



Contents lists available at ScienceDirect

Computer Languages, Systems & Structures

journal homepage: www.elsevier.com/locate/cl

Program transformations for information personalization

Saverio Perugini^{a,*}, Naren Ramakrishnan^b^a Department of Computer Science, University of Dayton, 300 College Park, Dayton, OH 45469-2160, USA^b Department of Computer Science, Virginia Tech, Blacksburg, VA 24061-0106, USA

ARTICLE INFO

Article history:

Received 19 December 2008

Received in revised form

5 August 2009

Accepted 3 September 2009

Keywords:

Hierarchical hypermedia
 Information personalization
 Navigation
 Out-of-turn interaction
 Program transformations
 Partial evaluation
 Program slicing
 Web interaction
 Web mining
 Website transformation

ABSTRACT

Personalization constitutes the mechanisms necessary to automatically customize information content, structure, and presentation to the end-user to reduce information overload. Unlike traditional approaches to personalization, the central theme of our approach is to model a website as a program and conduct website transformation for personalization by program transformation (e.g., partial evaluation, program slicing). The goal of this paper is study personalization through a program transformation lens, and develop a formal model, based on program transformations, for personalized interaction with hierarchical hypermedia. The specific research issues addressed involve identifying and developing program representations and transformations suitable for classes of hierarchical hypermedia, and providing supplemental interactions for improving the personalized experience. The primary form of personalization discussed is *out-of-turn interaction*—a technique which empowers a user navigating a hierarchical website to postpone clicking on any of the hyperlinks presented on the current page and, instead, communicate the label of a hyperlink nested deeper in the hierarchy. When the user supplies out-of-turn input we personalize the hierarchy to reflect the user's informational need. While viewing a website as a program and site transformation as program transformation is non-traditional, it offers a new way of thinking about personalized interaction, especially with hierarchical hypermedia. Our use of program transformations casts personalization in a formal setting and provides a systematic and implementation-neutral approach to designing systems. Moreover, this approach helped connect our work to human-computer dialog management and, in particular, mixed-initiative interaction. Putting personalized web interaction on a fundamentally different landscape gave birth to this new line of research. Relating concepts in the web domain (e.g., sites, interactions) to notions in the program-theoretic domain (e.g., programs, transformations) constitutes the creativity in this work.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

“The important thing in science is not so much to obtain new facts as to discover new ways of thinking about them.”

Sir William Lawrence Bragg, the youngest-ever recipient of the Nobel Prize

* Corresponding author. Tel.: +1937 229 4079; fax: +1937 229 2193.

E-mail addresses: saverio@udayton.edu (S. Perugini), naren@cs.vt.edu (N. Ramakrishnan).

URLs: <http://academic.udayton.edu/SaverioPerugini> (S. Perugini), <http://people.cs.vt.edu/~naren> (N. Ramakrishnan).

Information personalization constitutes the mechanisms necessary to automatically customize information content, structure, and presentation to the end-user to reduce information overload [1]. Personalization technologies are now ubiquitous on the web and critical to retaining customers (e.g., eBay, Amazon).

Our view of personalization is oriented toward personalizing user *interaction*. Specifically, we have developed an interaction technique which empowers a user navigating a hierarchical website to postpone clicking on any of the hyperlinks presented on the current page and, instead, communicate the label of a hyperlink nested deeper in the hierarchy. We call this technique *out-of-turn interaction* and when the user supplies out-of-turn input (i.e., a hyperlink label) we re-organize (or, in other words, personalize) the hierarchy to reflect the user's informational need.

Consider a user, interacting with an automobile website, such as Edmunds, interested in manufacturers offering hybrid automobiles. If the site's hierarchical structure requires the user to select a manufacturer at the top level, make at the following level, and so on, then to fulfill the information-seeking goal, this user would need to drill-down through each manufacturer and manually aggregate all hybrid automobiles discovered at the lower levels of the site. However, using out-of-turn interaction, this user could say 'hybrid' at the top level of the site and in response the system would prune out all manufacturer hyperlinks on the root page which do not lead to hybrid automobiles and, therefore, only present hyperlink representing manufacturers which offers hybrid models. With this set of reduced manufacturers the user has the option of browsing (i.e., clicking on one of the presented hyperlinks) or, again, interacting out-of-turn (e.g., by saying 'manual transmission').

Out-of-turn interaction permits the user to circumvent any intended flows of navigation hardwired into a hyperlink structure by the designer and, in this manner, flexibly reconciles any mismatch between the site's one-size-fits-all organization and the user's model of information seeking. Out-of-turn input can be communicated to a site either through text using a browser toolbar plugin [2] or through speech using a voice user interface [3].

Unlike traditional approaches to personalization, the central theme of our approach is to model information-seeking interactions with hierarchical hypermedia explicitly in a programmatic representation and use program transformations (e.g., partial evaluation, program slicing) to stage the interaction [4]. A *program transformation* is an automatic, closed operation mapping one program to another. Converting a program which computes x^n to one which computes x^2 is a simple example of a transformation (in this case, partial evaluation). *Program slicing* [5] is a program transformation used to extract statements which may affect or be affected by the values of variables from a program.

A website such as the *Yahoo!* directory may be viewed as a DAG, with vertices representing webpages, edges representing hyperlinks, edge labels representing hyperlink labels or search terms, and leaves representing content pages (destinations) (e.g., see Fig. 1). We support the user in experiencing a personalized traversal of such a website by permitting her to enter a search term out-of-turn and adjusting the graph accordingly, e.g., by retaining only the subgraph leading to leaves which have an occurrence of the search term on a path from the root to each leaf (e.g., see Fig. 2). A DAG may also be represented by a program, where each search term corresponds to a boolean variable in a branch in a nested conditional representation (e.g., see Table 4). In that representation, the out-of-turn adjustment described above is modeled by slicing the program. Thus, the essence of this paper is an equivalence between rooted DAGs and programs, such that some operations of interest on a DAG (such as adjustment to out-of-turn search terms) correspond to operations of interest on a program (such as slicing).

In summary, the central theme of our research is to pose website personalization and, particularly, website transformation, as the application of a program transformation technique to a programmatic representation of interaction based on (often partial) user input. This approach offers a new way of thinking about personalized interaction, especially with hierarchical hypermedia. Decoupling the logic into a program (representation, transformation) pair¹ provides a clean separation of concerns and allows us to personalize a hierarchy to individual users without explicitly enumerating a specialized hierarchy (or building a user model) for each individual user. It also fosters the attractive possibility of exploring alternate representations and transformations and studying the resulting forms of personalization enabled. The creativity in this research arises from relating concepts in the web domain (e.g., sites, links) to notions in the program-theoretic domain (e.g., programs, transformations) (see Table 1).

1.1. Objectives

We have built a software framework based on the theoretical ideas presented in this paper [3] and have conducted human-computer interaction studies with users to evaluate specific systems designed with it [6, Chapter 6; 7]. The goal of this paper is study personalization through a program transformation lens, and develop a formal model, based on program transformations, for personalized interaction with hierarchical hypermedia. The primary form of personalization discussed here is out-of-turn interaction. The objectives of this paper are to

1. develop graph-theoretic interpretations of out-of-turn interaction with a general class of websites,
2. illustrate how these interpretations can be supported by a program (representation, transformation) pair (called a *model*), often involving program slicing,

¹ The representation in a program (representation, transformation) pair specifies how to model a website as a program, such as the nested conditionals representation used in the programs shown in Table 4, while the transformation is a program transformation, i.e., a closed, source-to-source operation mapping a program to another program, such as partial evaluation or program slicing.

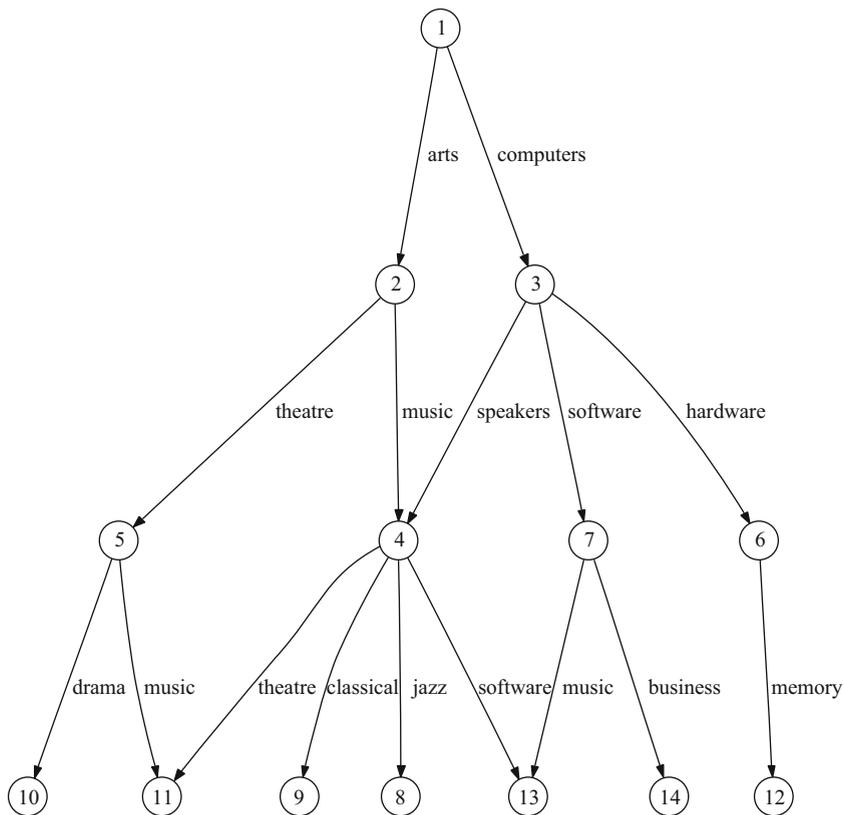


Fig. 1. Example of a DAG model of a hypothetical hierarchical website with characteristics similar to those in the *Yahoo!* directory.

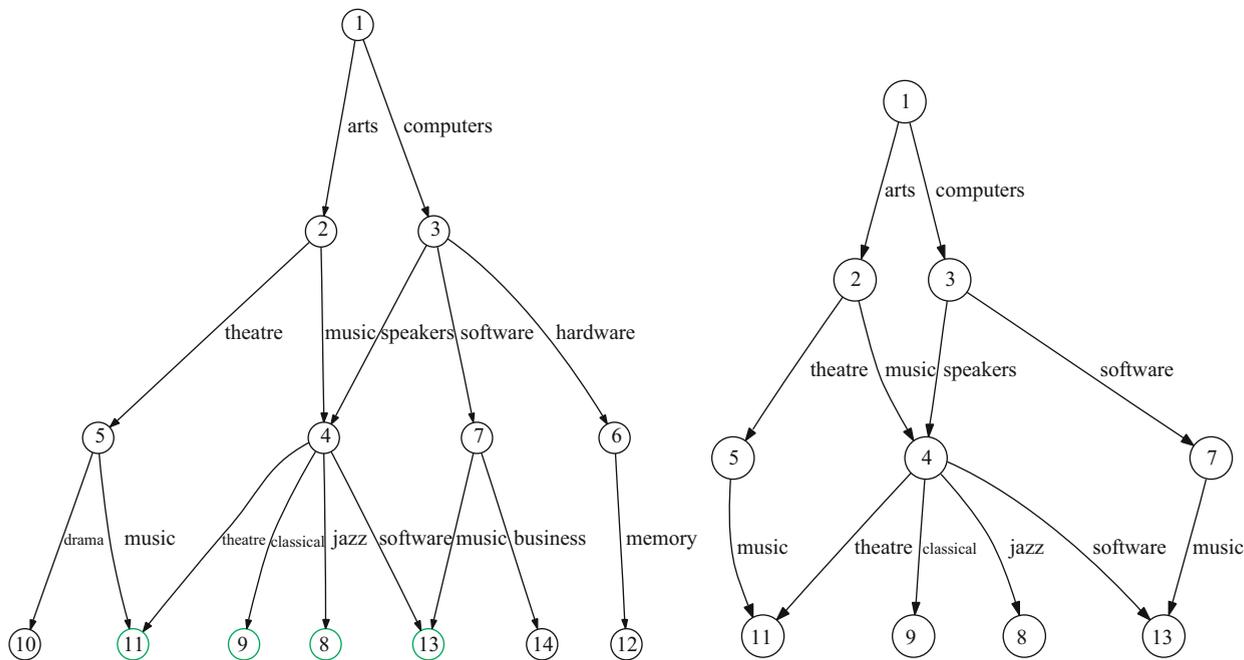


Fig. 2. Illustration of forward-propagation (FP) followed by back-propagation (BP) on the DAG in Fig. 1. (Left) Forward-propagation wrt the term 'music,' i.e., $FP(D, \text{music})$. (Right) Back-propagation wrt the $\{8, 9, 11, 13\}$ set of leaf vertices (highlighted green, or lighter in black and white rendering) in left, i.e., $BP(D, FP(D, \text{music}))$.

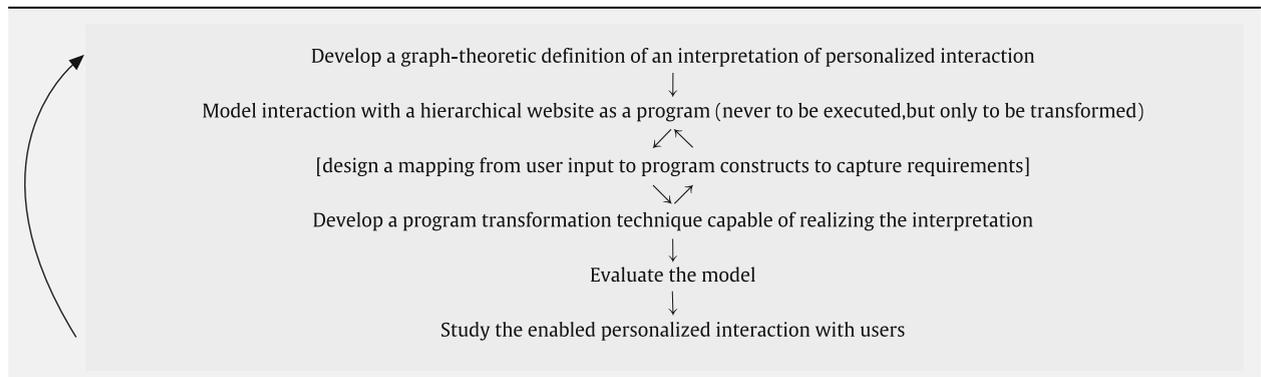
Table 1

Analogies between the web interaction and program-theoretic domains.

<i>Web interaction:</i>	Transformation(... Transformation	(Website, Hyperlink label),	..., Hyperlink label)	⇒	Personalized website
⋮	⋮	⋮	⋮	⋮	⋮
<i>Program-theoretic:</i>	Transformation(... Transformation	(Program, Program construct),	..., Program construct)	⇒	Specialized program

Table 2

Our research methodology.



3. evaluate the soundness and completeness (as well as other properties) of a model wrt a target interaction paradigm (e.g., browsing or out-of-turn interaction),
4. identify a partial order of classes of hierarchical hypermedia and explain its implications on the program transformation approach to personalization,
5. introduce web functional dependencies and describe how they can be mined from websites and used for automatic query expansion to personalize the user experience further,
6. illustrate program transformation techniques based on program slicing to mine web functional dependencies,
7. demonstrate that an alternate program transformation technique, which employs web functional dependencies, can achieve the same effect as the original transformation technique,
8. develop specialized program transformation techniques for some specific classes of hierarchical hypermedia, and
9. demonstrate how the program transformation formalism can be used to support supplementary personalized interactions.

1.2. Research methodology

When a user says something out-of-turn, we ask: what can be reasonably pruned out of the website? Answers to this question lead to interpretations of out-of-turn interaction, e.g., when a user says something out-of-turn,

1. (a) first identify leaf webpages reachable by a path involving a hyperlink labeled with the out-of-turn input, and
(b) prune all paths through the site that do not lead to any of these leaf pages and
2. prune all paths through the site which do not involve a hyperlink labeled with the out-of-turn input.

Both interpretations assume that the input supplied by the user is legal. An interaction using illegal input may either be undefined or cause an error. Interpretation 2 entails interpretation 1 because every path retained by interpretation 2 is retained under interpretation 1. In other words, interpretation 2 prunes all paths pruned under interpretation 1, but the converse does not hold. Next we ask: how can we support (i.e., model and realize) the interpretation of out-of-turn interaction using program-theoretic principles.

This suggests an iterative process, illustrated in Table 2, of developing a programmatic representation of interaction with an instance of hierarchical hypermedia and developing a program transformation technique, often involving a composition of program transformations, capable of supporting the desired personalized interactions from the model. Moreover, we must design a mapping from user requests (often partial input) to program constructs (often variables) to direct the transformation. We next evaluate the model by computing various metrics. We evaluate the personalized interaction

enabled by conducting studies with users [6, Chapter 6; 7] which often reveal insights into new interpretations, interaction paradigms, and interaction techniques and, thus, help close the loop.

While some parts of this paper have been previously reported upon by the authors [3,4], the present paper builds upon these foundations to develop a unifying theory for personalization using program transformations. Where necessary, key results from the above papers (e.g., Table 1, taken from [4]; Table 5, taken from [3]; Table 6, taken from [4], and Section 8, taken from [4]) are reported to ensure that the discussion here is self-contained.

2. Graph-theoretic view of personalized interaction

We begin by developing syntactic notions from graph theory and progressively attach web interaction semantics to develop a theory of representing and reasoning about interaction with hierarchical hypermedia.

2.1. Syntactic and semantic notions

Fig. 1 illustrates a DAG model of a hierarchical website with characteristics similar to web directories such as *Yahoo!* at <http://dir.yahoo.com> or the *Open Directory Project* (ODP) at <http://dmoz.org>. Table 3 is an abridged mapping from graph-theoretic notions to web analogs. Edges help model paths through a website a user follows to access leaf vertices. Leaf vertices model leaf webpages which contain content. We refer to a leaf content page as *terminal information* and the terms therein as *units of terminal information*. Edge-labels, which we refer to as *structural information*, model hyperlink labels or, in other words, choices made by a navigator en route to a leaf. An edge-label, a *unit of structural information*, is therefore a *term of information-seeking* (simply a *term* hereafter) which a user may bring to bear upon information seeking. Structural information thus helps make distinctions among terminal information.

A set of terms is *complete* when it determines a particular leaf webpage; otherwise it is *partial*. An *interaction set* of a DAG D is the complete set of the terms along a path from the root of D to a leaf vertex of D . An interaction set constitutes complete information; any proper subset of it is partial information. An interaction set of D classifies a leaf vertex of D , but does not capture any order of the terms.

We now provide definitions which pertain to a user's interaction with a website. A term is *in-turn* information if it appears as a hyperlink label on the user's current webpage and is, thus, currently solicited by the system. On the other hand, a term is *out-of-turn* information if it represents a hyperlink label nested somewhere deeper in the site and is, thus, currently unsolicited from the system, but relevant to information seeking. Each term from D which is not in-turn information is out-of-turn information.

Several partial orders can be defined over an interaction set wrt the time at which the user communicates the term to the system, called *arrival time*. When a user clicks on a hyperlink, she implicitly communicates the hyperlink label to the underlying system. For instance, when a user clicks on the hyperlink labeled 'arts' followed by that labeled 'music,' she communicates the $\langle \text{arts}, \text{music} \rangle$ terms, in that order. Similarly, when the user supplies out-of-turn input (using a textual or speech modality), he is communicating terms to the system. These partial orders can be summarized with partially ordered sets or posets. Each linear extension of such a poset is a total order called an *interaction sequence*. A *browsing interaction sequence* of D is a total order on an interaction set of D wrt the parenthood relation of D . An *out-of-turn interaction sequence* of D is a total order on an interaction set of D wrt the arrival time relation implied by out-of-turn interaction. Interestingly, both interpretations of out-of-turn interaction introduced above imply the same arrival time relation—a partial order containing only the reflexive tuples of terms from the interaction set. In other words, none of the terms from the interaction set are required to be ordered. The linear extensions of the posets associated with these partial orders are out-of-turn interaction sequences.

An *interaction paradigm* \mathcal{P} for D is given by the union of all linear extensions of posets defined over interaction sets of D . In other words, an interaction paradigm is a complete set of realizable interaction sequences from D wrt an interaction technique. When an edge-label labels more than one edge in a path from the root of D to a leaf vertex of D , it is advantageous to think of an interaction sequence as a finite effective enumeration of an interaction set of D , where the order of the terms in the enumeration corresponds to the arrival time relation afforded by the interaction technique wrt D .

Table 3
Graph-theoretic constructs and web analogs.

Graph-theoretic construct	Web analog
Graph	Website
Vertex	Webpage
Edge	Hyperlink
Edge-label	Hyperlink label
Root	Homepage

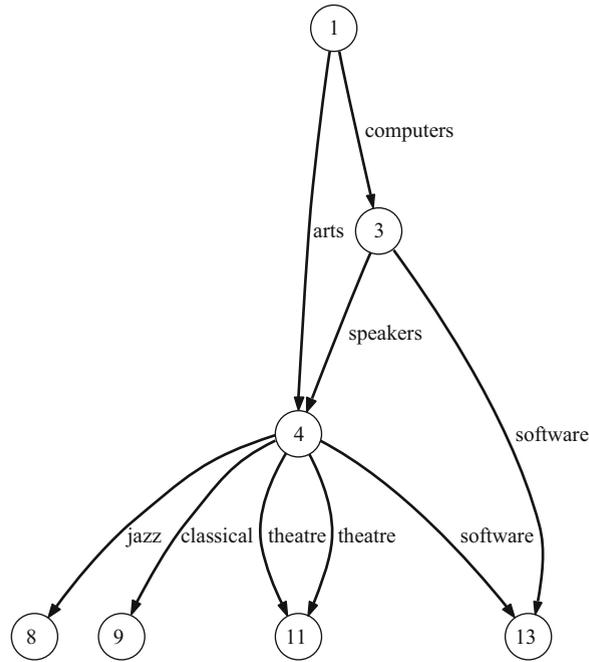


Fig. 3. Result of interpretation 1 of out-of-turn interaction with the DAG D shown in Fig. 1 wrt the term 'music,' i.e., $OOT_1(D, \text{music})$. Alternatively, one can think of this DAG as the result of shrink edges with the DAG D' in Fig. 2 (right), i.e., $SE(D', \text{music})$.

The browsing paradigm of D in Fig. 1 is $\{ \langle \text{arts, music, jazz} \rangle, \langle \text{arts, music, classical} \rangle, \langle \text{arts, music, theatre} \rangle, \dots, \langle \text{computers, hardware, memory} \rangle \}$. Likewise, an out-of-turn paradigm is

$\{ \langle \text{arts, music, jazz} \rangle, \langle \text{music, arts, jazz} \rangle,$
 $[\text{the remaining 4 permutations of } \{ \text{arts, music, jazz} \}],$
 $\langle \text{arts, music, classical} \rangle, \langle \text{music, arts, classical} \rangle,$
 $[\text{the remaining 4 permutations of } \{ \text{arts, music, classical} \}],$
 $\dots,$
 $\langle \text{computers, hardware, memory} \rangle, \langle \text{hardware, computers, memory} \rangle,$
 $[\text{the remaining 4 permutations of } \{ \text{computers, hardware, memory} \}].$

While there can be only one browsing paradigm, there are multiple out-of-turn paradigms. Moreover, the browsing paradigm for a DAG D is a subset of any out-of-turn paradigm for D .

2.2. Support terms and tools

We use the symbol \mathcal{D} to represent the universal set of DAGs, the symbol \mathcal{T} to represent the universal set of terms, and \mathcal{L} to denote the universal set of leaf webpages. Before we can expand this discussion to functions over $\mathcal{D} \times \mathcal{T}$ to realize the sequences of a particular interaction paradigm, we must develop some support terms and tools.

Sequencize is a total function $SQ : \mathcal{D} \rightarrow P(\mathcal{I})$ which given D returns the browsing paradigm of D . We use the symbol \mathcal{I} to represent the universal set of interaction sequences. $P(\cdot)$ denotes the power set function.

Term extraction is a total function $TE : \mathcal{D} \rightarrow P(\mathcal{T})$ which given D returns the set of all unique terms in D . A *term-co-occurrence set* of D is a set $T \subseteq TE(D)$. Let the *level* of an edge-label in D be the depth of the source vertex of the edge it labels. If a given edge-label occurs multiple times in D , a level is associated with every occurrence. A *term-level set* of D is a term-co-occurrence set comprising all unique terms in D with the same level. *Term-level extraction* is a total function $TLE : (\mathcal{D} \times \mathcal{N}) \rightarrow P(TE(D))$ which given D and a level $l (\geq 1) \in \mathcal{N} = \{1, 2, \dots, M\}$ returns the set of all unique terms in D with level l , i.e., a term-level set. We use the variable M to represent the maximum depth of D . If D represents the DAG in Fig. 1, $TLE(D, 2) = \{ \text{theatre, music, speakers, software, hardware} \}$. In any DAG, $TLE(D, 1)$ returns the set of terms available to supply through browsing.

Get sequences is a partial function $GS : (\mathcal{T} \times P(\mathcal{I})) \rightarrow P(\mathcal{I})_{\perp}$ which given a term t and a set of interaction sequences I_S returns the set of all interaction sequences in I_S which each contain t as a member. The symbol \perp denotes the partial nature of the function, i.e., the value of GS is undefined for some inputs. *Forward propagate* is a total function $FP : (\mathcal{D} \times \mathcal{T}) \rightarrow P(\mathcal{L})$ which given D and a term $t \in \mathcal{T} = TE(D)$ returns a set of leaf vertices L of D , where L contains each leaf vertex reachable from all paths of D containing an edge labeled t . *Collect results* is a total function $CR : \mathcal{D} \rightarrow P(\mathcal{L})$ which given D returns a set of all the leaf vertices in D . Collect results wrt the sub-DAG rooted at vertex 2 in Fig. 1 is the $\{8, 9, 10, 11, 13\}$ set of vertices. *Back*

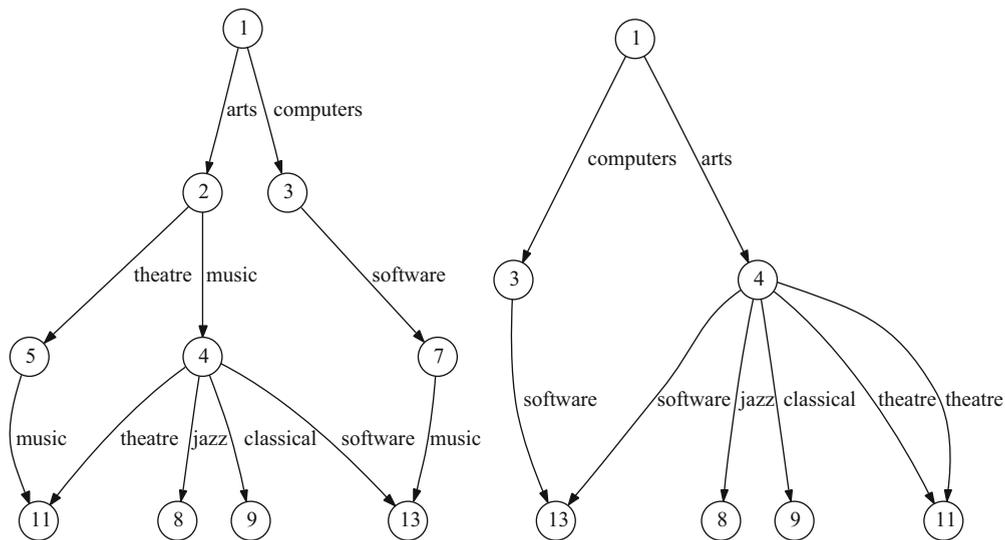


Fig. 4. (Left) Result of applying select paths to the DAG D' shown in Fig. 2 (right) wrt the term 'music,' i.e., $SP(D', \text{music})$. (Right) Result of interpretation 2 of out-of-turn interaction with the DAG D shown in Fig. 1 wrt the term 'music,' i.e., $OOT_2(D, \text{music})$. Alternatively, one can think of this DAG as the result of shrink edges with the DAG D'' (left), i.e., $SE(D'', \text{music})$.

propagate is a total function $BP : (D \times P(L)) \rightarrow D$ which given D and a set of leaf vertices L returns a DAG D' , where D' contains only browsing interaction sequences which classify the leaf vertices in L . Fig. 2 illustrates forward-propagation (left) on the DAG D in Fig. 1 followed by back-propagation (right) yielding D' .

Notice that D also can be represented as a $|TE(D)| \times |CR(D)|$ term-document matrix, where rows correspond to terms (structural information or edge-labels) and columns correspond to webpages (terminal information or leaf vertices). However, such a matrix is insufficient to reconstruct D .

2.3. Interpretations of out-of-turn interaction

Prior to providing graph-theoretic interpretations for out-of-turn interaction, we formalize browsing over DAG models of websites. *Browse* is a partial function $B : (D \times T) \rightarrow D_{\perp}$, which given D and a term $t \in TLE(D, 1)$ returns the sub-DAG rooted at the target vertex of the edge in D labeled with edge-label t whose source vertex is the root of D . If the DAG in Fig. 1 is D , $B(D, \text{computers})$ returns the sub-DAG rooted at vertex 3, which represents the result of a user clicking on the hyperlink labeled 'computers.'

We formally cast the above two interpretations for out-of-turn interaction in graph-theoretic terms:

Interpretation 1 of out-of-turn interaction is a partial function $OOT_1 : (D \times T) \rightarrow D_{\perp}$ which given D and a term $t \in TE(D)$ returns D' . It is defined as

$$OOT_1(D, t) = SE(\overbrace{BP(D, \underbrace{FP(D, t)}_{\text{Fig. 2 (left)}})}_{\text{Fig. 2 (right)}}, t), \tag{1}$$

where *SE* (*shrink edges*) is a partial function $SE : (D \times T) \rightarrow D_{\perp}$ which given D and a term $t \in TE(D)$ returns D' , where any edge e in D labeled with t is removed in D' , the source v_s of e is replaced with its target v_t in D' , and v_t becomes the new target of any edge e' with target v_s in D' .

Interpretation 2 of out-of-turn interaction is a partial function $OOT_2 : (D \times T) \rightarrow D_{\perp}$ which given D and a term $t \in TE(D)$ returns D' (see Fig. 4). It is defined as

$$OOT_2(D, t) = SE(\overbrace{SP(BP(D, \underbrace{FP(D, t)}_{\text{Fig. 4 (left)}}), t)}_{\text{Fig. 4 (right)}}), \tag{2}$$

where *SP* (*select paths*) is a partial function $SP : (D \times T) \rightarrow D_{\perp}$ which given D and a term $t \in TE(D)$ returns D' , where D' contains only the browsing interaction sequences from D involving t (i.e., $SQ(SP(D, t)) = GS(t, SQ(D))$). All browsing interaction sequences pruned under OOT_1 also are pruned under OOT_2 , but the converse does not hold. Formally, $SQ(OOT_2(D, t)) \subseteq SQ(OOT_1(D, t))$. Interpretation 2 retains shrunken versions of only browsing interaction sequences

involving the out-of-turn input, while interpretation 1 retains shrunken versions of all browsing interaction sequences which classify the leaf vertices classified by sequences involving the out-of-turn input.

The result of FP is the set of all leaf vertices classified by the browsing interaction sequences involving the out-of-turn input. We back-propagate from this set of leaves up to the root of the DAG with BP . To generalize this approach, we can replace FP with SL (*select leaves*)—any total function $SL : (D \times T) \rightarrow \mathcal{L}$ which given D and a term from D returns a set of leaf vertices of D . FP is an instance of SL . The use of a function such as SL allows us to marry our approach with standard techniques from information retrieval (IR) [8]. Moreover, our inclusion of a function which returns a set of leaf vertices leads to the possibility of bringing units of terminal information (additional terms modeled in the leaf documents and not explicitly used in the classification), in replacement of or in addition to structural information, to bear upon the interaction. For instance, we might perform a query (e.g., ‘Picasso’) in a vector-space model over the set of leaf webpages (documents) using cosine similarity to arrive at a target set of leaves from which to back-propagate. We do not study this extension here and present SL primarily as a technique which works with structural information.

Generalized interpretation 1 of out-of-turn interaction is a partial function $GOOT_1 : (D \times T) \rightarrow \mathcal{D}_\perp$ which given D and a term $t \in TE(D)$ returns D' . It is defined as

$$GOOT_1(D, t) = SE(BP(D, SL(D, t)), t). \quad (3)$$

Generalized interpretation 2 of out-of-turn interaction is a partial function $GOOT_2 : (D \times T) \rightarrow \mathcal{D}_\perp$ which given D and a term $t \in TE(D)$ returns D' . It is defined as

$$GOOT_2(D, t) = SE(SP(BP(D, SL(D, t)), t), t). \quad (4)$$

Notice that Eqs. (3) and (4) are the analogs of Eqs. (1) and (2), respectively, in that FP in the latter is replaced by SL . Note also that the out-of-turn paradigm studied here is supported by all four interpretations of out-of-turn interaction.

Notice also that SE can be made optional in each interpretation of out-of-turn interaction. The absence of SE however requires the user to ultimately browse to reach leaf content pages. This idea is generalized below through a concept we call *web functional dependencies*.

Lemma 2.1. *Any interpretation of out-of-turn interaction is commutative, assuming both sides are defined, i.e.,*

$$\begin{aligned} OOT_1(OOT_1(D, x), y) &= OOT_1(OOT_1(D, y), x), \\ OOT_2(OOT_2(D, x), y) &= OOT_2(OOT_2(D, y), x), \\ GOOT_1(GOOT_1(D, x), y) &= GOOT_1(GOOT_1(D, y), x), \quad \text{and} \\ GOOT_2(GOOT_2(D, x), y) &= GOOT_2(GOOT_2(D, y), x), \end{aligned}$$

where x and y represent terms. A sketch of the proof of Lemma 2.1 is given in [6, Chapter 4]. While we do not study this extension here, Lemma 2.1 is important to supporting multiple terms per utterance, i.e., the ability to communicate more than one term to the system in a single interaction, an important ingredient for personalized interaction.

3. Program-theoretic view of personalized interaction

We now view the above graph-theoretic interpretations of out-of-turn interaction through a programmatic lens and, specifically, relate out-of-turn interaction, and generalizations thereof, to *program slicing*. We begin by illustrating how to model interaction with a hierarchical website as a program.

3.1. Modeling interaction programmatically

Researchers have predominately modeled web interaction programmatically to maintain state across and within sessions [9,10]. Here we model interaction programmatically to personalize information access. Table 4 (left) illustrates a programmatic model of a user’s browsing interactions with Fig. 1. Notice our use of procedures to model common sub-trees induced by symbolic links. A *symbolic link* is a hyperlink to a webpage which has an existing incoming hyperlink [11]. For example, the edge from vertex 3 to 4 in Fig. 1 labeled ‘speakers’ is a symbolic link. The absence of symbolic links from a website make its DAG model a tree. In voluminous web directories, such as *Yahoo!* and *ODP*, symbolic link labels are augmented with ‘@’ and transport a user from one sub-branch of the directory (e.g., News) to another (e.g., Sports).

The expressive constructs, most notably, nested conditionals, of programming languages make such languages an attractive vocabulary of discourse for modeling interaction with hierarchical hypermedia. While we can model interaction with (imperative, functional, or logic) programming languages, we use C here for purposes of presentation. At the NATO Conference on Software Engineering Techniques held in Rome, Italy in 1969, A. Perlis said, ‘a good programming language is a conceptual universe for thinking about programming’ [12]. Similarly here, a good programming language is a conceptual universe for thinking about interacting with hierarchical hypermedia.

3.2. Program slicing

We relate the four interpretations of out-of-turn interaction to the application of program slicing, a transformation common in debuggers, to suitably selected programmatic representations of interaction, such as that shown Table 4 (left).

Table 4
Modeling interaction programmatically.

P_D	$P_{D'} = \llbracket \text{zoom} \rrbracket [P_D, \text{music}]$	$P_{D''} = \llbracket \text{mix} \rrbracket [P_{D'}, \text{music} = 1]$
1 if (arts)	if (arts)	if (arts)
2 if (theatre)	if (theatre)	if (theatre)
3 if (drama)		
4 page = 10;		
5 if (music)	if (music)	
6 f1();	f1();	f1();
7 if (music)	if (music)	
8 f2();	f2();	f2();
9 if (computers)	if (computers)	if (computers)
10 if (speakers)	if (speakers)	if (speakers)
11 f2();	f2();	f2();
12 if (software)	if (software)	if (software)
13 if (music)	if (music)	
14 f3();	f3();	f3();
15 if (business)		
16 page = 14;		
17 if (hardware)		
18 if (memory)		
19 page = 12;		
20		
21 void f1() {	void f1() {	void f1() {
22 page = 11;	page = 11;	page = 11;
23}	}	}
24		
25 void f2() {	void f2() {	void f2() {
26 if (theatre)	if (theatre)	if (theatre)
27 f1();	f1();	f1();
28 if (classical)	if (classical)	if (classical)
29 page = 9;	page = 9;	page = 9;
30 if (jazz)	if (jazz)	if (jazz)
31 page = 8;	page = 8;	page = 8;
32 if (software)0	if (software)	if (software)
33 f3();	f3();	f3();
34 }	}	}
35		
36 void f3() {	void f3() {	void f3() {
37 page = 13;	page = 13;	page = 13;
38 }	}	}

(Left) P_D , programmatic representation of interaction with the website modeled by the DAG D in Fig. 1. (Center) $P_{D'}$, result of applying the zoom transformation to the program on the left wrt `music`. This program is the representation of interaction with the website modeled by the DAG D' in Fig. 2 (right). (Right) $P_{D''}$, result of applying partial evaluation to (center) wrt `music = 1`. This program is the representation of interaction with the website modeled by the D'' in Fig. 3.

Program slicing [5,13,14], originally introduced by [15–17], is a technique used to extract statements, which may affect or be affected by the values of variables of interest computed at some point of interest, from a program. A slice of a program is taken wrt a (point of interest, variable of interest) pair, referred to as the *slicing criterion*. The point of interest may be specified with a line number from the program. The resulting slice consists of all program statements which may affect or be affected by the value of the variable at the specified point.

Slices such as that shown in Table 5 (center), which Weiser first articulated, are called *executable backward static slices* [5,18,19] (referred to here as simply *backward slices*). They are *executable* because the slice is required to be an executable program. The slice is *backward* since this is the direction in which dependencies are followed to sources in the program. Contrast this with a *forward slice* [18] which consists of the program statements affected by the value of a particular variable at a particular statement (see Table 5, right). Backward slices contain data and control predecessors, while forward slices consist of data and control successors. The slice is *static* because it is computed without consideration of program input. Dynamic slicing techniques are covered in [6, Appendix B]. For an introduction to program slicing, techniques for computing slices, and applications, we refer the reader to [5,13,14].

We relate FP to the program intersection of two unions of forward slices. Researchers in the programming languages community consider a union of slices as a slice and trivial to compute. The slicing criteria ‘can be easily extended to slicing wrt a collection of locations and a collection of variables at each location by taking the union of the individual slices’ [5]. The first union consists of the forward slices of the program wrt the variable modeling the out-of-turn input at every occurrence of it in the program. The second union consists of the forward slices of the program wrt the variable indexing the leaf vertices at every occurrence of it in the program. The intersection of these two unions results in several occurrences of the variable indexing the leaf vertices (e.g., `page` in Table 4). Each occurrence is at a point in the program which is

Table 5
Illustration of program slicing (simplified for purposes of presentation).

Original program	backward(6,volume)	forward(1,h)
<pre> 1 read(r,h); 2 cArea = PI * r * r; 3 sArea = 2 * cArea + 2 * r * PI * h; 4 volume = cArea * h; 5 print(sArea); 6 print(volume); </pre>	<pre> read(r,h); cArea = PI * r * r; volume = cArea * h; print(volume); </pre>	<pre> read(r,h); sArea = 2 * cArea + 2 * r * PI * h; volume = cArea * h; print(sArea); print(volume); </pre>

(Left) A program which accepts the radius and height of a cylinder as input and computes and prints the cylinder's surface area and volume. (Center) A static backward slice wrt (6, vol). (Right) A static forward slice wrt (1,h) (variable key: r = radius; h = height; $cArea$ = circle area; $sArea$ = surface area).

Table 6
Comparison of partial evaluation and program slicing along a syntax- vs. semantic-preserving dichotomy.

	Syntax-preserving	Semantic-preserving
Partial evaluation	×	✓
Program slicing	✓	×

affected by the variable corresponding to the out-of-turn input. In graph terms, this procedure results in the set of leaf vertices classified by all browsing interaction sequences involving the out-of-turn input.

The set of program fragments, thus obtained from FP , can be thought of as slicing criteria input to BP which is then the union of backward slices, each wrt every (point, variable) pair resulting from the initial forward slicing procedure. Intuitively, given valid input, a forward slice is performed wrt the corresponding program variable to determine the terminal webpages which are reachable from that point. These webpages are collected and back-propagated with backward slicing so that only those paths that reach these pages are retained. Notice that these two operations implicitly capture exclusions among program variables, e.g., when the user says 'Honda' the slices remove any program segments which involve Toyotas. Such a combination of forward and backward slicing is similar to the zoom operator for interactively pruning information hierarchies [20]. Thus, we call this composite program transformation technique *zoom*.

The idea of performing set-theoretic operations on forward and backward slices is closely related to the concepts of *program dicing* [5,21] and *program chopping* [22]. Performing set-theoretic operations on one or more backward program slices yields a *program dice* [5,21]. Originally program dicing was limited to backward slices. Program chopping, on the other hand, which also is a generalization of slicing [22], is an extension of dicing to forward slices. Forward slices increase the usefulness of dicing [5]. Chopping identifies the statements which transmit values from one statement to another. In other words, it shows all the ways which one set of program points affect another set of points. A *program chop* [22–24] therefore consists of all program points affected between one point (the chop source) and another (the chop target) [25]. It also is the subset of the intersection of a forward and backward slice (Anderson, P., personal communication, November 10, 2003). In the absence of procedures, a chop 'can be viewed as a generalized kind of program dice' [5].

Table 6 compares partial evaluation and program slicing. It reinforces that while partial evaluation is semantic-preserving, it is not syntax-preserving. Conversely, while the variants of program slicing considered in this article are syntax-preserving, they are not semantic-preserving. However, there are variants of program slicing (e.g., *amorphous slicing*, also known as *semantic slicing*) which are the reverse, i.e., semantic-preserving, but not syntax-preserving (see [6, Appendix B] for more details).

3.3. Notation: programs as data objects

To succinctly describe our program transformation techniques (e.g., the above notion of zoom) we adopt a slightly modified notation for describing the semantic function of a programming language used in a textbook on partial evaluation [26]. A specification language, defined by a context-free grammar, for program transformations is introduced in [27]. While the language is imperative, here, for purposes of presentation, we use a declarative style. GrammarTech, Inc., the company which develops and produces a program slicer for ANSI C called *CodeSurfer* is currently working on a textual representation, employing a Lisp-like syntax, for set-theoretic operations over program slices for a future release, e.g., (intersect (slice A) (slice B)) (Anderson, P., personal communication, April 8, 2004):

- $\llbracket \text{int} \rrbracket [P_D, x_1 = 1, x_2 = 0, \dots, x_n = 1]$ denotes 'partially, interpret the programmatic representation of DAG D (P_D) wrt the partial assignment of variables $x_1 = 1, x_2 = 0, \dots, x_n = 1$.' $\llbracket \text{int} \rrbracket$ denotes the application of a stepwise interpreter. We use interpretation here as a program transformation.

Table 7

Relating interaction techniques in DAG models of a websites to compositions of program transformations.

Interaction technique	Program transformation technique
Browsing	$\llbracket \text{int} \rrbracket_{[P_D, \text{input} = 1]}$
Interp. 1 of OOT Interaction	$\llbracket \text{mix} \rrbracket_{\llbracket \llbracket \text{zoom} \rrbracket_{[P_D, \text{input}], \text{input} = 1}]}$
Interp. 2 of OOT Interaction	$\llbracket \text{mix} \rrbracket_{\llbracket \llbracket \text{sp} \rrbracket_{[P_D, \text{input}], \text{input} = 1}]}$
Gen. Interp. 1 of OOT Interaction	$\llbracket \text{mix} \rrbracket_{\llbracket \llbracket \text{backward} \rrbracket_{[P_D, \llbracket \text{SL} \rrbracket_{[P_D, \text{input}], \text{input} = 1}]}]}]$
Gen. Interp. 2 of OOT Interaction	$\llbracket \text{mix} \rrbracket_{\llbracket \llbracket \text{sp} \rrbracket_{\llbracket \llbracket \text{backward} \rrbracket_{[P_D, \llbracket \text{SL} \rrbracket_{[P_D, \text{input}], \text{input} = 1}]}]}]}]$

Notice that $\llbracket \text{SL} \rrbracket$ is a meta-program-transformation, i.e., it represents any program transformation which returns a set of program points containing the variable `page`.

- $\llbracket \text{mix} \rrbracket_{[P_D, x_1 = 1, x_2 = 0, \dots, x_n = 1]}$ denotes ‘partially evaluate (non-sequentially interpret) the programmatic representation of DAG $D(P_D)$ wrt the partial assignment of variables $x_1 = 1, x_2 = 0, \dots, x_n = 1$.’ $\llbracket \text{mix} \rrbracket$ is the conventional way to denote a partial evaluator [26,28,29]. It refers to *mixed computation* since a partial evaluator performs a mixture of interpretation and code-generation [28].
- $\llbracket \text{forward} \rrbracket_{[P_D, x]}$ denotes ‘union each forward slice of the programmatic representation of DAG $D(P_D)$ wrt the variable x at every program point containing x .’
- $\llbracket \text{backward} \rrbracket_{[P_D, x_1, x_2, \dots, x_n]}$ denotes ‘union each backward slice of the programmatic representation of DAG $D(P_D)$ wrt the variable x at program points $1, 2, \dots, n$.’

Using this notation, $\llbracket \text{zoom} \rrbracket$ is formally defined as

$$\llbracket \text{zoom} \rrbracket_{[P_D, \text{input}]} = \llbracket \text{backward} \rrbracket_{[P_D, \llbracket \text{forward} \rrbracket_{[P_D, \text{input}]} \cap \llbracket \text{forward} \rrbracket_{[P_D, \text{page}]}]},$$

where `input` is the program variable representing the out-of-turn input and `page` is the variable indexing the leaf vertices. We define $\llbracket \text{sp} \rrbracket$, the programmatic analog to *select paths (SP)*, as

$$\llbracket \text{sp} \rrbracket_{[P_D, x]} = \llbracket \text{forward} \rrbracket_{[P_D, x]} \cup \llbracket \text{backward} \rrbracket_{[P_D, x]}.$$

We define $\llbracket \text{te} \rrbracket_{[P_D]}$, the programmatic analog to *term extraction (TE)*, as the union of all program data successors of each structural variable at its declaration. A program data successor is a restriction to a forward slice in that rather than including transitive dependencies, it just contains the immediate dependency of a program point. Specifically, ‘a program point’s data successors are the points where the variables that were modified at that point are used’ [25]. Since structural variables’ sole presence in these programs arises in the context of an `if (...)` expression, we might just as easily think of conducting $\llbracket \text{te} \rrbracket_{[P_D]}$ with a regular expression.

Armed with these formalisms, we can relate interpretations of interaction techniques, including browsing, to classical program transformations. In Table 7 we present program transformation techniques to realize browsing and the four interpretations of out-of-turn interaction given above. Program slicing is typically used for debugging, safety, and security. Only few have used slicing for web applications [30]. Our use of program slicing helps marry it with information personalization.

4. Evaluation

A model $\mathcal{M} = (P_D, \mathcal{X})$ is a program (representation, transformation) pair. We would like to evaluate a model by assessing its capability to realize a desired set of interaction sequences. To do so we measure how close the model comes to realizing the targeted interaction paradigm \mathcal{P} . Ideally, we would like to have

$$\mathcal{M} \rightsquigarrow I \leftrightarrow I \in \mathcal{P},$$

where \rightsquigarrow denotes *stages*, i.e., \mathcal{M} stages interaction sequence I iff I is in the interaction paradigm \mathcal{P} . A model stages an interaction sequence if successive applications of its transformation technique to its programmatic representation, given user input, realize the interaction sequence. In this manner, the model *stages* the user’s interaction.

Soundness of a model: $\mathcal{M} \rightsquigarrow I \rightarrow I \in \mathcal{P}$. A model \mathcal{M} is *sound* for an interaction paradigm \mathcal{P} if each interaction sequence which \mathcal{M} stages is in \mathcal{P} . In other words, if the model can stage an interaction sequence, then the sequence is in the paradigm.

Completeness of a model: $I \in \mathcal{P} \rightarrow \mathcal{M} \rightsquigarrow I$. A model \mathcal{M} is *complete* for an interaction paradigm \mathcal{P} if \mathcal{M} can stage each interaction sequence in \mathcal{P} . In other words, if an interaction sequence is in the paradigm, then the model can stage it.

The models presented in Table 7 are each sound and complete for the corresponding out-of-turn paradigm presented in this article. The completeness of each model holds under the assumption that no path from the root of the website to each leaf contains more than one hyperlink with the same label. Intuitively, this is because communicating (in-turn or out-of-turn) information initiates a transformation which simplifies the site wrt all hyperlinks labeled with that input, some of which may lie along the same path. This assumption is captured by our definition of interaction set and (browsing)

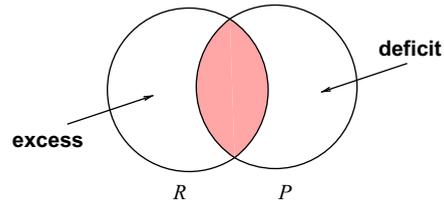


Fig. 5. Venn diagram highlighting the intersection between the set of sequences R (left) staged by an incomplete, unsound model and its intended interaction paradigm \mathcal{P} (right).

interaction sequence. As a result, our definition of browsing is slightly different than its traditional interpretation. The soundness of the model holds under the assumption that all terms supplied in any interaction sequence lie along a single path. This is always true in a tree, but may not be true in a DAG. This assumption is captured by the definition of interaction set and (out-of-turn) interaction sequence.

We can identify both the excess and deficit of a model. A complete, but unsound, model has *excess*—interaction sequences not in its intended paradigm which it stages. On the other hand, a sound, but incomplete, model has *deficit*—sequences in its intended interaction paradigm which it fails to stage. An unsound and incomplete model exhibits both excess and deficit wrt its targeted paradigm. A *sufficiency* metric for a model is formulated akin to the recall measure in IR [8]:

$$\text{sufficiency} = \frac{|R \cap \mathcal{P}|}{|\mathcal{P}|} = \frac{|R \cap \mathcal{P}|}{SOP(D)},$$

where R represents the set of sequences staged by the model and

$$SOP(D) = \sum_{I_i \in SQ(D)} |gIS(I_i)|,$$

where gIS (*get Interaction Set*) is a total function $gIS: \mathcal{I} \rightarrow \mathcal{S}$ which given an interaction sequence I_i returns the interaction set over which it is defined. We use \mathcal{S} to denote the universal set of sets. Fig. 5 illustrates how excess and deficit arise.

Lemma 4.1. *Any complete model for an out-of-turn interaction paradigm is also complete for a browsing paradigm.*

A sketch of the proof of Lemma 4.1 is given in [6, Chapter 4]. Intuitively, this lemma indicates any browsing paradigm is a proper subset of the corresponding out-of-turn interaction paradigm or, in other words, out-of-turn interaction subsumes browsing. *This is a significant result as it means that we can support the union of these two interaction paradigms with a single program transformation technique. In other words, no anticipation of in-turn or out-of-turn input is necessary to discern a program transformation technique. Thus, we achieve uniform processing of input.* For interaction this means that the user can interleave hyperlink clicks (in-turn browsing inputs) and voice utterances (out-of-turn inputs) in any order she desires [2] to achieve a *mixed-initiative* mode of information seeking [31]. *Mixed-initiative interaction* is a flexible interaction strategy where each participant can play an equal role in steering the progression of the interaction [32].

Beyond soundness and completeness, we developed a measure which estimates the *compression* achieved in a model—the ratio of interaction sequences realizable through interpretation (and hence directly stated in the representation) to the total number of sequences realizable through transformation:

$$\text{compression} = \frac{|R - E|}{|R|},$$

where E represents the set of interaction sequences stageable from P_D with an interpreter ($\llbracket \text{int} \rrbracket$). Intuitively, the compression ratio quantifies the percentage of sequences which we get ‘for free’ by using the program transformation technique.

An effective model is one which maximizes both sufficiency and compression. However, these measures are bipolar and exhibit a tradeoff similar to that between *precision* and *recall* in IR [8]. To maximize sufficiency, we might choose to explicitly model each interaction sequence in the representation, affecting the compression ratio negatively.

There are alternate applicable evaluation criteria for this research [33]. Evaluating personalization applications for traditional user-satisfaction and task completion metrics is another practice. In addition to measuring satisfaction, studies with users can improve our understanding of an extant interaction paradigm or yield new paradigms to model [6, Chapter 6; 7]. Personalization applications also can be evaluated wrt classical IR metrics [34]. For example, Sacco [20] studies how the application of zoom as well as bucket size (i.e., number of documents classified under each terminal concept) affect the maximum resolution of a taxonomy. Maximum resolution, which is a measure of retrieval effectiveness, is the average minimum number of documents the user has to manually inspect.

5. Graph-theoretic classes of hierarchical hypermedia

While the transformation techniques in Table 7 work on any DAG, we identify classes of hierarchical hypermedia for insight into the possibility of using specialized program transformation techniques to realize out-of-turn interaction for each class. We begin by defining a few terms.

The *maximum depth* of a DAG D is the level of an edge-label in D which is greater than or equal to the level of all other edge-labels in D . A *cluster* c is a term-level set such that no edge label in it labels an edge in a different term-level set. If the maximum depth of D equals the number of clusters in D , then D is *levelwise* DAG. Intuitively, a levelwise DAG is one where each level of the DAG corresponds to a *facet* [35] of information assessment in the website it models or, in other words, each term $t_i \in TE(D)$ resides at exactly one level. Therefore, sometimes such sites are referred to as *faceted*.

Each DAG shown in Fig. 6 is levelwise. The DAG in Fig. 6 (left) models a simplified path through the online Virginia Tech timetable of courses. Its clusters are {Computer Science, Mathematics} and {Ramakrishnan, Adjerid, Beattie, Green} implying that the two levels shown correspond to *department* and *instructor* facets, respectively. Notice that symbolic links, such as those illustrated in Fig. 6 (left and center) by the edge labeled ‘Adjerid’ emanating from the target of the edge labeled ‘Computer Science,’ are necessary to model the presence of cross-listed courses. The DAG in Fig. 6 (right), which models the Congressional portion of the Project Vote Smart (PVS at <http://votesmart.org>) website, also is levelwise, albeit without symbolic links. Its clusters are {California, Virginia} and {House, Senate} implying that the two levels shown correspond to *state* and *branch of Congress* facets, respectively. In contrast to those shown in Fig. 6, notice that the DAG in Fig. 1 is non-levelwise, or sometimes called *unfaceted* [36].

We now introduce the concept of mutual exclusivity in DAG models of websites. If no leaf vertex of D lies at the end of two paths from D which each involve a distinct edge-label from a term-co-occurrence set T , then we say that T is a *mutually exclusive* term-co-occurrence set. While there are multiple mutually exclusive term-co-occurrence sets (e.g., {Computer Science, Green}) in the DAGs shown in Fig. 6 (left and center), none are clusters. On the other hand, the term-level sets {California, Virginia} and {House, Senate} of the DAG shown in Fig. 6 (right) are mutually exclusive and clusters. If D is levelwise and has at least one mutually exclusive cluster, then D is a *weak-mutually exclusive* DAG. If a DAG D is levelwise and no leaf vertex of D lies at the end of two distinct paths from D , where each path contains a distinct term from the same cluster, then D is *strong-mutually exclusive* DAG. Note that the mutual exclusivity of DAGs subsumes the levelwise property by definition. Notice further that replacing only one edge-label (i.e., ‘Beattie’ to ‘Adjerid’) in the DAG shown in Fig. 6 (left) makes it a weak-mutually exclusive DAG (see Fig. 6, center); its cluster at level-two is mutually exclusive, but that at level one is not due to the symbolic link. This DAG models a course in the online Virginia Tech timetable which presumably has two sections, each taught by a different instructor, from different departments. If there is a unique simple path from the root of a DAG D to each vertex in D , then D is an *edge-labeled, rooted tree* (hereafter referred to as a *tree*), e.g., Fig. 6 (right). Observe also that we develop the above classes of hierarchical hypermedia without reference to semantics; these notions are purely syntactic, e.g., the levelwise property does not take into account polysemy of terms.

Lemma 5.1. A DAG D is levelwise and each term-level set of D is a cluster and mutually exclusive iff D is strong-mutually exclusive.

Lemma 5.2. A strong-mutually exclusive DAG which is not a tree does not exist, given our definition of a DAG.

Lemma 5.3. Interpretation 1 of out-of-turn interaction over a tree D is functionally equivalent to interpretation 2 of out-of-turn interaction over D .

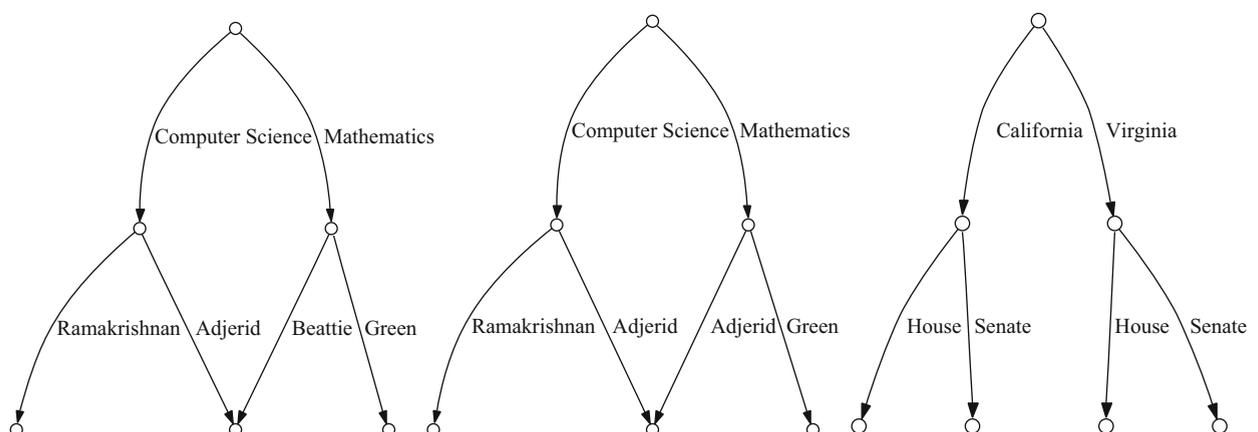


Fig. 6. Simple levelwise DAGs. (Left) non-mutually exclusive. (Center) weak-mutually exclusive. (Right) strong-mutually exclusive.

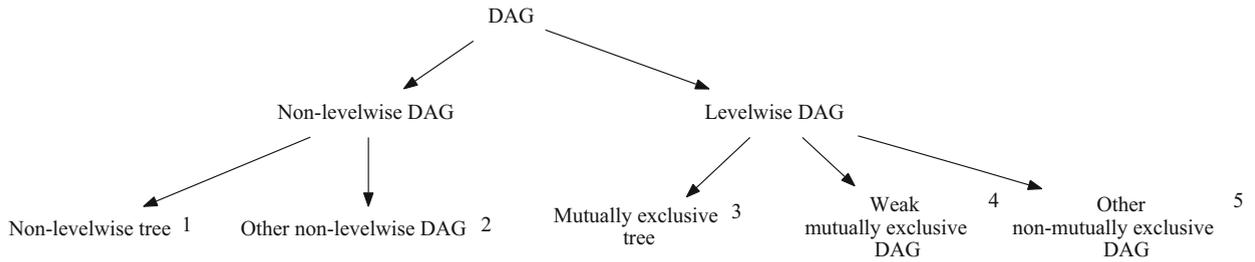


Fig. 7. Partial order of classes of hierarchical hypermedia. Superscripts help connect these classes to those in Table 8.

Table 8

Alternate view of classes of hierarchical hypermedia making the five considered classes (leaves of the tree in Fig. 7) salient.

DAG	non-me	me	Tree	non-me	me
non-lw	$\sqrt{2}$	\perp	non-lw	$\sqrt{1}$	\perp
lw	$\sqrt{5}$	$\sqrt{4}$	lw	\perp	$\sqrt{3}$

Lemma 5.4. Generalized interpretation 1 of out-of-turn interaction over a tree D is functionally equivalent to generalized interpretation 2 of out-of-turn interaction over D .

Lemma 5.5. The interpretations of out-of-turn interaction considered in this article preserve the levelwise property in DAGs.

Lemma 5.6. If a DAG D is a levelwise tree, then D is a strong-mutually exclusive tree.

Corollary 1. The interpretations of out-of-turn interaction considered in this article preserve the strong-mutually exclusive property in trees.

Lemmas 5.2 and 5.6 make the type of mutually exclusivity unequivocal in DAGs. Thus, unless stated otherwise, we refrain from qualifying mutually exclusivity as *weak* or *strong* in such contexts.

Sketches of the proofs of Lemmas 5.1–5.6 are given in [6, Chapter 4]. The proof for 5.4 follows analogously from Lemma 5.3. We can construct its outline by replacing all occurrences of *FP* in the proof of Lemma 5.3 with *SL*, since *FP* is an instance of *SL*. The proof of Corollary 1 follows directly from Lemmas 5.5 and 5.6.

A Hasse diagram for a partial order of these classes of hierarchical hypermedia is shown in Fig. 7. The leaf vertices in this Hasse diagram are the five classes considered here. Examples of sites modeled after non-levelwise trees are GAMS at <http://gams.nist.gov> and the Google directory at <http://directory.google.com>. Examples of other non-levelwise DAGs are the Yahoo! directory and ODP.² PVS and the Kelley Blue Book at <http://kbb.com> each can be modeled as a mutually exclusive tree. The Virginia Tech online timetable of classes can be modeled as a weak mutually exclusive DAG. The mutually exclusive tree class is the primary focus of our user studies [6, Chapter 6; 7]. Note that the other non-mutually exclusive DAG class does not contain any trees, by Lemma 5.6, as shown in Fig. 7. An alternate view of these five classes is given in Table 8.

We describe methods for automatically identifying these classes of hierarchical hypermedia in [6, Chapter 4]. The application of any interpretation of out-of-turn interaction to a DAG may result in a mutually exclusive DAG. Therefore, to take advantage of program transformation techniques specialized for these classes, which we develop below, we need to apply these identification procedures after each out-of-turn interaction.

6. Mining web functional dependencies for automatic query expansion

Before re-entering the program-theoretic domain, we must develop a few more graph-theoretic terms and tools which help provide intuition for the relations to some of the program transformation techniques to follow. Specifically, we illustrate how mining dependencies underlying websites helps expand input for out-of-turn interaction.

A *path-term-co-occurrence set* is a term-co-occurrence set which only contains the terms from D which co-occur with a particular term along paths through D . Build *path-term-co-occurrence set* is total function $BPTC : (D \times T) \rightarrow P(T)$ which given

² Formally, the Yahoo! directory and ODP are not DAGs due to the presence of cycles induced by symbolic links.

D and a term $t \in T = TE(D)$ returns a path-term-co-occurrence set. It is defined as

$$BPTC(D, t) = \bigcup_{I_{t_i} \in GS(t, SQ(D))} GIS(I_{t_i}) - \{t\},$$

where I_{t_i} is an interaction sequence from D containing term t_i . The set {Adjerid, Green} is the path-term-co-occurrence set in the DAG illustrated in Fig. 6 (center) wrt the term ‘Mathematics.’

Given a set of leaves from D reachable through paths containing a particular term t (i.e., $L_t = FP(D, t)$), a *leaf-term-co-occurrence set* is a term-co-occurrence set which only contains the terms from D , excluding t , which lie along all of the paths from D which reach all of the leaves in L_t . *Build leaf-term-co-occurrence set* is total function $BLTC : (D \times T) \rightarrow P(T)$ which given D and term $t \in T = TE(D)$ returns a leaf-term-co-occurrence set. It is defined as

$$BLTC(D, t) = TE(BP(FP(D, t))) - \{t\}.$$

The set {Computer Science, Adjerid, Green} is the leaf-term-co-occurrence set in the DAG illustrated in Fig. 6 (center) wrt the term ‘Mathematics.’

Lemma 6.1. *If a DAG D is a tree, then any path-term-co-occurrence set wrt a term t in D is a leaf-term-co-occurrence set wrt t in D . In other words, if a DAG D is a tree, then $BPTC(D, t) = BLTC(D, t)$.*

A sketch of the proof of Lemma 6.1 is given in [6, Chapter 4].

Corollary 2. *If a DAG D is a tree, then the complete set of path-term-co-occurrence sets in D equals the complete set of leaf-term-co-occurrence sets in D . In other words, if a DAG D is a tree, then*

$$\bigcup_{t_i \in TE(D)} \{BPTC(D, t_i)\} = \bigcup_{t_i \in TE(D)} \{BLTC(D, t_i)\}.$$

The proof for Corollary 2 follows directly from Lemma 6.1.

Identifying structural relationships in data-intensive websites, in domains such as e-commerce (e.g., Amazon), digital libraries (e.g., CITIDEL at <http://citidel.org>), and scientific computing (e.g., GAMS at <http://gams.nist.gov>) is becoming a precursor to personalizing information access [37]. Deploying out-of-turn interaction involves addressing some practical considerations, including the identification of such relationships. Specifically, when the targeted website contains dependencies between and across the levels of its DAG model, an out-of-turn interaction can result in the system soliciting information from the user which can be inferred from the previous input. For instance, consider a user’s interaction with the Kelley Blue Book, an automobile website which progressively solicits for automobile attributes, in an order pre-determined by the site designer, and forces users to communicate those attributes in this manner to access a vehicle webpage. KBB is modeled as mutually exclusive tree. If a user communicates an out-of-turn input which causes all paths save for one to be pruned from the site, then that user is required to click through a series of webpages each only containing one link (leading to the next page). For instance, if saying ‘Civic’ out-of-turn results in one car, then the user may still have to click the ‘sedan’ hyperlink on the resulting page and the ‘Honda’ hyperlink on the penultimate page en route to the leaf webpage.

Classic menu design research states that no menu should contain less than two items. Therefore, in such cases, we would like to consolidate the path and lead the user directly to the terminal webpage in one stroke. Thus, when the user says ‘Civic’ out-of-turn, we can infer ‘sedan Honda’ by functional dependency and safely expand the input to ‘Civic sedan Honda’ without losing any information. The ‘Civic → Honda’ web FD asserts that all of the paths through the site involving ‘Civic’ also involve ‘Honda.’ Clearly, the reverse does not hold as Honda makes several automobile models. We call such a dependency a *positive-path web FD*. We detect such dependencies between the site’s levels (in this case, vehicle-type, -maker, and -model) and leverage these positive web FDs for query expansion. When such dependencies are employed, it is important to provide real-time feedback to users so that the information contained on the rhs of the web FD is not lost. This can be done by collecting the rhs of each web FD triggered and augmenting the leaf webpage with this information. Alternatively, we could present an ‘Input so far:’ label in the browser’s status bar and, at each step in the interaction, incrementally include the rhs of any fired web FDs. Notice that triggering such web FDs affects the stageable interaction sequences and thus the realizable (browsing and out-of-turn) interaction paradigms. Thus, we may have to update our definition of an interaction sequence so that it can be defined over a subset of an interaction set.

Using web FDs to expand user out-of-turn input ultimately creates invisible shortcuts through the website for the user. For example, saying ‘Washington, DC’ out-of-turn at the top level of PVS expands to ‘Washington, DC House Democrat District-at-large’ (because only one congressperson represents Washington, DC) and directly reaches the webpage of an individual congressional official. In levelwise sites, a cursory understanding of the underlying domain is necessary to manually identify a majority of the web FDs. For instance, in KBB, several web FDs fit the *model → make* template (e.g., ‘Civic → Honda’). Web FDs are less salient in sites which are non-levelwise such as large web directories such as *Yahoo!* and *ODP*. Thus, techniques from association rule mining, especially those designed for the web [38–41], become important and applicable in both classes.

6.1. Mining positive web functional dependencies

6.1.1. Positive-path web functional dependencies

Discovering positive-path web FDs entails identifying pairs of hyperlink labels where all of the paths through the site involving one also involve the other. Then when a user supplies out-of-turn input, we consult the precomputed set of web FDs to answer the question: what terms in the complete set of terms, $TE(D)$, lie along each path involving a hyperlink labeled with the out-of-turn input? An instance of such a web FD in KBB is 'Civic → Honda.'

Since the number of paths remaining through the site is reduced as a result of each (in-turn or out-of-turn) interaction, the number of positive-path web FDs satisfied by a site dynamically changes after each interaction. Therefore, while the number of such web FDs (say p) in a site is potentially exponential in its size, just as in traditional data mining algorithms, the number of web FDs reduces dramatically as we drill through the levels. We have developed an algorithm, which is output-sensitive with complexity $O(p)$, to mine all of the possible positive-path web FDs a priori [42].

An alternate approach is to mine a new set of web FDs at each step in the interaction. However, since positive-path web FDs can only exist among terms which co-occur on a path, we need not examine $TE(D)$ terms per term. Rather we can make use of a path-term-co-occurrence set and examine only the terms which co-occur on a path with the particular term (corresponding to the out-of-turn input). This approach implies computing a path-term-co-occurrence set after the user supplies an out-of-turn input and using it to discover any positive-path web FDs relevant to the user's interaction. *Mine positive-path-FD set* is total function $mFDs_{pp} : (D \times T) \rightarrow P(T)$ which given D and a term $t_i \in T = TE(D)$ returns a term-co-occurrence set. It is defined as

$$mFDs_{pp}(D, t_i) = \bigcup_{t_j \in BPTC(D, t_i)} \{t_j\} \text{ such that } GS(t_i, SQ(D)) \subseteq GS(t_j, SQ(D)).$$

This function finds web FDs of the form $t_i \rightarrow t_j$. A *positive-path-FD set* of D is any term-co-occurrence set returned from $mFDs_{pp}$. The complete set of positive-path web FDs satisfied by the DAG in Fig. 6 (left) is

{Ramakrishnan	→	Computer Science
Adjerid	→	Computer Science
Beattie	→	Mathematics
Green	→	Mathematics

No positive-path web FDs hold in the DAG shown in Fig. 6 (right).

6.1.2. Positive-leaf web functional dependencies

Positive-leaf web FDs are more general than positive-path web FDs in that all positive-path web FDs also are positive-leaf web FDs. Discovering positive-leaf web FDs entails identifying pairs of hyperlink labels where all of the leaves classified under paths involving one label also are classified under paths involving the other. Then when a user supplies out-of-turn input, we consult the set of web FDs to answer the question: what terms in the complete set of terms, $TE(D)$, lie along paths which classify all leaves reachable by paths involving a hyperlink labeled with the out-of-turn input? The number of positive-leaf web FDs in a site changes after every interaction because the number of leaves remaining in the site is reduced as a result of each interaction. Thus, we can mine a new set of positive leaf web FDs at each step. Since positive-leaf web FDs can only exist among terms which co-occur wrt leaves, we need not examine $TE(D)$ terms per term. Rather we can make use of a leaf-term-co-occurrence set and examine only the terms which co-occur with a particular term (corresponding to the out-of-turn input) wrt leaves. This approach implies computing a leaf-term-co-occurrence set after the user supplies an out-of-turn input and using it to discover any positive-leaf web FDs relevant to the user's interaction. *Mine positive-leaf-FD set* is total function $mFDs_{pl} : (D \times T) \rightarrow P(T)$ which given D and a term $t_i \in T = TE(D)$ returns a term-co-occurrence set. It is defined as

$$mFDs_{pl}(D, t_i) = \bigcup_{t_j \in BLTC(D, t_i)} \{t_j\} \text{ such that } FP(D, t_i) \subseteq FP(D, t_j).$$

This function mines web FDs of the form $t_i \rightarrow t_j$. A *positive-leaf-FD set* of D is any term-co-occurrence set returned from $mFDs_{pl}$. The complete set of positive-leaf web FDs satisfied by the DAG in Fig. 6 (left) is

{Ramakrishnan	→	Computer Science,
Adjerid	→	Computer Science,
Adjerid	→	Mathematics,
Adjerid	→	Beattie,
Beattie	→	Computer Science,
Beattie	→	Mathematics,
Beattie	→	Adjerid,
Green	→	Mathematics

The DAG shown in Fig. 6 (right) does not satisfy any positive-leaf web FDs.

Lemma 6.2. Any positive-path web FD in a DAG D is a positive-leaf web FD in D .

Corollary 3. Any positive-path web FD set wrt a term t_i in a DAG D is a subset of the positive-leaf web FD set wrt t_i in D . In other words, $mFDS_{pp}(D, t_i) \subseteq mFDS_{pl}(D, t_i)$.

We leave it to the reader to convince thyself that the converse of Lemma 6.2 (i.e., any positive-leaf web FD is a positive-path web FD) is not true. However, we have

Lemma 6.3. If a DAG D is a tree, then any positive-leaf web FD in D is a positive-path web FD in D .

Corollary 4. If a DAG D is a tree, then any positive-path web FD set wrt a term t_i in D equals the positive-leaf web FD set wrt t_i in D . In other words, if D is a tree, then $mFDS_{pp}(D, t_i) = mFDS_{pl}(D, t_i)$.

Corollary 5. If a DAG D is a tree, then the complete set of positive-path web FDs in D equals the complete set of positive-leaf web FDs in D .

Notice that Lemma 6.3 uses *if*, not *iff*; we leave it to the reader to convince thyself that the converse of Lemma 6.3 (i.e., if any positive-leaf web FD in a DAG D is a positive-path web FD in D , then DAG D is a tree) is not true.

Identifying and employing instances of positive web FDs is helpful for usability purposes, but not necessary for realizing out-of-turn interaction. What is non-intuitive, however, is that a related concept, *negative web FDs*, is helpful in developing an alternate program transformation technique for out-of-turn interaction and specializations of it for the mutually exclusive classes of hierarchical hypermedia (when we use the phrase *mutually exclusive classes* we refer to a *mutually exclusive tree* and a *weak mutually exclusive DAG*, not a general connotation to two (or more) classes which are mutually exclusive with each other). Specifically, partially evaluating a programmatic representation of a website wrt the variables representing the terms on the rhs of a *negative-path* or *-leaf* FD set each assigned zero prunes the site.

6.2. Mining negative web functional dependencies

6.2.1. Negative-path web functional dependencies

Identifying negative-path web FDs entails identifying pairs of hyperlink labels where none of the paths through the site involving one label involve the other. Identifying negative-path web FDs involves answering the question: what terms in the complete set of terms, $TE(D)$, *never* lie along any of the paths involving a hyperlink labeled with a particular term? An example of such a web FD in KBB is ‘Honda \rightarrow Toyota.’ Intuitively, this means that none of the paths through the site involving the term ‘Honda’ involve the term ‘Toyota.’ *Mine negative-path-FD set* is total function $mFDS_{np} : (D \times T) \rightarrow P(T)$ which given D and a term $t_i \in T = TE(D)$ returns a term-co-occurrence set. It is defined as

$$mFDS_{np}(D, t_i) = TE(D) - \{t_i\} - BPTC(D, t_i).$$

A *negative-path-FD set* of D is any term-co-occurrence set returned from $mFDS_{np}$. The complete set of negative-path web FDs satisfied by the DAG in Fig. 6 (left) is

```
{Computer Science → ¬ Mathematics,
Computer Science → ¬ Beattie,
Computer Science → ¬ Green,
Mathematics → ¬ Computer Science,
Mathematics → ¬ Ramakrishnan,
Mathematics → ¬ Adjerid,
Ramakrishnan → ¬ Mathematics,
Ramakrishnan → ¬ Adjerid,
Ramakrishnan → ¬ Beattie,
Ramakrishnan → ¬ Green,
Adjerid → ¬ Mathematics,
Adjerid → ¬ Ramakrishnan,
Adjerid → ¬ Beattie,
Adjerid → ¬ Green,
Beattie → ¬ Computer Science,
Beattie → ¬ Ramakrishnan,
Beattie → ¬ Adjerid,
Beattie → ¬ Green,
Green → ¬ Computer Science,
Green → ¬ Ramakrishnan,
Green → ¬ Adjerid,
Green → ¬ Beattie}.
```

Likewise, the complete set of negative-path web FDs satisfied the DAG in Fig. 6 (right) is

```
{California → ¬ Virginia,
Virginia → ¬ California,
House → ¬ Senate,
Senate → ¬ House}.
```

6.2.2. Negative-leaf web functional dependencies

Identifying negative-leaf web FDs entails identifying pairs of hyperlink labels where none of the leaves classified under paths involving one label are classified under paths involving the other. Identifying negative-leaf web FDs involves answering the question: what terms in the complete set of terms, $TE(D)$, do not lie along paths from the root of D to its leaves which classify any of the leaves reachable by paths involving a hyperlink labeled with a particular term? An example of such a web FD in KBB is 'Honda \rightarrow \neg Toyota.' Unlike the discussion of negative-path web FDs, this type of web FD must be interpreted as 'none of the leaves classified by paths involving the term 'House' are classified by the paths involving the term 'Senior seat.' *Mine negative-leaf-FD set* is total function $mFDS_{nl} : (D \times T) \rightarrow P(T)$ which given D and a term $t_i \in T = TE(D)$ returns a term-co-occurrence set. It is defined as

$$mFDS_{nl}(D, t_i) = TE(D) - \{t_i\} - BLTC(D, t_i).$$

A *negative-leaf-FD set* of a DAG D is any term-co-occurrence set returned from $mFDS_{nl}$. The complete set of negative-leaf web FDs satisfied by the DAG in Fig. 6 (left) is

{Computer Science	\rightarrow	\neg Green,
Mathematics	\rightarrow	\neg Ramakrishnan,
Ramakrishnan	\rightarrow	\neg Mathematics,
Ramakrishnan	\rightarrow	\neg Adjerid,
Ramakrishnan	\rightarrow	\neg Beattie,
Ramakrishnan	\rightarrow	\neg Green,
Adjerid	\rightarrow	\neg Ramakrishnan,
Adjerid	\rightarrow	\neg Green,
Beattie	\rightarrow	\neg Ramakrishnan,
Beattie	\rightarrow	\neg Green,
Green	\rightarrow	\neg Computer Science,
Green	\rightarrow	\neg Ramakrishnan,
Green	\rightarrow	\neg Adjerid,
Green	\rightarrow	\neg Beattie).

The complete set of negative-leaf web FDs satisfied by the DAG D in Fig. 6 (right) is (the complete set of negative-path web FDs satisfied by the D)

{California	\rightarrow	\neg Virginia,
Virginia	\rightarrow	\neg California,
House	\rightarrow	\neg Senate,
Senate	\rightarrow	\neg House}.

Note that the \neg symbol is always on the rhs of a negative web FD. Notice further that negative (-path and -leaf) and positive (-path and -leaf) web FDs are not complements of each other, i.e., the presence of $x \rightarrow y$ does not necessarily imply the presence of $x \rightarrow \neg\{TE(D) - y\}$. Similarly, a site which satisfies $x \rightarrow \neg y$ may not satisfy $x \rightarrow \{TE(D) - y\}$.

Lemma 6.4. Any negative web FD $x \rightarrow \neg y$ considered here also implies $y \rightarrow \neg x$.

Lemma 6.5. Any negative-leaf web FD in a DAG D is a negative-path web FD in D .

Notice that the claim in Lemma 6.5 is the contrapositive of the claim in Lemma 6.2.

Corollary 6. Any negative-leaf-FD set wrt a term t_i in D is a subset of the negative-path-FD set wrt t_i in D . In other words, $mFDS_{nl}(D, t_i) \subseteq mFDS_{np}(D, t_i)$.

We leave it to the reader to convince thyself that the converse of Lemma 6.5 (i.e., any negative-path web FD is a negative-leaf web FD) is not true.

Lemma 6.6. If a DAG D is a tree, then any negative-path web FD in D is a negative-leaf web FD in D .

Corollary 7. If a DAG D is a tree, then any negative-path-FD set wrt a term t_i in D equals the negative-leaf-FD set wrt t_i in D . In other words, if D is a tree, then $mFDS_{np}(D, t_i) = mFDS_{nl}(D, t_i)$.

Corollary 8. If a DAG D is a tree, then the complete set of negative-path web FDs in D equals the complete set of negative-leaf web FDs in D .

Sketches of the proofs of Lemmas 6.2–6.6 are given in [6, Chapter 4]. The proof of Corollary 3 follows directly from Lemma 6.2. The proof of Corollary 4 follows directly from Lemmas 6.2 and 6.3, and the proof of Corollary 5 follows directly from Corollary 4. The proof of Corollary 6 follows directly from Lemma 6.5, and the proof of Corollary 7 follows directly from Lemmas 6.5 and 6.6.

Any DAG without an interaction set containing all terms from the site it model satisfies negative web FDs [36]. Moreover, any DAG with more than one distinct term satisfies either positive or negative web FDs or both [36].

We have developed algorithms for mining web FDs as well as a theory of reasoning with web FDs. We have used the algorithm and theory to mine web FDs from a variety of websites to demonstrate that web hierarchies are teeming with

dependencies reflecting the nature of the underlying domain and how they can be exploited in a user information-seeking interface for automatic query expansion [42].

7. Mining web functional dependencies by program transformation

Interestingly, the classes of web FDs defined above not only can be exploited by program transformation techniques, but also can be mined through program transformation.

The following program transformation technique can mine a positive-path-FD set wrt a particular term:

$$\llbracket \text{ppfd} \rrbracket [P_D, \text{input}] = \bigcup_{t_i \in \llbracket \text{bptc} \rrbracket [P_D, \text{input}]} \{t_i\}, \text{ if } (\llbracket \text{zoom} \rrbracket [P_D, \text{input}] \cap \llbracket \text{zoom} \rrbracket [P_D, t_i]) = \llbracket \text{zoom} \rrbracket [P_D, \text{input}],$$

where $\llbracket \text{bptc} \rrbracket = \llbracket \text{te} \rrbracket [\llbracket \text{sp} \rrbracket [P_D, \text{input}] - \text{input}]$ and is the program transformation technique we associate with *BPTC*. Notice that the above expression closely reflects the textual definition of a positive-path web FD. Intuitively, this program transformation technique mines positive-path web FDs of the form $\text{input} \rightarrow t_i$. We can mine a positive-leaf-FD set analogously:

$$\llbracket \text{plfd} \rrbracket [P_D, \text{input}] = \bigcup_{t_i \in \llbracket \text{bltc} \rrbracket [P_D, \text{input}]} \{t_i\}, \text{ if } (\llbracket \text{zoom} \rrbracket [P_D, \text{input}] \cap \llbracket \text{zoom} \rrbracket [P_D, t_i]) = \llbracket \text{zoom} \rrbracket [P_D, \text{input}],$$

where $\llbracket \text{bltc} \rrbracket = \llbracket \text{te} \rrbracket [\llbracket \text{zoom} \rrbracket [P_D, \text{input}] - \text{input}]$ and is the program transformation technique we associate with *BLTC*. Similarly, this program transformation technique mines positive-leaf web FDs of the form $\text{input} \rightarrow t_i$. Since identifying negative web FDs does not involve containment, program transformation techniques for doing so are more simplistic than those for mining positive web FDs. Mining a negative-path-FD set entails using $\llbracket \text{te} \rrbracket$ to extract all of the conditional variables from the programmatic complement of $\llbracket \text{sp} \rrbracket [P_D, \text{input}]$:

$$\llbracket \text{npfd} \rrbracket [P_D, \text{input}] = \llbracket \text{te} \rrbracket [P_D - \llbracket \text{sp} \rrbracket [P_D, \text{input}]],$$

where the $-$ symbol (i.e., minus sign) means complement. Similarly, discovering a negative-leaf-FD set entails using $\llbracket \text{te} \rrbracket$ to extract all of the conditional variables from the complement of $\llbracket \text{zoom} \rrbracket [P_D, \text{input}]$:

$$\llbracket \text{nlfld} \rrbracket [P_D, \text{input}] = \llbracket \text{te} \rrbracket [P_D - \llbracket \text{zoom} \rrbracket [P_D, \text{input}]].$$

8. A duality in uses of program slicing

We have used partial evaluation and program slicing primarily as pruning operators. There is a relationship and tradeoff between these two program transformations in the context of our research. Specifically, one can think of program slicing as a transformation for

1. directly pruning a website or
2. extracting information about what to prune from the website and partially evaluating this information away to conduct the same site pruning in (1).

Table 7 (second row) illustrates how program slicing can play role (1) above to realize interpretation 1 of out-of-turn interaction. In the previous section, we show that program slicing can also be used to mine web FDs. The members of the negative-leaf-FD set wrt some relevant term t label only the edges to be pruned from a DAG model of a website when the user communicates term t out-of-turn. Therefore, partially evaluating wrt each variable modeling each term in the negative-leaf-FD set of a particular term set to zero is sufficient to prune the undesired paths from a website to realize interpretation 1 of out-of-turn interaction.

The following expression captures these two methods of realizing interpretation 1 of out-of-turn interaction and also captures the dual role played by program slicing, i.e., role (1) above is played by $\llbracket \text{zoom} \rrbracket$ on the lhs and role (2) above is played by $\llbracket \text{nlfld} \rrbracket$ on the rhs:

$$\llbracket \text{mix} \rrbracket [\llbracket \text{zoom} \rrbracket [P_D, \text{input}], \text{input} = 1] \equiv \llbracket \text{mix} \rrbracket [P_D, \llbracket \text{nlfld} \rrbracket [P_D, \text{input}] = 0, \text{input} = 1],$$

where the notation

$$\llbracket \llbracket \text{nlfld} \rrbracket [P_D, \text{input}] = 0, \rrbracket$$

denotes ‘the assignment of zero to each variable modeling each member of the negative-leaf-FD set of the term modeled by the variable input .’

We use the word *sufficient* above, rather than *necessary*. This is because all terms labeling edges along the tails of paths to prune, when the user communicates term t out-of-turn, are members of t 's negative-leaf-FD set. However, to remove that tail, we need not partially evaluate wrt all variables modeling the terms along that tail set to zero. We need only partially evaluate wrt the variable modeling the term with the minimum depth set to zero. Such a partial evaluation removes the remainder of the path by default. Intuitively, this means that we need only consider the members of the negative-leaf-FD which have the minimum depth in each path through a DAG. Considering this optimization compromises

the clarity of the equivalence which we have just outlined as well as the duality in uses of program slicing. Therefore, we do not incorporate it into our presentation and expression. Also notice that when dealing with trees, we can replace $\llbracket \text{n1fd} \rrbracket$ with $\llbracket \text{npfd} \rrbracket$ in the above expression by Corollary 7.

Now let us consider how we might develop an analogous equivalence to the program transformation technique given in Table 7 (third row) for interpretation 2 of out-of-turn interaction. Recall to support interpretation 2 we replace $\llbracket \text{zoom} \rrbracket$ with $\llbracket \text{sp} \rrbracket$ on the lhs of the above equivalence expression. Thus, it is natural to consider replacing $\llbracket \text{n1fd} \rrbracket$ with $\llbracket \text{npfd} \rrbracket$ on the rhs of the above expression to form an analogous equivalence expression for interpretation 2 of out-of-turn interaction. However, upon careful examination we see that the rhs is not equal to the lhs and, thus, does not realize interpretation 2 of out-of-turn interaction. This is because if P_D models a non-tree DAG, then our use of partial evaluation, as described above, to remove undesired paths may also remove leaves which lie at the end of desired paths! For example, consider partially evaluating a programmatic representation of the website illustrated in Fig. 6 (left) wrt the negative-path-FD set of 'Adjerid.' This transformation removes the source (leaf) of the edge labeled 'Adjerid' which is unfaithful to interpretation two. However, the equivalence expression for interpretation 1 can serve as a surrogate equivalence expression for interpretation 2 of out-of-turn interaction with trees (by Corollary 4), where $\llbracket \text{zoom} \rrbracket$ and $\llbracket \text{n1fd} \rrbracket$ can be replaced by $\llbracket \text{sp} \rrbracket$ and $\llbracket \text{npfd} \rrbracket$, respectively. We are unable to develop equivalence expressions for the generalizations of these interpretations of out-of-turn interaction here due to their template-oriented nature, i.e., the presence of $\llbracket \text{SL} \rrbracket$.

Studying this duality reveals that there might be simpler methods for discerning which branches must be removed in specialized DAG classes. The following program transformation technique realizes out-of-turn interaction, but is specifically tailored toward taking advantage of the mutually exclusive property in DAGs and trees:

$$\llbracket \text{dead-code} \rrbracket \llbracket \llbracket \text{mix} \rrbracket [P_D, \llbracket \llbracket \text{tle} \rrbracket [P_D, \text{input}] - \text{input}] = 0, \text{input} = 1 \rrbracket,$$

where $\llbracket \text{dead-end} \rrbracket$ represents a dead-code detection and elimination transformation [43,44]. Static slicing methods can detect dead-code [45]. Detection is typically conducted in a debugging context since such statements are usually unexecutable due to the presence of a bug [14]. This approach illustrates that partial evaluation is a specialization of program slicing wrt programmatic representations of levelwise, mutually exclusive DAG models of websites. Here $\llbracket \text{tle} \rrbracket$ supports *TLE* programmatically and involves extracting all the structural variables at a particular level of nesting. It is intended to be polymorphic in that its input can be either given extensionally as a level number or intensionally as a term which occurs at the desired level. Notice that the latter usage is unambiguous only in the case of mutually exclusive classes (which are the only classes it is used for here).

Other researchers have echoed similar connections between partial evaluation and slicing; the two techniques have been shown to yield similar results in some situations and different results in others [46].

9. Program transformations for supplemental personalized interactions

To target our formal model for representing and reasoning about personalized interaction in hierarchical websites, we illustrate how the formalisms can be augmented to support supplementary personalized interactions. Moreover, the success of a personalization system relies on those finer touches which deliver a compelling experience to the end-user. Studies of out-of-turn interaction have revealed that users desire supplemental interactions to enhance the personalized experience while interacting with hierarchical hypermedia [6, Chapter 6]. We showcase a suite of such interactions and relate them to program transformation techniques.

9.1. Meta-enquiry: What may I say?

To employ an interactive information system effectively, users need a mechanism to stay abreast of its underlying vocabulary during the interaction. Keeping users attuned to the communicable information is an issue all information systems must address. Yankelovich echoes this issue as 'how do users know what to say?' [47]. While each hyperlink label is always available, a new (or causal) user of a site may be unfamiliar with or unaware of subsequent solicitations for input and, hence, the terms of information seeking the designer has modeled which are available to supply out-of-turn. This is a classic problem in IR research, has been identified and described by many [20,48–51], and is endemic to all IR systems. Users typically have a 'limited knowledge of a given database' [51] and thus experience 'difficulty expressing their information need' [48].

The *naive what may I say?* interaction permits the user to determine what partial information remains unspecified thus far in her interaction with a system. It can be trivially supported by *TE*. Supporting this interaction programmatically entails using $\llbracket \text{te} \rrbracket$ to extract all unique structural variables, which correspond to the partial information available to commit the system, from the representation.

In a voluminous space such a set of terms may be large and overload the user and, therefore, we provide a similar interaction, called *what may I say?*, which entails clustering terms by level (facet), using *TLE*, to help orient users. This approach is, of course, applicable only to levelwise sites. When used in this manner, this interaction also updates the user on how their previous interactions have affected the remaining choices and, thus, provides context. Supporting this interaction programmatically entails using $\llbracket \text{tle} \rrbracket$ to extract all unique structural variables at a desired level of nesting from the representation.

9.2. Restructure classification

Notice that while an out-of-turn interaction paradigm subsumes the corresponding browsing paradigm (see Lemma 4.1), interaction sequences *derivable* through out-of-turn interaction are not *describable* by browsing. The *restructure classification* interaction, which is only applicable to levelwise sites, enables the creation of a personalized browsing hierarchy, which can be further navigated by browsing or out-of-turn interaction. For example, a website organized along an author-journal-title motif could be restructured into a journal-author-title organization to support interactive aggregation scenarios.

Approaches to this interaction are to permit a user to restructure an entire enumerative classification a priori or incrementally [52,53]. The semistructured data community [37,54] has advocated restructuring websites through declarative queries [55,56]. Such an approach restructures an entire organization in one-stroke. A similar approach is taken in *user-defined hierarchies* [57]. An alternate approach is to permit the user to define the ensuing interaction incrementally while browsing. This interaction style has been echoed metaphorically as ‘magically [laying] down track to suggest useful directions to go based on where [one has] been so far and what [one is] trying to do’ [58]. These ideas have been studied in the *Flamenco Search System Project* [59]. *Flamenco* explores faceted classification in various catalogs and websites. The adaptive hypermedia community [48,60] is a proponent of such an incremental approach to classification specification. In adaptive hypermedia, links are dynamic [61] and lead to different destinations for different users.

Counterintuitively, our nested-conditional representation of interaction need not be restructured at all to accommodate this interaction! A method to extract structural variables corresponding to a particular level of conditional nesting is sufficient. These variables can be used to create hyperlinks on a generated webpage, each initiating an interaction which appears to be browsing, but is actually out-of-turn interaction wrt the program representation. This, yet again, emphasizes the importance of investing in representation, and the flexibility in our programmatic representations.

Thus, support for this interaction requires level-order, edge-label extraction, *TLE*. Supporting this interaction programmatically entails the same code extraction used for automatic query expansion in mutually exclusive classes of hierarchical hypermedia, i.e., $\llbracket \text{tle} \rrbracket$. Notice that the edge-label extraction for this interaction is wrt one particular level and is, thus, a specialization of the edge-label extraction required for (*naive*) *what may I say?* which is wrt all levels. In other words, the terms returned from *TLE* are a subset of those returned from *TE*.

9.3. Collect results

This is a termination interaction which allows the user to request that the classification be flattened and re-presented as a flat list of hyperlinks to relevant content pages. For instance, while drilling-down a hierarchical website a user may wish to curb further interaction at a particular point and retrieve a flat list of the terminal webpage reachable from that point. At such a point, the user may not care to pursue further distinctions or dependencies. Effectively, the user has prematurely declared that the interaction is over. This interaction involves the forward-propagation used to support out-of-turn interaction and is defined by *CR* introduced above. We support *CR* programmatically with the $\llbracket \text{cr} \rrbracket = \llbracket \text{forward} \rrbracket_{[P_D, \text{page}]}$ program transformation technique—a specialization of $\llbracket \text{zoom} \rrbracket$ which only extracts terminal variables from the representation. Notice that *collect results* suggests the need to invert the factoring of terminal information which we built into the programmatic representation with nesting. Observe that the function which supports *collect results* also can be used to generate *information previews*—hyperlink annotations, typically parenthesized and containing a frequency count of the documents reachable from the hyperlink it decorates—which have been shown important in situations where the user is confronted with a decision regarding where to go next [62]. Such previews are now a de facto standard in most faceted browsing and search systems and web directories.

9.4. Inverse personalization

Inverse personalization is so named not because its goals run contrary to personalization, but because in a given representation it conducts a mapping from terminal (e.g., leaf pages) to structural information (e.g., hyperlink labels). This interaction helps support situations where the user knows what information she wants, but is interested in determining how to retrieve that information and, in this manner, permits users to conduct ‘what if’ analyzes. For instance, using *inverse personalization* with an online apartment recommender system, a user can issue a request such as ‘under what conditions will the Georgetown apartments be the only choice?’ Such a user is interested in the terms along the interaction sequences leading to that terminal information (e.g., if you want a swimming pool, covered parking, and free wi-fi). Such personalized interactions are supported in our approach by back-propagating from the leaf webpages returned by *SL* and optionally extracting each term in each path from the root as a whole, using *TE*, or per interaction sequence, using *TE o SP*.

Inverse personalization is similar to generalized interpretation 1 of out-of-turn interaction, as evidenced by its use of *SL*. However, akin to *collect results*, this interaction is a specialized form of out-of-turn interaction. Honing in on leaves (i.e., forward-propagation) is unnecessary; back-propagation is sufficient and, therefore, defines this interaction. In a nested conditional representation, terminal information is indexed by a terminal variable (e.g., *page*) which is not user-modifiable (as is structural information, e.g., *software*). Therefore, we need a transformation capable of exploiting terminal variables

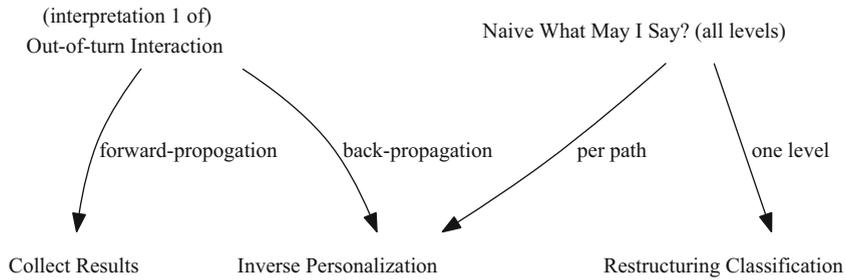


Fig. 8. A taxonomy of supplementary interactions. Directed arrows represent specialization relations.

Table 9

Program transformation techniques for and observations on the supplemental personalized interactions.

Sup Int	Read-only	Closed	Ind	Pgm Trans Tech
NWhat?	✓	✓	×	$\cup\{\llbracket te \rrbracket\llbracket P_D, x \rrbracket\}$, for every structural variable x
What?	✓	✓	×	$\cup\{\llbracket t1e \rrbracket\llbracket P_D, x \rrbracket\}$, for every level x
RC	✓	✓	✓	$\llbracket t1e \rrbracket\llbracket P_D, x \rrbracket$, for one level x
CR	×	×	✓	$\llbracket forward \rrbracket\llbracket P_D, page \rrbracket$
P^{-1}	✓	✓	✓	$\llbracket te \rrbracket\llbracket sp \rrbracket\llbracket backward \rrbracket\llbracket SL \rrbracket\llbracket P_D, input \rrbracket\rrbracket$

Sup Int: supplemental interaction, Ind: independent, Pgm Trans Tech: program transformation technique, NWhat: *naive what may I say?*, What?: *what may I say?*, RC: *restructuring classification*, CR: *collect results*, and P^{-1} : *inverse personalization*.

such as backward slicing. Collecting the terms along the resulting paths requires an extraction similar to that required for *naive what may I say?*, i.e., $\llbracket te \rrbracket$.

Fig. 8 illustrates the relationships between these supplemental personalized interactions.

Table 9 contains program transformation techniques for these supplemental personalized interactions. Notice that the *what may I say?* and *restructuring classification* interactions require slight variations of the same transformation technique. Since *collect results* and *inverse personalization* are each specializations of out-of-turn interaction, they relate to forms of forward and backward slicing, respectively (and code extraction) and, thus, when combined come close to realizing out-of-turn interaction.

Table 9 also summarizes our observations on these personalized interactions. The top row of the matrix lists attributes of the interactions listed on the far left. A read-only interaction is one which only manipulates the representation rather than modifying it. An interaction is closed if it accepts a representation and returns a representation and, thus, does not curb further interaction. Independent interactions are those which not only complement out-of-turn interaction, but are also defined, applicable, and useful in its absence.

10. Related research

Three main ideas underly the work presented here:

- explicit modeling of interaction,
- capture of partial information, and
- use of program transformations.

Any information systems context where one or more of the above ideas apply is ground for applying the techniques presented in this article.

Modeling interaction has been recognized as important in multiple interactive information retrieval projects, especially to manage interaction and maintain state. Belkin et al. [63] introduced the idea of an *interaction script* which can be thought of as a program for interaction, though expressed in English rather than program codes, and is intended to be interpreted. Graunke et al. [9] describe an approach to automatically restructure batch programs for interactive use on the web. An important issue addressed is maintaining state across web interactions which use the stateless HTTP protocol. The approach involves capturing first-class continuations [64] through the `call/cc` (call-with-current-continuation) facility in Scheme. Since first-class continuations can be saved and resumed, they are an ideal construct for capturing and restoring state between user interactions over the web. Using a similar idea based on continuations, Queinnee [10] developed a model for a web server intended to address state maintenance problems caused by connections terminated prematurely,

pressing the back button, and window cloning. Moreover, Quan [65] use continuations and currying [66] with explicit representations to postpone, save, and resume interactions with intrusive dialog boxes, including partially-filled boxes, in traditional application software, such word processors and e-mail clients. The common theme of these efforts, including our research, is the appeal to concepts from programming languages to achieve a rich and expressive form of a human-computer interaction. Our work differs from these efforts in its focus on out-of-turn interactions. Nevertheless, these projects reinforce our viewpoint that information system design can benefit from investing in representations of interaction, and exploring concepts from programming languages.

Until recently, the issue of context in information systems has received comparatively little attention. With the proliferation of mobile environments and information appliances [67], coupled with an improved understanding of their usages [68], capturing and reusing context has become important. Context can be viewed as a rich form of partial information, allowing the work presented here to be applied toward exploiting it. This property is attractive because the partial information can potentially be multi-faceted—information about a user's preferences (e.g., in a recommender system), a user's location (e.g., in a mobile environment; [69]), or a user's partially completed interaction (e.g., a shopping cart at Amazon). The representations studied in this article can be generalized to accommodate such richer forms of partial information.

Finally, the use of program transformations as presented here is relatable to the larger community which aims to systematize and automate the software engineering of complex, interactive web applications [9]. It finds relevance in many website re-structuring and re-engineering efforts [30,70–72], especially in the adaptive and semantic web contexts, declarative site specification [56], and improving site usability [73]. Work [74,75] has also been done to incorporate program transformations for personalization into models for digital libraries, e.g., 5S [76]. While there are established and effective models for classical information retrieval, e.g., vector-space [8], models for solutions to information personalization problems are in their infancy. We believe that program transformations, which are under-explored in the personalization community, suggest helpful metaphors for developing such models.

11. Discussion

We implemented the graph-theoretic transformations described here in ML (*Meta-Language*) using SML/NJ [66] and explored the feasibility of the website-program associations by conducting verification experiments on programmatic representations of interaction with PVS, ODP, and the Virginia Tech online timetable of courses using program transformation software systems [3]. The *set calculator of Code-Surfer* [25], a program slicing system, was especially helpful for performing set-theoretic operations over slices. We discuss the details of the program transformation systems we used in [6, Appendix C].

11.1. Putting it all together

We have developed a tool and markup language for automatically generating web user interfaces with support for a combination of the supplemental personalized interactions introduced in Section 9 [77]. Moreover, the following three software components, and an interaction manager for coordination, constitute a customizable software framework for creating web personalization systems which support mixed-initiative interaction [3]:

- *Web transformation engine*: A web service which prunes a hierarchical website when given out-of-turn input. The engine is based on the program transformation techniques in Table 7 and implements a form of forward followed by backward program slicing using XSLT. Since it handles in-turn and out-of-turn inputs (captured by interaction interfaces below) in a uniform manner, it enables out-of-turn interaction without requiring the site designer to anticipate the points at which out-of-turn interaction can happen. Using a single transformation technique this engine supports out-of-turn interaction with both faceted and un-faceted sites.
- *Speech UI*: A voice interface, implemented with VoiceXML and X+V, which permits the user to supply out-of-turn inputs through speech and enables multimodal interaction when used in conjunction with hyperlinks.
- *Extempore*: A cross-platform toolbar plugin embedded into the Mozilla Firefox web browser; the analog of the speech user interface using a textual modality.

Extempore [2] and a variety of alternate interaction interfaces [52,53] are available for demonstration from our group's webpage at <http://oot.cps.udayton.edu>. We have instantiated this framework for the following case studies which afford interactions initiated by these interfaces, coordinated by this interaction manager, and staged by this transformation engine:

- GAMS (Guide to Available Mathematical Software) at <http://gams.nist.gov> [33,78],
- Project Vote Smart at <http://votesmart.org> [2,3],
- Pigments through the Ages at <http://webexhibits.org/pigments/>,
- CITIDEL (computing and information technology interactive digital educational library) at <http://citidel.org> [75,77],

- Open Directory Project at <http://dmoz.org> [3], and
- Online Virginia Tech Timetables of Classes accessible through <http://vt.edu> [42].

These sites are an assortment of the classes of hierarchical hypermedia introduced in Section 5.

11.2. Contributions

Our research, and specifically out-of-turn interaction, is part of a larger research effort which seeks to marry navigational (e.g., the *Yahoo!* directory) and direct search (e.g., *Google*) [79]. Our research has made the following contributions to the interactive information retrieval, personalization, dialog management, and programming languages communities.

- *New approach to personalization*: We applied program transformations, seemingly unrelated techniques predominately used in compilers and debuggers, to a problem in interactive information retrieval. While viewing a website as a program and website transformation for personalization as program transformation is unorthodox, it offers a new way of thinking about personalized interaction, especially with hierarchical hypermedia. While others have studied personalization in information hierarchies using more traditional approaches, e.g., Dalamagas et al. [80] take a data mining approach to this problem, to the best of our knowledge, no one has taken a program transformation approach to personalizing interaction in hierarchical domains. This new approach lead to the following contributions.
- *Unifies search in faceted and unfaceted domains*: Our nested conditionals representation of hierarchical hypermedia can be used to model faceted sites (e.g., PVS) and unfaceted sites (e.g., ODP). As a result, a single program transformation technique, e.g., any technique in Table 7, can be used to personalize interaction in both domains and, thereby, unifies search with two fundamentally distinct forms of hypermedia. Currently, there are toolkits available for building faceted browsing and search applications (e.g., *Flamenco* at <http://flamenco.berkeley.edu>). However, search over unfaceted spaces such as the *Yahoo!* directory or ODP is different. We can personalize both spaces with the same transformation using our approach.
- *New line of dialog management research*: It is important to note that our research aims to model a user's *interaction* with a site, rather than simply a site schema. While traditional data-oriented approaches, such as XML or a relational database, might be appropriate for the latter, programming languages offer attractive constructs (e.g., booleans, conditionals, functions, `goto`s) for modeling interaction as we have demonstrated above. While it is easy to presume that modeling a user's interaction with a hierarchical website as a program is equivalent to modeling it as an XML document or in a database relation, there is an important distinction between the two. An imperative program, like the only possible style of interaction with the static, rigid, one-size-fits-all organization of an information hierarchy, has a default order of execution; there is no such analog in an XML document³ or database relation. This distinction is significant because it compelled us to explicitly recognize orders of execution and identify alternate domains where order is relevant. One such domain is human-computer dialogs; often a dialog script also has a default order of evaluation. Thus, our experience using programs to represent interaction with hierarchical sites helped connect our work to human-computer dialog management and, in particular, mixed-initiative interaction [31]. Specifically, we recognized that supporting the user in interleaving in-turn and out-of-turn inputs while browsing a website is simple form of mixed-initiative interaction [2] and were compelled to study computational models for human-computer dialogs [81]. This connection may not have been fostered had we approached personalized interaction from a more traditional perspective. Dialogs are now pervasive in computer systems and used in automated teller machines and airport kiosks. While the fluidity inherent in dialog is essential to provide a natural experience to the user, it is also a vice for the implementor due to the directions in which the user might steer the dialog, which must all be captured in the implementation. These problems are difficult since dialogs range in complexity from those modeled after a simple pre-defined series of questions and answers (sometimes referred to as slot-fillers) to those which give the user a great deal of control over the direction in which to steer the dialog [31,82]. The connection we have fostered is that a dialog script, like a program, has a default order of execution, and applying program transformations alters that order, thereby, making the dialog flexible without having to explicitly hardcode all supported deviations from the default order within the script itself. This is a primary contribution of our research. As a result, we have also discovered that program transformations are helpful for specifying dialogs [83]. Putting personalized web interaction on a fundamentally different landscape (i.e., program transformation) gave birth to this new line of research [7,83,84] which may not have come to fruition had we started with more traditional approaches (XML, databases).
- *Technology assimilation by reduction to theoretical principals*: Using the ideas described here, many relevant interactive applications and dialog standards can be re-visited and re-studied by their support for personalizing interaction. For instance, consider the VoiceXML markup language designed to simplify the construction of interactive voice response

³ We do not intend to imply that in using the program transformation approach we cannot or should not implement the transformation techniques using tools such as XML/XSLT. Rather, we are saying that we may not have made the connection to dialogs had we started out with such a database approach.

systems [85]. VoiceXML markup tags describe prompts, forms, and fields which constitute a dialog, and support both directed and mixed-initiative dialogs. Ramakrishnan et al. [84] have shown that VoiceXML's form-interpretation algorithm is a partial evaluator in disguise. Such connections are noteworthy because they reduce rapidly emerging and constantly evolving technologies to fundamental theoretical operations and, therefore, improve assimilation of such technologies.

- *Towards a model for information personalization*: Our use of program transformations casts personalization in a formal setting and provides a systematic and implementation-neutral approach to designing information personalization systems, especially those in hierarchical hypermedia domains with support for mixed-initiative interaction. Moreover, flexibility is built into the approach. When a given program (representation, transformation) pair is deemed inadequate to achieve the desired form of personalization, we have two choices: keep the representation fixed and investigate or design alternative program transformation techniques capable of realizing the desired interaction or, alternatively, develop representations wrt a fixed suite of transformation techniques. Conversely, we can develop new forms of personalization by creatively applying program transformations to representation. This ability to vary the elements of a model not only suggests the extensibility of the methodology, but also constitutes the creativity in our research.

11.3. Future work

The interpretations for out-of-turn interaction presented here are not exhaustive, especially since we defined open-ended generalized interpretations of out-of-turn interaction. For instance, we might expand the scope of addressable out-of-turn information by modeling the terms on the leaf webpages (i.e., terminal information) and making them available as out-of-turn input for the user to supply. In such case, we might implement *SL* using a text-based search engine. Nevertheless, use of a general program transformation, such as slicing, suggests that alternate interpretations also can be accommodated. For instance, we intend to explore additional program transformations, such as generalized partial evaluation [86,87] and parameterized partial evaluation [88], as well as other related programming language concepts, including reflection [89], concept assignment [90], and program schemas [91], and study the personalized interactions these techniques might enable; see [6, Appendix B] for more information.

We also intend to develop a computational model, based on first-class closures and continuations, for simplifying the implementation of mixed-initiative human-computer dialogs, and develop algorithms for automatically generating an optimal dialog script from a high-level specification of a dialog. While first-class closures and continuations are powerful programming constructs, they are under-utilized outside of compilers and interpreters.

Our long term research goal is to formalize and simplify the design and implementation of dialog-based systems, especially those involving mixed-initiative human-computer interactions. The approach involves empirically studying user interactions with systems and formally casting those interactions through a programming languages lens.

Acknowledgments

We thank Travis Suel at the University of Dayton for re-designing and -implementing the out-of-turn framework and multiple interaction interfaces, including *Extempore*, and for his design and implementation of new interfaces. John Cresencia at the University of Dayton also provided software support in a variety of ways.

References

- [1] Perugini S, Ramakrishnan N. Personalizing interactions with information systems. In: Zelkowitz M, editor. *Advances in Computers*, vol. 57. Amsterdam: Academic Press; 2003. p. 323–82.
- [2] Perugini S, Ramakrishnan N. Personalizing web sites with mixed-initiative interaction. *IEEE IT Professional* 2003;5(2):9–15.
- [3] Narayan M, Williams C, Perugini S, Ramakrishnan N. Staging transformations for multimodal web interaction management. In: Najork M, Wills C, editors. *Proceedings of the 13th international ACM world wide web conference (WWW13)*. New York, NY: ACM Press; 2004. p. 212–23.
- [4] Perugini S, Ramakrishnan N. A generative programming approach to interactive information retrieval: insights and experiences. In: Glück R, Lowry M, editors. *Proceedings of the fourth international ACM conference on generative programming and component engineering (GPCE)*. Lecture notes in computer science, vol. 3676. Berlin: Springer; 2005. p. 205–20.
- [5] Binkley D, Gallagher K. Program slicing. In: Zelkowitz M, editor. *Advances in Computers*, vol. 43. Amsterdam: Academic Press; 1996. p. 1–50.
- [6] Perugini S. *Program transformations for information personalization*. Ph.D. thesis, Department of Computer Science, Virginia Tech; 2004.
- [7] Perugini S, Anderson T, Moroney W. A study of out-of-turn interaction in menu-based, IVR, voicemail systems. In: Rosson M, Gilmore D, editors. *Proceedings of the 25th international ACM conference on human factors in computing systems (CHI)*. New York, NY: ACM Press; 2007. p. 961–70.
- [8] Baeza-Yates R, Ribeiro-Neto B. *Modern information retrieval*. New York: ACM Press; 1999.
- [9] Graunke P, Findler R, Krishnamurthi S, Felleisen M. Automatically restructuring programs for the web. In: *Proceedings of the 16th IEEE international conference on automated software engineering (ASE)*. Los Alamitos, CA: IEEE Computer Society; 2001. p. 211–22.
- [10] Queinsec C. The influence of browsers on evaluators or, continuations to program web servers. *ACM SIGPLAN Notices* 2000;35(9):23–33.
- [11] Perugini S. Symbolic links in the open directory project. *Information Processing & Management* 2008;44(2):910–30.
- [12] Randell B, Buxton J, editors. *Software engineering techniques: report of a conference sponsored by the NATO Science Committee*. Brussels, Scientific Affairs Division, NATO (1970), Rome, Italy; 1969. 164pp.
- [13] Harman M, Hierons R. An overview of program slicing. *Software Focus* 2001;2(3):85–92.
- [14] Tip F. A survey of program slicing techniques. *Journal of Programming Languages* 1995;3(3):121–89.

- [15] Weiser M. Program slices: formal, psychological, and practical investigations of an automatic program abstraction method. PhD dissertation, University of Michigan; 1979.
- [16] Weiser M. Programmers use slices when debugging. *Communications of the ACM* 1982;25(7):446–552.
- [17] Weiser M. Program slicing. *IEEE Transactions on Software Engineering* 1984;10(4):352–7.
- [18] Horwitz S, Repts T, Binkley D. Interprocedural slicing using dependency graphs. *ACM Transactions on Programming Languages and Systems* 1990;12(1):26–60.
- [19] Venkatesh V. The semantic approach to program slicing. *ACM SIGPLAN Notices* 1991;26(6):107–19.
- [20] Sacco G. Dynamic taxonomies: a model for large information bases. *IEEE Transactions on Knowledge and Data Engineering* 2000;12(3):468–79.
- [21] Lyle J, Weiser M. Automatic program bug location by program slicing. In: *Proceedings of the second international conference on computers and applications*. Los Alamitos, CA: IEEE Computer Society; 1987. p. 877–82.
- [22] Jackson D, Rollins E. Chopping: a generalisation of slicing. Technical Report CMU-CS-94-169, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA; 1994.
- [23] Jackson D, Rollins E. A new model of program dependencies for reverse engineering. *ACM SIGSOFT Software Engineering Notes* 1994;19(5):2–10.
- [24] Repts T, Rosay G. Precise interprocedural chopping. *ACM SIGSOFT Software Engineering Notes* 1995;20(4):41–52.
- [25] Anderson P, Repts T, Teitelbaum T. Design and implementation of a fine-grained software inspection tool. *IEEE Transactions on Software Engineering* 2003;29(8):721–33.
- [26] Jones N, Gomard C, Sestoft P. *Partial evaluation and automatic program generation*. Prentice Hall International; 1993.
- [27] Hildum D, Cohen J. A language for specifying program transformations. *IEEE Transactions on Software Engineering* 1990;16(6):630–8.
- [28] Jones N. An introduction to partial evaluation. *ACM Computing Surveys* 1996;28(3):480–503.
- [29] Jones N. *Computability and complexity from a programming perspective*. In: *Foundations of computing series*. Cambridge, MA: MIT Press; 1997.
- [30] Ricca F, Tonella P. Web application slicing. In: *Proceedings of the international conference on software maintenance (ICSM)*. Los Alamitos, CA: IEEE Computer Society; 2001. p. 148–57.
- [31] Allen J. Mixed-initiative interaction. *IEEE Intelligent Systems* 1999;14(5):14–6.
- [32] Hearst M. Mixed-initiative interaction. *IEEE Intelligent Systems* 1999;14(5):14.
- [33] Ramakrishnan N, Perugini S. The partial evaluation approach to information personalization. Technical Report cs.IR/0108003, Computing Research Repository (CoRR); 2001.
- [34] Dhyani D, Ng W, Bhowmick S. A survey of web metrics. *ACM Computing Surveys* 2002;34(4):469–503.
- [35] Taylor A, editor. *Waynar's introduction to cataloging and classification*, 9th ed. Englewood, CO: Libraries Unlimited, Inc.; 2000.
- [36] Perugini S. Supporting multiple paths to objects in information hierarchies: faceted classification, faceted search, and symbolic links. *Information Processing & Management* 2009. doi:10.1016/j.ipm.2009.06.007.
- [37] Abiteboul S, Buneman P, Suciu D. *Data on the web: from relations to semistructured data and XML*. San Francisco, CA: Morgan Kaufmann; 2000.
- [38] Eirinaki M, Vazirgiannis M. Web mining for web personalization. *ACM Transactions on Internet Technology* 2003;3(1):1–27.
- [39] Mobasher B, Cooley R, Srivastava J. Automatic personalization based on web usage mining. *Communications of the ACM* 2000;43(8):142–51.
- [40] Mulvenna M, Anand S, Büchner A. Personalization on the net using web mining. *Communications of the ACM* 2000;43(8):122–5.
- [41] Srivastava J, Cooley R, Deshpande M, Tan P-N. Web usage mining: discovery and applications of usage patterns from web data. *ACM SIGKDD Explorations* 2000;1(2):12–23.
- [42] Perugini S, Ramakrishnan N. Mining web functional dependencies for flexible information access. *Journal of the American Society for Information Science and Technology* 2007;58(12):1805–19.
- [43] Chen Z, Xu B, Yang H. Detecting dead statements for concurrent programs. In: *Proceedings of the international conference on software maintenance (SCAM)*. Los Alamitos, CA: IEEE Computer Society; 2001. p. 67–72.
- [44] Wegman W, Zadeck F. Constant propagation with conditional branches. *ACM Transactions on Programming Languages and Systems* 1991;13(2):181–210.
- [45] Bergeretti J-F, Carré B. Information-flow and data-flow analysis of while-programs. *ACM Transactions on Programming Languages and Systems* 1985;7(1):37–61.
- [46] Repts T, Turnidge T. Program specialization via program slicing. In: Danvy O, Glück R, Thiemann P, editors. *Proceedings of the Dagstuhl seminar on partial evaluation. Lecture notes in computer science*, vol. 1110. Wadern, Germany: Springer, Schloss Dagstuhl; 1996. p. 409–29.
- [47] Yankelovich N. How do users know what to say? *ACM Interactions* 1996;3(6):32–43.
- [48] Bodner R, Chignell M. Dynamic hypertext: querying and linking. *ACM Computing Surveys* 1999;31(4es):31(4es) Article No. 15.
- [49] Croft W, Thompson R. I³R: a new approach to the design of document retrieval systems. *Journal of the American Society for Information Science* 1987;38(6):389–404.
- [50] Marchionini G. *Information seeking in electronic environments*. In: *Cambridge series on human-computer interaction*. Cambridge, UK: Cambridge University Press; 1997.
- [51] Williams M. What makes RABBIT run?. *International Journal of Man-Machine Studies* 1984;21(4):333–52.
- [52] Perugini S, Ramakrishnan N. Interacting with web hierarchies. *IEEE IT Professional* 2006;8(4):19–28.
- [53] Perugini S. Real-time query expansion and procedural interfaces for information hierarchies. In: Broder A, Maarek Y, editors. *Proceedings of the international ACM SIGIR workshop on faceted search*; 2006.
- [54] Florescu D, Levy A, Mendelzon A. Database techniques for the world-wide web: a survey. *SIGMOD Record* 1998;27(3):59–74.
- [55] Fernández M, Florescu D, Levy A, Suciu D. A query language for a web-site management system. *SIGMOD Record* 1997;26(3):4–11.
- [56] Fernández M, Florescu D, Kang J, Levy A, Suciu D. Catching the boat with strudel: experiences with a web-site management system. *ACM SIGMOD Record* 1998;27(2):414–25.
- [57] Wilson R, Bergeron R. Dynamic hierarchy specification and visualization. In: *Proceedings of the IEEE symposium on information visualization (INFOVIS)*. Los Alamitos, CA: IEEE Computer Society; 1999. p. 65–72.
- [58] Hearst M. Next generation web search: setting our sites. *IEEE Data Engineering Bulletin* 2000;23(3):38–48.
- [59] Hearst M, Elliott A, English J, Sinha R, Swearingen K, Yee K-P. Finding the flow in web site search. *Communications of the ACM* 2002;45(9):42–9.
- [60] Brusilovsky P. Adaptive hypermedia. *User Modeling and User-Adapted Interaction* 2001;11(1–2):87–110.
- [61] Pokorny J. Static pages are dead: how a modular approach is changing interaction design. *ACM Interactions* 2001;8(5):19–24.
- [62] Plaisant C, Shneiderman B, Doan K, Bruns T. Interface and data architecture for query preview in networked information systems. *ACM Transactions on Information System* 1999;17(3):320–41.
- [63] Belkin N, Cool C, Stein A, Thiel U. Cases, scripts, and information-seeking strategies: on the design of interactive information retrieval systems. *Expert Systems with Applications* 1995;9(3):379–95.
- [64] Friedman D, Wand M, Haynes C. *Essentials of programming languages*, 2nd ed. Cambridge, MA: MIT Press; 2001.
- [65] Quan D, Huynh D, Karger D, Miller R. User interface continuations. In: Konstan J, editor. *Proceedings of the 16th annual ACM symposium on user interface software and technology (UIST)*. New York, NY: ACM Press; 2003. p. 145–8.
- [66] Ullman J. *Elements of ML programming*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall; 1997.
- [67] Bergman E, editor. *Information appliances and beyond*. In: *The Morgan Kaufmann series on interactive technologies*. San Francisco, CA, Morgan Kaufmann; 2000.
- [68] Perry M, O'Hara K, Sellen A, Brown B, Harper R. Dealing with mobility: understanding access anytime, anywhere. *ACM Transactions on Computer-Human Interaction* 2001;8(4):323–47.
- [69] Hightower J, Borriello G. Location systems for ubiquitous computing. *IEEE Computer* 2001;34(8):57–66.

- [70] Ricca F, Tonella P, Baxter I. Restructuring web applications via transformation rules. In: Proceedings of the first international workshop on source code analysis and manipulation (SCAM). Los Alamitos, CA: IEEE Computer Society; 2001. p. 150–60.
- [71] Ricca F, Tonella P. Understanding and restructuring web sites with ReWeb. *IEEE MultiMedia* 2001;8(2):40–51.
- [72] Ricca F, Tonella P, Baxter I. Web application transformations based on rewrite rules. *Information and Software Technology* 2002;44(13):811–25.
- [73] Spiliopoulou M. Web usage mining for web site evaluation. *Communications of the ACM* 2000;43(8):127–34.
- [74] Gonçalves M, Zafer A, Ramakrishnan N, Fox E. Modeling and building personalized digital libraries with PIPE and 5SL. In: Smeaton A, Callan J, editors. Proceedings of the joint DELOS-NSF workshop on personalisation and recommender systems in digital libraries. Dublin, Ireland; 2001.
- [75] Perugini S, McDevitt K, Richardson R, Pérez-Quiñones M, Shen R, Ramakrishnan, N, et al. Enhancing Usability in CITIDEL: multimodal, multilingual, and interactive visualization interfaces. In: Chen H, Wactlar H, Chen C, Lim E-P, Christel M, editors. Proceedings of the fourth international ACM/IEEE-CS joint conference on digital libraries (JCDL). New York, NY: ACM Press; 2004. p. 315–24.
- [76] Gonçalves M, Fox E, Watson L, Kipp N. Streams, structures, spaces, scenarios, societies (5S): a formal model for digital libraries. *ACM Transactions on Information Systems* 2004;22(2):270–312.
- [77] Perugini S, Ramakrishnan N, Fox E. Automatically generating interfaces for personalized interaction with digital libraries. Technical Report cs.DL/0402022, Computing Research Repository (CoRR); 2004.
- [78] Perugini S, Lakshminarayanan P, Ramakrishnan N. Personalizing the GAMS cross-index. Technical Report TR-00-01, Department of Computer Science, Virginia Tech; 2000.
- [79] Sacco G, Tzitzikas Y, editor. *Dynamic taxonomies and faceted search: theory, practice, and experience*. Berlin: Springer; 2009.
- [80] Dalamagas T, Bourros P, Galanis T, Eirinaki M, Sellis T. Mining user navigation patterns for personalizing topic directories. In: Fundulaki I, Polyzotis N, editors. Proceedings of the ninth annual ACM international workshop on web information and data management (WIDM). New York, NY: ACM Press; 2007. p. 81–8.
- [81] Haller S, McRoy S, editor. *Computational models for mixed initiative interaction: papers from the 1997 AAAI spring symposium*, vol. SS-97-04. Menlo Park, CA: AAAI Press; 1997.
- [82] Ferguson G, Allen J. TRIPS: an integrated intelligent problem-solving assistant. In: Mostow J, Rich C, Buchanan B, editors. Proceedings of the 15th national conference on artificial intelligence (AAAI). Menlo Park, CA: AAAI Press; 1998. p. 567–72.
- [83] Capra R, Narayan M, Perugini S, Ramakrishnan N, Pérez-Quiñones M. The staging transformation approach to mixing initiative. In: Tecuci G, editor. Working notes of the IJCAI 2003 workshop on mixed-initiative intelligent systems. Menlo Park, CA: AAAI/MIT Press; 2003. p. 23–9.
- [84] Ramakrishnan N, Capra R, Pérez-Quiñones M. Mixed-initiative interaction = mixed computation. *ACM SIGPLAN Notices* 2002;37(3):119–30.
- [85] McGlashan S, Burnett D, Danielsen P, Ferrans J, Hunt A, Karam G, et al. Voice extensible markup language: voiceXML. Technical Report, VoiceXML Forum, version 2.0; 2001.
- [86] Futamura Y, Nogi K, Takano A. Essence of generalized partial computation. *Theoretical Computer Science* 1991;90(1):61–79.
- [87] Takano A. Generalized partial computation for a lazy functional language. *ACM SIGPLAN Notices* 1991;26(9):1–11.
- [88] Consel C, Khoo S. Parameterized partial evaluation. *ACM Transactions on Programming Languages and Systems* 1993;15(3):463–93.
- [89] Maes P. Concepts and experiments in computational reflection. In: Conference proceedings on object-oriented programming systems, languages and applications (OOPSLA). New York, NY: ACM Press; 1987. p. 147–55.
- [90] Harman M, Gold N, Hierons R, Binkley D. Code extraction algorithms which unify slicing and concept assignment. In: Proceedings of the ninth IEEE working conference on reverse engineering (WCRE). Los Alamitos, CA: IEEE Computer Society; 2002. p. 11–21.
- [91] Ianov Y. The logical schemes of algorithms. In: *Problems of cybernetics*, vol. 1. New York, NY: Pergamon Press; 1960. p. 82–140.

Saverio Perugini is an Assistant Professor in the Department of Computer Science at the University of Dayton. His research interests include information personalization, web mining, and functional programming. He has a Ph.D. in Computer Science from Virginia Tech.

Naren Ramakrishnan is a Professor and the Associate Head for graduate studies in the Department of Computer Science at Virginia Tech. He received his Ph.D. in Computer Sciences from Purdue University. His research interests include mining scientific data, recommender systems, and information personalization.