is not in the collection even though it is. This issue is returned to later in more detail.
- **Successful objectives**: the agent achieves their objectives. For example, finding a required text in a collection.

3.3 Symptoms And Causes: Identifying Trouble And Understanding Why It Occurs

In attempting to achieve their objectives the agent's interaction trajectory would ideally be *canonical* [4] *i.e.* involving no unnecessary events. However, it is often the case that trajectories are not canonical; symptoms of non-canonical trajectories are discussed in this section and are currently categorised as *blind alleys*, *interactional traps*, and *interactional detours*. Such symptoms of non-canonical trajectories and their implications for design of information management systems, are discussed in the remainder of this section and illustrated through examples from the NZDL. These examples are further used to illustrate how causes for these symptoms can be postulated and used to inform (re)design.

Figure 4 motivates the discussion of non-canonical trajectory symptoms by illustrating a set of possible interaction trajectories required to meet a user's objective in finding the text of Jules Verne's *Survivors of the Chancellor*. The figure illustrates the changes in state for the clients within the computer system (*i.e.* new web pages being displayed, indicated by circled text in the figure), and the interaction events that lead to such changes in state (*i.e.* user selection of links, entering of text, selection of menu items, or clicking of buttons, indicated by text enclosed in a square). The gray boxes indicate sets of events which can be performed in any order. In moving towards their objective the user interacts with the computer system initially by selecting the Gutenberg collection (computer system moves from state (a) to state (b)). They may then generate five different communication events (the event potential of this part is high compared with other user interfaces which only allow searching) which result in the computer system responding by moving to one of the states (c), (d), (e), (f), or (g), and so on. The changes in the user's state, such as a result of learning, are considered in the following discussions.

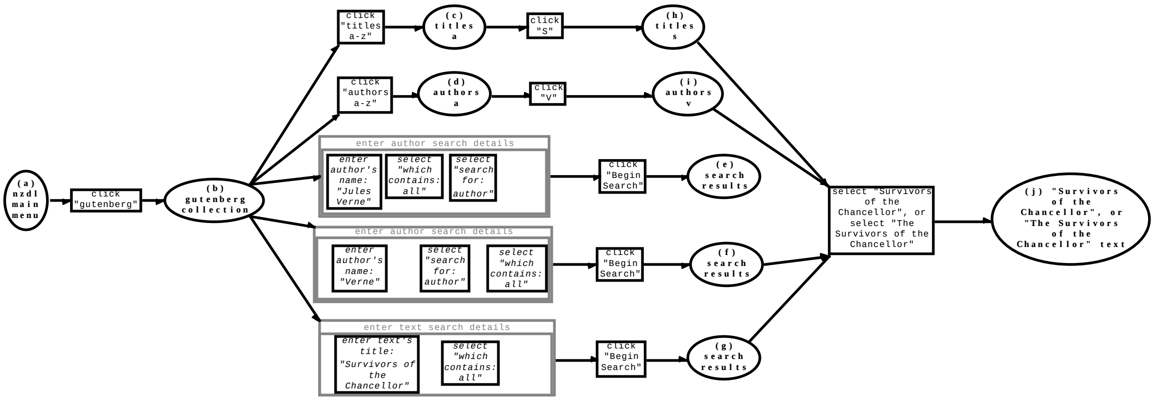

Figure 4: Possible interaction trajectories from the client agent perspective

In terms of symptoms of interactional trouble, a trajectory may involve a *blind alley* in which the agent proceeds for some time before they realise that their objective is unachievable from the current position in the interaction. The agent must then backtrack, or restart the interaction in an attempt to find an appropriate interaction trajectory to meet their objective. In the example using the NZDL this is illustrated by attempts to find *Survivors of the Chancellor* by searching for texts written by Jules Verne. At the time of writing, the Gutenberg collection's versions of *Survivors of the Chancellor* are associated with Verne, not Jules Verne. So a user could arrive at the client state (e) in the figure with a set of texts not including *Survivors of the Chancellor* as illustrated in the screen shot of figure 5. At this point a user may assume that the text does not exist (see later discussion), or may assume that they need to reformulate their search, for instance by searching for Verne only. In terms of user state this results in the user learning that within this collection, and probably within other collections accessible through the NZDL, it is best to search for authors' works by surname rather than full name. Alternatively, they may attempt to search for Jules *or* Verne resulting in a large set of search results. Either way they have been waylaid in their attempt to reach their objective and have been involved in some developments of new understandings of the use of the collection(s).

Figure 5 also illustrates a potential cause of interactional trouble in the form of multiple entries for the same text *e.g.* two entries for *De la Terre à la Lune*, and two entries for *Classic Books*. Throughout the collection there are multiple entries with the same titles, typically representing texts submitted by different contributors, or from different sources. The presence of these identical entries raises an issue of *discriminability* - how easy it is to discriminate between different possible events. This issue could be addressed in future (re)designs by providing details of the difference between entries in the search results, or by collapsing texts of the same title into one entry.

Figure 5: Results of searching for Jules Verne

Alternatively, an agent may encounter an *interactional trap* whereby the agent believes that the objective is not achievable even though it is. This contrasts with a blind alley in which the agent realises that the objective is unachievable from the current point, resets (returns to the start of the interaction), and attempts to find a different interaction trajectory to meet the objective. In the NZDL example we might consider a novice user who searched for the author Jules Verne and was presented with the set of texts illustrated in figure 5. From an unfamiliarity with the interface, and possibly search interfaces in general, the user in this case may assume that the text *Survivors of the Chancellor* is not held within the collection. As far as such as user is concerned, the computer system has presented evidence that it does not exist. This leads to the user incorrectly abandoning their objective, and moreover leads to them maintaining the belief that searching on author's full names is the most appropriate form of search within the collection, possibly generalising such a belief to other collections.

Finally, there may be an *interactional detour* whereby an agent performs interaction which does not directly move them towards their objective, for example repairing a mistake. In the running NZDL example, selecting the wrong letter from the alphabetical list of authors' surnames would constitute an interactional detour. Figure 6 illustrates a page in the collection which provides such an alphabetical list of authors' surnames. In trying to reach the set of authors including Jules Verne the user may accidentally select W. This is then repaired by correctly selecting U-V. They key is that after communicating to the computer system their intention and interpreting the computer system's response the user realises their mistake, or detour, and rectifies it. This contrasts with blind alleys where the user proceeds for some time before realising their objective is not achievable along that trajectory and so must reset the interaction before finding an appropriate trajectory as opposed to retracing a few steps to carry on to the objective.

Figure 6: Access to authors and their works *via* author surname

Understanding the causes of interactional detours can help to inform (re)design. IF allows us to use other theories, such as theories of human error (*e.g.* Reason [11]), to provide the basis for understanding causes. Using Reason's work on human error we suggest two causes of detours: *slips* and *mistakes*. Slips are due to 'actions not going according to plan' *i.e.* the user accidentally issued the wrong event such as accidentally clicking W rather than U-V, maybe the interface could be redesigned to reduce the possibility of accidental selection. Mistakes, on the other hand, are where 'the plan itself is inadequate to achieve its objectives' *e.g.* the user incorrectly thinking that the on-screen letter W would take them to texts by Verne. Maybe the interface could be redesigned to provide better visual discrimination between the letters W and V. The key here is that IF provides a means for modelling the interaction and identifying (potential) interactional troubles. Explanations for why these troubles occur are drawn from other frameworks and theories.

3.4 Properties Of Interaction

So far we have considered basic IF expressiveness from the fundamentals through to possible causes of interactional trouble. The aim of further work is to develop IF into a framework which can provide useful insights into interaction not only at this basic level, but also at more abstract levels such as the interactional requirements mentioned earlier: familiarisation, the IKIWISI principle, interactional lucidity, and canonical interactional trajectories. With regard to this aim this paper reports on work in progress towards relating interactional requirements to basic events and trajectories. Indeed, we can already provide indications of how this could be achieved. For instance, the previous example of blind alley and interactional traps alluded to the idea that these occur due to users' poor familiarity with the computer system (the user is unaware of the materials in the collection and how they are structured).

This notion of familiarity can be pursued further by considering the following brief interaction sequence with the NZDL. As mentioned previously, immediately upon entry to the NZDL, the user is shown what collections are available: nothing about the details, but very short descriptors of each. Because this entry page is very long, the user will have to scroll down to see all the page. Below the collections is a short introduction to the library and to the library access software (called Greenstone). When the user selects a collection, an introductory page about that collection is shown, together with a set of search options. In the example used here, the collection accessed in the Nutrition Collection (as opposed to the Gutenberg Collection used previously). This page, illustrated in figure 7, includes an immediately visible paragraph 'About this collection' and one on 'How to find information'. Some of the options – for example, keyword search – may give the user little further information that supports familiarisation. Others – notably view by topic – give substantial further information on what is in this collection and how it is structured. While just viewing one collection in this way does not give a detailed view of the character of the entire library, it provides a flavour. Since a uniform structure – if not a uniform content type – is imposed on all the collections, this supports familiarisation well. Therefore we would expect to see less interactional trouble as the user can easily become familiar with the structure and so not develop symptoms of interactional trouble such as pursuing blind alleys of interaction.

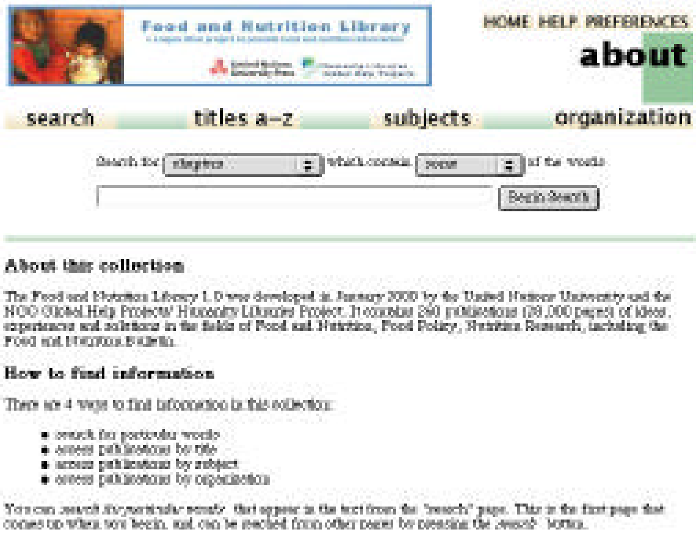

Figure 7: Introductory page to the Nutrition Collection

On the basis of this single example, we can propose a preliminary definition for the familiarisation interaction property in terms of interactional events as follows:

- Every event initiated by the device (*i.e.* system output) has a post-condition that is a significant increment on the pre-condition. (*i.e.* the user receives information that can be interpreted in the context of their existing understanding of the structure of the library and improves that understanding).
- There is a point early in the interaction where the options available to the user all result in a device event that has the same properties with respect to familiarisation (*i.e.* it does not matter which collection the user chooses: the structure and nature of the information then presented will be reliably similar).
- The length of the interaction to achieve familiarisation is short.

This definition needs further refinement: it is not yet presented elegantly or sufficiently precisely to be really useful, but indicates the direction we are heading in. However, we can use this definition as it stands to assess other libraries.

4 Discussion

The Interaction Framework's approach shows promise in shedding light on interaction issues, particularly in information management systems where typically there are, or should be, several ways of achieving objectives. Even in the simple examples considered in this paper several interaction trajectories can be employed in meeting objectives. A more developed IF will come into its own when considering the interactional requirements of more advanced information seeking and management patterns such as those discussed by O'Day and Jeffries [10]. In their work they discussed 'information orienteering' and the various techniques people employ in getting from one piece of information to another. IF's current notions of blind alley interactions, interactional traps, and detours will provide purchase on the issues of how users move around information spaces. Similarly, IF shows potential in understanding the problems of structuring and accessing various forms of information and the interaction requirements these will generate. As digital libraries become more sophisticated and provide not only different media, but also greater structuring of their meta-data (*e.g.* structured repositories of video data [5]), IF's notions of familiarity, lucidity, and others to be developed will support deeper reasoning about such issues than is currently possible.

This paper has outlined the Interaction Framework in its current state and used it to illustrate some interactional issues with online collections of text. However, since the use of the artefact depends not only on the closed world of the user and device, but on other factors that are less well understood, there is a need for empirical data to inform and extend the modelling. For this, we plan to use two longitudinal case studies that deal with different application contexts: education and medical practice. Both case studies permit direct comparison between traditional and digital approaches to information provision.

The project reported here is, as yet, at an early stage of development. We anticipate significant challenges in this work, particularly in relating local user-device interactions to the broader context of use, and in defining interactional requirements in a way that can inform design. However, while these challenges are likely to be substantial, they are a necessary step on the way towards developing applicable theory of interactions.

5 Acknowledgements

This work is supported by EPSRC Grant GR/M81748.

6 References

1. ACM (1998) Communications of the ACM, 41(4), April 1998. Special issue on Digital Libraries.

2. Barnard, P. & May, J. (1999) Representing cognitive activity in complex tasks. Human Computer Interaction, 14, 93-158.

3. Barnard, P.J., Blandford, A.E. and Harrison, M.D. (1994). The Unselected Window Scenarios: Analysis Based on the Interaction Framework. Proceedings of the CHI'94 Basic Research Symposium,.

4. Blandford, A. E., Harrison, M. D. & Barnard, P. J. (1995) Using Interaction Framework to guide the design of interactive systems International Journal of Human-Computer Studies, 43, 101-130.

5. Bryan-Kinns, N. J. (2000) VCMF: A Framework for Video Content Modelling. Multimedia Tools and Applications, Vol. 10, No. 1, pp. 23-45.

6. Hall, A. (1990) Seven Myths of Formal Methods IEEE Software, September. 11-19.

7. Harrison, M.D., Blandford, A.E. & Barnard, P.J. (1994). The Requirements Engineering of User Freedom. in F. Paternò, Ed. The Design, Specification and Verification of Interactive Systems, CNUCE-CNR.

8. Lesk, M. (1997) Practical Digital Libraries. Morgan Kaufmann.

9. Levy, D. M., & Marshall, C. C. (1995) Going Digital: A Look at Assumptions Underlying Digital Libraries. Communications of the ACM, Vol. 38, No. 4, pp. 77-84.

10. O' Day, V. L., & Jeffries, R. (1993). Orienteering in an Information Landscape: How Information Seekers Get From Here to There. Proceedings of ACM InterCHI '93, pp. 438-445.

11. Reason, J. (1990). Human Error. Cambridge University Press.

12. Smith, W., Hill, B., Long, J. & Whitefield, A. (1997) A design-oriented framework for modelling the planning and control of multiple task work in secretarial office administration. Behaviour and Information Technology. 16.3. pp. 161-183.

13. Witten I.H., Nevill-Manning C.G., McNab R. and Cunningham S.J. (1998) A public library based on full-text retrieval. Communications of the ACM, 41(4) 71_75, April. *NZDL can be accessed at http://www.nzdl.org/*.