

Adaptive Hypermedia Courses:

Qualitative and Quantitative Evaluation and Tool Support

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Technische Universiteit Eindhoven,
op gezag van de rector magnificus, prof.dr.ir. C.J. van Duijn,
voor een commissie aangewezen door het College voor Promoties
in het openbaar te verdedigen
op Vrijdag 25 April 2014 om 16.00 uur

door

Vinicius Faria Culmant Ramos

geboren te Porto Velho, Brazilië

Dit proefschrift is goedgekeurd door de promotoren:

prof.dr. P.M.E. De Bra

en

prof.dr. G.B. Xexéo

Copromotor:

dr. M. Pechenizkiy

Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool

Support by V.F.C. Ramos.

Eindhoven: Technische Universiteit Eindhoven, 2014

A catalogue record is available from the Eindhoven University of Technology Library.

ISBN: ???-??-???-????-?

Copyright © 2014 by V.F.C. Ramos. All rights reserved.



ADAPTIVE HYPERMEDIA COURSES: QUALITATIVE AND QUANTITATIVE EVALUATION AND TOOL SUPPORT

Vinicius Faria Culmant Ramos

Tese de Doutorado apresentada ao Programa de Pós-graduação em Engenharia de Sistemas e Computação, COPPE, da Universidade Federal do Rio de Janeiro, como parte dos requisitos necessários à obtenção do título de Doutor em Engenharia de Sistemas e Computação.

Orientadores: Geraldo Bonorino Xexéo
Paul Maria Emile De Bra

Rio de Janeiro
Novembro de 2013

ADAPTIVE HYPERMEDIA COURSES: QUALITATIVE AND
QUANTITATIVE EVALUATION AND TOOL SUPPORT

Vinicius Faria Culmant Ramos

TESE SUBMETIDA AO CORPO DOCENTE DO INSTITUTO ALBERTO LUIZ
COIMBRA DE PÓS-GRADUAÇÃO E PESQUISA DE ENGENHARIA (COPPE)
DA UNIVERSIDADE FEDERAL DO RIO DE JANEIRO COMO PARTE DOS
REQUISITOS NECESSÁRIOS PARA A OBTENÇÃO DO GRAU DE DOUTOR
EM CIÊNCIAS EM ENGENHARIA DE SISTEMAS E COMPUTAÇÃO.

Examinada por:

Prof. Geraldo Bonorino Xexéo, D.Sc.

Prof. Paul Maria Emile De Bra, Ph.D.

Prof. Daniel Schwabe, Ph.D.

Prof. Geraldo Zimbrão da Silva, D.Sc.

Profa. Miriam Struchiner, Ph.D.

Dr. Mykola Pechenizkiy, Ph.D.

RIO DE JANEIRO, RJ - BRASIL

NOVEMBRO DE 2013

Ramos, Vinicius Faria Culmant

Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support/ Vinicius Faria Culmant Ramos. — Rio de Janeiro: UFRJ/COPPE, 2013.

xviii, 135 p.: il.; 29,7 cm

Orientadores: Geraldo Bonorino Xexéo

Paul Maria Emile De Bra

Tese (doutorado) — UFRJ/ COPPE/ Programa de Engenharia de Sistemas e Computação, 2013.

Referências Bibliográficas: p. 127-135.

1. Adaptive Hypermedia. 2. Evaluation. 3. Adaptive Course I. Xexéo, Geraldo Bonorino, *et al.* II. Universidade Federal do Rio de Janeiro, COPPE, Programa de Engenharia de Sistemas e Computação. III. Título.

To my wife Marina and
my wonderful children
Júlia and Maurício,
you make my life easier
and happier.

Acknowledgement

First of all, I would like to thank my parents. This work would not have been possible without their support and assistance during my whole life. Thanks Dad (Sr. Claudio) and Mom (D. Sonia). Both of you deserve all my love and gratitude.

I would also like to thank my brother and sisters: Carlos, Cássia and Isabela. Thank you for your care, laughs and support in the moments I needed, even though sometimes I did not deserve that much. Many thanks to my sister Cássia and her beautiful family: Felipe (husband), Nathália (daughter) and Maria Eduarda (daughter). It is an easy task to talk about my youngest sister (Isabela), since the youngest is always the one loved by everyone, and she is certainly no exception. My brother, his wife (Ana) and their amazing happy daughter (Clara) are more than special to me. I want to thank you so much for all the funny and lovely moments we spent together.

My happiness and my life would not have been complete without these three wonderful people: Marina (wife), Júlia (daughter) and Maurício (son). They are the most special people in the world to me. I do not have words to thank you all. Júlia, you are 5 years old now and I spent most of the happiest moments of my life at your side. Even if you do not know yet, you have given me the incentives to face my challenges and to keep going over the problems every day. Maurício, you are less than a year old, and you, like your sister, also made my life happier. You are a beautiful baby and your constant smile makes my life easier. It is not easy to thank my wonderful wife, Marina, using only words, since she is the person that supports me in all my decisions, and shares all the easy and difficult moments of life at my side. Marina, you are amazing! Thank you for all your unending support, and for all the nights I had to stay awake to write the thesis or to write research papers. Thank you for helping me during such stressful moments. Thank you for sharing the greatest moments in my life. THANK YOU, Marina!

Thank you Vera (mother-in-law), Pedro (brother-in-law) and Juliana (wife's cousin) for the support and funny moments in the last three years we stayed closer.

Thank you Alexandre Stauffer. You have been a great friend for years. You are not only a good friend, but you are also the person who gave me support for continuing to study, making research and also developing all the academic skills I

have today. You read each part of my thesis and papers, not because you had to, but because you are an amazing friend (and researcher). Your criticism improved all my academic qualities, and, consequently, my research. In my place you have free beer for the rest of your life, you only have to ask for them. Thank you very very very much.

There are a lot of friends I would like to thank. First of all, I would like to thank my friends from my undergraduate at UFRJ: Calisto, Izaías, Paulo Nunes, Raphael (Dill), Taísa and Targino. I would like to thank the secretaries of PESC/COPPE: Ana Paula, Claudia, Guty, Mercedes, Patrícia and Solange. Thank you, Mrs. Riet and Mrs. Ine, you were amazingly nice and kind during my stay at TU/e. Many thanks to my friends from Eindhoven: Antonino Simone, Alberto Perrota, Camille Carcoute, Evegny Knutov, Gierry, Ivelina, Jorge, Marcos, Mayla Brusio and Nicolas Monti.

I am very thankful for Dr. David Smits for all your help. Since I have started my research with the AHA! system (and later the GALE system), Dr. David Smits gave me a lot of support in the development of new code to the AHA! and GALE system.

Thank you Prof. Miriam Struchiner for being in my thesis committee at UFRJ. Prof. Miriam has played a major role in my academic life. She was responsible for my first steps in research, accepting me as an undergraduate student into her lab in 2002. Prof. Miriam deserves all my gratitude and admiration.

Thank you Prof. Dr. Marcus Specht and Prof. Dr. Wim Jochems for accepting to be in my thesis committee at TU/e, and thank you Prof. Daniel Schwabe and Prof. Geraldo Zimbrão for accepting to be in my thesis committee at UFRJ and TU/e.

I am thankful to Dr. Mykola Pechenizkiy to be in my thesis committee at UFRJ and TU/e, and also for giving me fundamental and important feedback on my research. He was also very kind and nice during my stay at TU/e. Thank you very much.

At last, but not less important, I am thankful to my supervisors: Prof. Dr. Paul De Bra and Prof. Geraldo B. Xexéo.

Prof. Xexéo gave me the opportunity to work with him in his research group and was the first one to encourage my mobility to TU/e in 2009. His friendship

and kindness brought us into a closer contact than that between a supervisor and a student. I know I can count on him.

I have all my gratitude to Prof. Dr. De Bra. You are one of the greatest researchers I have ever met. You are more than a researcher, you are a nice and generous person. Your academic and personal support before, during and after my stay in Eindhoven were invaluable to me, and have made me have a great admiration for you as a researcher and a person. I have learned a lot with you, I am thank you for every single word, method, technique, academic stuffs etc you taught me. Thank you very very very much!

Many thanks to CNPq for the financial support during most part of this thesis.

Resumo da Tese apresentada à COPPE/UFRJ como parte dos requisitos necessários para a obtenção do grau de Doutor em Ciências (D.Sc.)

ADAPTIVE HYPERMEDIA COURSES: QUALITATIVE AND QUANTITATIVE EVALUATION AND TOOL SUPPORT

Vinicius Faria Culmant Ramos

Novembro/2013

Orientadores: Geraldo Bonorino Xexéo
Paul Maria Emile De Bra

Programa: Engenharia de Sistemas e Computação

O foco deste trabalho é a avaliação de cursos adaptativos criados e entregues pelo AHA! (Adaptive Hypermedia Architecture) and GALE (Generic Adaptation Language and Engine desenvolvido dentro do projeto EU FP7 GRAPPLE). O objetivo destas avaliações é entender a influência da adaptação no aprendizado dos alunos de um curso adaptativo. Os métodos avaliativos são divididos em qualitativo e quantitativo. Os métodos quantitativos consistem na análise dos *logs*, testes e notas de avaliações dos alunos. As análises de questionários fazem parte dos métodos qualitativos. Também fizemos neste trabalho a avaliação da modularidade e extensibilidade do GALE como parte da preocupação em ter um sistema como esse com apenas um núcleo de adaptação que pode ser estendido para ser usado em diferentes tipos de sistemas adaptativos. Esta tese também apresenta ferramentas de análise de logs de navegação e de resultados de testes e questionários dos cursos adaptativos no GALE. O principal objetivo destas ferramentas é auxiliar os autores apresentando-os medidas estatísticas de seu próprio curso, permitindo a eles uma análise da estrutura do curso do ponto de vista da navegação dos alunos. No final, discutimos os trabalhos futuros, em especial as sugestões de mudanças na configuração do GALE (para desenvolvedores que precisam estender o sistema) e na estrutura de cursos baseado nas observações dos comportamentos dos alunos e nas sugestões apresentadas por eles.

Abstract of Thesis presented to COPPE/UFRJ as a partial fulfillment of the requirements for the degree of Doctor of Science (D.Sc.)

ADAPTIVE HYPERMEDIA COURSES: QUALITATIVE AND QUANTITATIVE EVALUATION AND TOOL SUPPORT

Vinicius Faria Culmant Ramos

November/2013

Supervisors: Geraldo Bonorino Xexéo
Paul Maria Emile De Bra

Department: Systems Engineering and Computer Science

The focus of this work is the evaluation of adaptive courses created and delivery by AHA! (the Adaptive Hypermedia Architecture) and GALE (the Generic Adaptation Language and Engine developed in the EU FP7 project GRAPPLE). The main goal of these evaluations is to understand the influence of adaptation on students' learning in an adaptive hypertext course. The evaluation methods are divided into qualitative and quantitative ones. The quantitative methods consist of analysis of the students' logs, the performed tests and assignment grades. The analysis of questionnaires is part of the qualitative method. In this work we also performed an evaluation of the modularity and extensibility of the GALE system as part of the concern in having such a system with a single core adaptation engine that can be extended in order to use it for different types of adaptation. This thesis also presents tools for the analysis of navigation logs and test and quiz results of adaptive courses in GALE. The main goal of these tools is to assist the courses' authors to retrieve statistical measurements for their own courses, allowing them to analyze the structure of the course from the point of view of the students' navigation. At the end we discuss future work, and in particular suggest changes to the setup of GALE (for developers that needs to extend the system) and to the structure of hypertext courses based on the observed student behavior as well as the student feedback.

Samenvatting van het proefschrift ingediend bij COPPE/UFRJ en de TU/e ter verkrijging van de graad van Doctor.

ADAPTIEVE HYPERMEDIA LEERMATERIAAL: KWALITATIEVE EN
KWANTITATIEVE EVALUATIE EN ONDERSTEUNENDE HULPMIDDELEN

Vinicius Faria Culmant Ramos

November/2013

Promotores: Geraldo Bonorino Xexéo en Paul Maria Emile De Bra

Department: Systems Engineering and Computer Science

De focus van dit werk is de evaluatie van adaptief leermateriaal dat gemaakt is voor en aangeboden door het AHA! system (Adaptive Hypermedia Architecture) en door GALE (Generic Adaptation Language and Engine, ontwikkeld in het EU FP7 project GRAPPLE). Het hoofddoel van deze evaluaties is om de invloed te begrijpen die het gebruik van adaptief leermateriaal heeft op het leergedrag van studenten. De evaluatie bestaat uit kwalitatieve en kwantitatieve methoden. De kwantitatieve methoden bestaan uit een analyse van de logbestanden met betrekking tot de navigatie door studenten, de uitgevoerde toetsen en de eindbeoordeling van opdrachten. De kwalitatieve methode bestaat uit de analyse van vragenlijsten. In dit werk hebben we ook een evaluatie uitgevoerd van de modulariteit en uitbreidbaarheid van het GALE systeem, bedoeld om met een uitbreidbare kern verschillende types adaptatie te ondersteunen. Dit proefschrift presenteert ook hulpmiddelen voor de analyse van navigatie-logbestanden en toets en quiz resultaten behaald in GALE leermateriaal. Het hoofddoel van deze hulpmiddelen is om de auteurs van leermateriaal te helpen om statistische metingen te verkrijgen voor hun leermateriaal en om hen toe te laten de structuur van het leermateriaal te analyseren met het oogpunt van navigatie door studenten. Aan het eind van het proefschrift bespreken we toekomstig werk, waarbij we in het bijzonder uitbreidingen voorstellen in de opstelling van GALE (voor ontwikkelaars die het systeem moeten uitbreiden) en in de structuur van hypertext leermateriaal, gebaseerd op observatie van het gedrag van studenten en op terugkoppeling door studenten.

Contents

1	Introduction	1
1.1	Motivation	2
1.2	Research Questions and Hypothesis	4
1.3	Research Goal and Approach	5
1.4	Important Definitions	6
1.5	Thesis Outline	8
1.6	Suggestion of Reading Paths	9
2	Evaluation of Adaptive Hypermedia Systems	10
2.1	Adaptive Systems Comparison	11
2.1.1	Domain Model Comparison	11
2.1.2	User Model Comparison	12
2.1.3	Adaptation Model and Adaptation Engine Comparison	13
2.1.4	Other “models” in Adaptive Systems	14
2.2	Evaluation of Adaptive Systems	15
2.2.1	Comparative Evaluation	16
2.2.2	Layered Evaluation	18
2.2.3	Empirical Evaluation	19
2.2.4	Evaluation Frameworks	20
3	Preliminary Evaluation of an Adaptive Course	23
3.1	Description of the Adaptive Course	23
3.2	Dataset	27
3.3	A Quantitative Evaluation Approach: Empirical Evidence	28
3.4	HUBS and Informative Pages	33
3.4.1	The Empirical Hub Coefficient	34
3.4.2	Informative Pages	36
3.5	Discussion and Conclusion	38
4	Second Evaluation of an Adaptive Course	39
4.1	Description of the Adaptive Course	40
4.2	Dataset	45

4.3	Qualitative and Quantitative Evaluation Approach: Empirical Evidence and Questionnaires	46
4.4	HUBS and Informative Pages	53
4.4.1	The Empirical Hub Coefficient	53
4.4.2	Informative Pages	55
4.5	Discussion and Conclusion	56
5	Continuous-Time Layered Evaluation in an Adaptive Hypermedia System	60
5.1	Introduction	60
5.2	Methodology	62
6	GALE: Open Generic Extensible Adaptive System	68
6.1	The Architecture of GALE	69
6.1.1	The GALE Processor Pipeline	72
6.1.2	GALE Flexibility through Configuration	77
6.1.3	Extending GALE: from generic towards general-purpose adaptation engine	83
6.2	Examples of GALE Applications	85
7	GALE Extensibility Evaluation: A Qualitative Approach	95
7.1	Introduction	95
7.2	Dataset	96
7.3	A Qualitative Evaluation Approach: Authoring in and creating extensions for GALE	98
7.4	Discussion and Conclusions	104
8	New Tools for Analyzing Navigation, Test and Quiz Logs	106
8.1	Technical Support	107
8.2	Statistical Measures	112
9	Conclusion and Future Work	119
9.1	Contributions on Adaptive Course Evaluation	119
9.1.1	Hubs and Informative Pages	120
9.2	Contributions on GALE Extensibility Evaluation	121

9.3	Limitations	122
9.4	Future Work	123

List of Figures

1	Screen-shot of the <i>welcome</i> page from the hypertext course, showing also <i>bad</i> , <i>good</i> and <i>neutral</i> links in the main view.	25
2	Total number of accesses performed by the students to each main concept of the course. The colors represent the presentation class of the link used to access the concept.	31
3	Number of students that accessed a concept via bad links. We show the five concepts with the largest number of student accesses via bad links.	32
4	Concepts whose links were most frequently used to access another concept.	34
5	Screenshot of a page from the hypertext course, showing also the navigation menu (left) and a header that gives access to settings and progress information.	43
6	Screenshot of the <i>welcome</i> page from the hypertext course, showing also <i>bad</i> and <i>neutral</i> links in the menu (left) and in the main view (right).	52
7	Adaptation decomposed (presented in [1])	61
8	Navigation Path sequences from students who fail in the test S. (1) represents the students' paths. (2) represents the common pattern extracted from (1), where * stands for visits to zero or more concepts.	64
9	Navigation Paths sequence from students that succeed in the test S. (3) represents the students' paths. (4) represents the common pattern extracted from (3), where * stands for visits to zero or more concepts.	65
10	The architecture of GALE [2]	70
11	Handling a request for a concept [2]	71
12	User model updates being handled by the adaptation engine and UM service(s) [3]	75
13	Creating pages separately [2]	86
14	Authoring through template pages [2]	86
15	Example of a page based on a template.	87
16	Link adaptation based on prerequisites.	90

17	Screenshot of the adaptive PhD thesis about and served by GALE. . .	91
18	Screenshot of a concept with links to the analysis tools within the menu in the top right position.	111
19	Screenshot of the menu options for the Log Analysis Tool.	112
20	Visits per Concepts: the number of visits for a concept and links to drill-down the measure per <i>Incoming</i> link or per <i>Outgoing</i> link. . . .	113
21	A screenshot of the test analysis tool. The test called <i>test-history</i> is shown as an example of the tool.	118

List of Tables

1	Number of pages with EHC according to the row and out-degree represented by the column.	35
2	First access information for some concepts. The table shows the out-degree, the number of students who accessed each concept, and three types of first access: <i>first access</i> , which consists of the very first access of a student to the concept provided she followed a link or pressed the back button of the browser after visiting the concept, <i>first click</i> , which is the first time a student followed a link from the concept, and <i>first access with click</i> , which is the intersection of the previous two. For each type of access, the table shows the number of students and the average time in seconds they spent on the page.	36
3	Number of time a student answered the final test, his highest score and the exam's grade	47
4	Summary of the answers about why students follow bad links.	49
5	Number of concepts with EHC according to the row and out-degree represented by the column.	55
6	First access information for some concepts. The table shows the out-degree, the number of students who accessed each concept, and three types of first access: <i>first access</i> , which consists of the very first access of a student to the concept provided she followed a link or pressed the back button of the browser after visiting the concept, <i>first click</i> , which is the first time a student followed a link from the concept, and <i>first access with click</i> , which is the intersection of the previous two. For each type of access, the table shows the number of students and the average time in seconds they spent on the page.	57
7	Adaptive applications developed by (groups of students)	94
8	Number of answers about the aspects of authoring regarding the ease of creation in a level from 1 (very easy) to 5 (very difficult).	99
9	Number of answers about the aspects of development regarding the ease of implementation in a level from 1 (very easy) to 5 (very difficult).102	

1 Introduction

In the e-learning field, Bates [4] says that today the researchers do not have to discuss whether the technology enhances learning; instead the discussion must be about the way the technology is used for learning. Besides, Bates discusses the use of learning environments through the Web.

Traditional e-learning systems were created to make learning material available and reachable, and to present tasks, activities and evaluations, but they did not allow for user interaction. Furthermore, these systems present their content (learning material, activities and evaluation) always in the same way not taking into account the skills and the characteristics of the users. For example, in a traditional system if we present a list with the main topics of a course, all students will see the same list, and in the same way (topics in the same order, links in the same color and style etc). It is often the case that different students in a course have different learning styles, goals, interests etc. Hence, the development of personalized, extensible and generic systems, which can be adapted to students' diversity, is a very powerful tool and a very big challenge.

E-Learning systems were initially developed to assist users to reach their goals and make these systems more attractive. In this way, many educational systems were developed using artificial intelligence techniques: for example, the tutor systems [5, 6] and adaptive systems [7, 8, 9]. An example of Intelligent Tutoring Systems (ITS) is presented by Brusilovsky in [6]. In this paper, the author refers to an on-line course that presents a button to the students, which we call *suggestion button*, that links the actual content to the most relevant topic (system's suggestion) in the course. This *suggestion button* intended to guide students that find themselves lost in the system, without knowing which content they should read in the course.

In the beginning, the main goal of these systems was to assist the users to solve proposed problems and tests. The authors assumed that the users learned the content of the course outside the on-line environment, by means of books and classes. In this way, the ITS developers (and authors) decided to combine the ITS assistance with learning materials, using hypertext and hypermedia environment. The union of the ITS and the hypermedia offered much more functionality to the system and to the users than the static educational environment. There were a

few system, which we now classify as adaptive hypermedia (AH) systems, whose development was influenced by the ITS [8, 10, 11].

1.1 Motivation

The concept of AH goes back a long way, all the way to Vannevar Bush' article "As We May Think" [12] if we interpret the text in a sufficiently creative way. The "Memex" device envisioned by Bush was a form of hypertext or even hypermedia because it was not limited to text. The user could link documents, which for the device would imply that retrieving one document would automatically also retrieve and display the other. The user could build "trails of his interest through the maze of materials available to him", which is a form of personalization (a personalized guided tour in fact). This personalization was aimed at coping with information overload, an all too common problem of Internet users today. Furthermore, the personalization also included adding some kind of annotation: "he inserts a page of longhand analysis of his own". This shows that just linking information items (possibly from different authors) may not constitute a coherent story, which reveals the importance of adding the annotations or what we would call content adaptation.

The personalization envisioned by Bush was aimed at *revisiting* information (finding it again through trails), and at *recalling* a previously discovered meaning. When Bush defined *trailblazing* as a possible new profession we can understand this as preparing personalization for others.

AH systems, summarized by Brusilovsky in 1996 [13] and 2001 [14], and further updated and detailed by Knutov et al. in 2009 [15], are systems that build, for each user, an individual model and apply it to adapt the application (e.g. an on-line course) to that user. The individual model is constructed based on the preferences, goals and evolving knowledge of the users. The AH aims at automating the Bush's "trailblazing" through *link adaptation* and the annotations through *content adaptation*. The AHAM reference model [16, 17] captured these methods and techniques into an abstract adaptive hypermedia architecture.

The AHA! system [18] was developed in parallel with the AHAM reference model [16, 17, 19]. Actually, AHA! was originally designed to serve a hypertext course taught through the Web [9]. The AHA! system was under development for

almost the last 20 years. The main goal of the AHA! developers was to make it as generic as possible, and to make it able to serve different adaptive methods and techniques and, consequently, different adaptive applications.

A lot has happened since AHAM and AHA! and the other models and systems of the early 21st century. Knutov et al. [15] describe many new adaptation techniques developed to provide a list of challenges for creating a new generic adaptive hypermedia system. That research has led to the GAF model (Generic Adaptation Framework) [20, 21], which has been designed to be capable of dealing with several types of adaptive systems and applications, from “traditional” adaptive hypermedia to personalized search and recommender systems. Designing a new adaptive system that encompasses all new techniques, from *open corpus adaptation*, through *domain models based on ontologies*, *group adaptation* (and *group formation*), *higher order adaptation* based on *web-log mining*, *context-awareness*, *information retrieval* and *recommender systems*, is an impossible task. GALE tackles this challenge through a very modular and extensible approach [2, 22, 3], consisting of an extensible generic and general purpose adaptive hypermedia engine. GALE was inspired and developed over the code of the AHA! system with a significant architectural difference: the separation of the *adaptive engine* from the *user modeling service* [23]. For this reason, we say that GALE is an evolution of the AHA! system.

It is very important to develop generic and general purpose systems, but it is also important to develop tools for the authors that create applications in such systems to evaluate them. For example, in an adaptive course it is important to the author to discover whether the course is navigated as he expected, and whether the students are learning from the proposed test.

The first works about evaluation of AH systems compared adaptive systems with non-adaptive versions of the same systems [24, 8, 25]. In general, the main goal of this kind of evaluation is to compare the efficiency between the two versions or, in some cases, to analyze the success and failure rates of users doing a task. A few years before these works, Totterdell and Boyle [26] already presented potential problems in this kind of evaluation, such as: to determine which system state the adaptive version must be chosen to compare with the non-adaptive version, and to specify in which point the comparison must start, since there is a gap in the adaptive version from the starting point of the user’s navigation and the apprehension of his needs

and characteristic. The authors also describe an evaluation of the AH systems that take into account the layers of these systems, originally, called as layered evaluation.

The adaptation of an AH system refers to the method or technique used to adapt the system or the content to the user. The most common ones are techniques used to adapt the presentation and/or the navigation of the system to the user based on the *user model*. A good taxonomy about these techniques can be found in [15, 21].

The layers of an AH system are the core of the system. They are responsible for each step of the adaptation. In fact, layers of an AH system are implicit logical divisions of the system, where each of them are responsible for different tasks. For example, in GALE there is an *User Model* (UM) and an *Adaptation Model* (AM) layer. The first one is responsible for the creation and update of the user knowledge, interests, preferences, goals and objectives, style, and other relevant properties that might be useful for adaptation. The AM is responsible to adapt the presentation, information content, and navigation structure to the user's knowledge, interests, goals and objectives, etc. In 2000 the layered evaluation became popular, even though Totterdell and Boyle had described this type of evaluation already in the 90s [26]. Brusilovsky et al. [27] propagated the technique of layered evaluation in their revision of their evaluation of an AH system [28] where they presented the benefits of this technique.

1.2 Research Questions and Hypothesis

This work provides answers to the following research questions:

1. Does the annotation/hiding adaptive technique influence the choices of the students? Does the annotation/hiding adaptive technique consists of an effective model of adaptation and students' guidance?

The hypothesis we have is that the adaptive technique is an effective model of adaptation and students' guide. To answer these questions we analyzed two versions of an adaptive course that we call *Hypermedia course*. The first analysis was made 3 years after the end of the course, and we used a quantitative approach to measure the influence of the annotation/hiding technique in the students' choice. We used new functions and extensions, as well as the results of this analysis, to carry out a second evaluation of the

Hypermedia course.

2. How does the interplay between the link structure of an adaptive course and the rules employed by the adaptation influence the choices of the students? The hypothesis we have is that the link structure influence in the choices of the students. Here we analyze the connection between the annotation/hiding adaptive technique and the link structure of the course, and their effect on the way the students follow the course.

In our analysis we focus on the concept of *hubs* and *authorities* defined by Kleinberg [29]. Intuitively, *hubs* are pages that contain a large number of links to others pages. We cannot apply this concept directly in our setting, since in an adaptive application the number of links of a web page may change over time.

3. Does GALE reach its goal as a generic and extensible adaptive system? Can GALE system effectively serve different kinds of adaptive applications?

Our hypothesis is that GALE reach its goal as a generic and extensible adaptive system serving different kinds of adaptive applications. In this study, master students were asked to develop new applications and extensions to GALE as part of their master coursework, and we put a questionnaire on-line for few weeks for the master students to answer regarding their experience. The answers were anonymous to give the students freedom to evaluate the system.

1.3 Research Goal and Approach

This thesis focuses on the evaluation of adaptive hypermedia courses created and developed with AHA! and GALE. We are also interested in the evaluation of the GALE system as a modular and extensible adaptive system. An important remark is that this research is focused on the evaluation of the AHA! and GALE systems and their applications; thus our goals and inferences are based in these systems. Our evaluations are intended to provide a better experience to students, authors and developers of adaptive applications in AHA! and GALE. Our research is divided in the following phases:

- Study and analyze the research methods and techniques used in empirical evaluations of adaptive systems, web-based information systems, intelligent tutoring systems, and other related fields, in order to apply (or create new ones and apply) them in the evaluation of adaptive courses.
- Evaluate an adaptive course created and delivered by the AHA! system using a selection of empirical and quantitative methods and techniques studied in the previous phase, present the results, problems and pitfalls found in the evaluation. The adaptive course was offered in 2006 and our analysis was carried out 3 years later. The data analyzed contains navigation and test logs.
- Evaluate an updated version of the adaptive course delivered by the GALE system with empirical, quantitative and qualitative methods and techniques. The methods and techniques are revisited and adapted from our first evaluation in order to eliminate the problems and pitfalls observed there. The empirical and quantitative evaluation was made using the navigation, test and quiz logs, while a questionnaire was used as a qualitative technique.
- Propose an evaluation framework to be implemented in GALE.
- Evaluate the modularity and extensibility of GALE, in order to validate the following goal of GALE: to be a generic and extensible adaptive system, which means that it can be used by researchers and educators without all being forced to use the same type of presentation and adaptation.

1.4 Important Definitions

In this section we introduce a few crucial definitions related to the AHA! and GALE systems, and that we use throughout the thesis. These definitions are related to the main ideas behind the adaptive techniques used in these systems.

- **Concept** — is an abstract representation of an information item from the application domain, e.g. subjects to be studied in a course, or artists, art styles, or art pieces (like paintings) in a museum. It has a unique identifier and any arbitrary Java data structure, where part of this code identify a resource file to be retrieved, adapted and presented to the user as a HTML

page. Some attributes have a meaning for the system, like *access* (a Boolean attribute that temporarily becomes true when a page is accessed), some have meaning for the author (and user), like *knowledge* or *interest*, and some have meaning for both, like *visited* (determining the link color) [30]. For example, the concept called *stratum* is composed by attributes, such as, *title*, *parent*, *suitability* and *availability*, and it refers to the file *stratum.xml*. The *suitability* and *availability* are Java variables. These variables are evaluated/manipulated by the system in the moment the *stratum* concept is requested. After this evaluation/manipulation, the system retrieve, adapt and present the *stratum* concept as an HTML page.

- **Page** — is the result (or representation) of a request for a concept in the system. It is typically an HTML page.
- **Link Hiding** — is an adaptive technique that can be found in Brusilovsky’s taxonomy [14]. The presentation of the links in a page is defined by an author-defined requirement. The requirements can express the common prerequisite relationships between concepts but can be used for any other condition. When a page is generated, links marked as conditional are displayed differently depending on the suitability of the link destination. Typically the system uses three link classes, named *good*, *neutral* and *bad*. The use of the name *bad* is only related to the link classes, it does not imply that it is a bad choice following these links. The link is shown in a color that depends on its class; the standard implementation of these colors are blue for good links, purple for neutral links and black for bad links. This realizes the link hiding technique because black color (associated to bad links) is the same color used for the text web page, making the link visually indistinguishable from the text.
- **Link Annotation** — is also an adaptive technique that can be found in Brusilovsky’s taxonomy [14]. It differs from link hiding because it is not “hidden” in the text. For example, a list of links in a page or in a menu is link annotation. The use of three colors for link classes that are all different from the color of the text of the web page are also considered part of the link annotation technique. Another strategy of link annotation is to associate icons to each class.

1.5 Thesis Outline

This thesis fits in the context of the evaluation of AH systems. Before we present an overview of the chapters presented in this work, we would like to remark that since this thesis is not an adaptive application, it is not possible to have adaptive techniques implemented in regular paper. Therefore, in Section 1.6 we would like to suggest you different reading paths depending on your background knowledge and interests. In Chapter 2 we briefly present the literature of the AH systems and their evaluations. After that, in Chapter 3, we present the evaluation of an adaptive course (to simplify we call it as the *Hypermedia* course) created for and served by AHA!. The Hypermedia course was offered in 2006 and our analysis was made 3 years later. The main goal of the evaluation is to study the combined effect of link structure and annotation/hiding on the navigation patterns of users. In Chapter 4 we present an evaluation of an updated version of the Hypermedia course, but at this time, it was served by GALE in 2012. There are differences between the two courses, an important one in the sense of the layout and navigation is the optional menu view in the updated version. This difference combined with a deadline of an exam significantly changed the findings versus what we found in the earlier analysis. The main goal of this evaluation is to identify which pages and links influence the choices of the students and to contrast this with the test logs, questionnaire answers and exam grades as well as the adaptation rules employed by the adaption engine. After this evaluation, the main goal of the Chapter 5 is to present a continuous evaluation with the objective to reveal possible problems in the adaptive courses or in the way students are using the course while the course design can still be modified. In Chapter 6 we present the GALE system and its modularity and extensibility. In this way, we describe the core functionality of GALE and its built-in elements. We also show in Chapter 6 a few extensions that have already been developed and used to make GALE more suitable for different types of applications. In Chapter 7 we present our findings about the evaluation of the GALE modularity and extensibility. After that, we present in Chapter 8 the tools developed for the authors of an adaptive application served by GALE. The main goal of these tools is to summarize the user's navigation, tests and quiz logs to present it to the authors. With this summary, the

authors can analyze the logs and draw their own conclusions. In the last chapter, we present our conclusions and future work.

1.6 Suggestion of Reading Paths

Since this thesis is not an adaptive application, we would like to suggest different reading paths depending on your interests and background. Regardless of your interests and background, we suggest the reader to start from Chapter 1 and to end with Chapter 9. For the remaining chapters, we suggest the following reading paths.

Newbies in AH systems. All Chapters. For people with little experience in the field of AH systems and/or evaluation of AH systems, we suggest to read this thesis in the order it is written. It is important to remark that Chapters 6 and 8 are more technical. Thus, a reader with not more than a minor interest in computer science, may skip these chapters.

Evaluation. Chapters 3, 4, 5, 7, and 8. For people interested in the evaluation of adaptive courses, we suggest to read Chapters 3, 4, 8 and 5, in this order. Afterwards, you may also want to read Chapter 7 if you are interested in the evaluation of the GALE system as a generic and extensible adaptive system.

Technical. Chapter 6. It contains technical details about GALE system and its extensions. This chapter can be read at any moment, since it is not a *prerequisite* of any other chapter. A possibility we find particularly good is to read this chapter right before reading Chapter 4 or Chapter 7.

2 Evaluation of Adaptive Hypermedia Systems

The main characteristic of an adaptive system is the ability of adapting itself to the needs of the users. For information services to comfortably replace human counterparts such as museum guides [31, 32, 33], teachers [10, 34] or personal tutors and guidance [10, 11, 35], the services need to take the characteristics of individual users and of user groups into account to decide what to present, how to present it and how to structure or order the presentation. For authors to easily (or at least easier than starting from scratch) create adaptive applications, like adaptive courses, it is important to have tools and frameworks to help them [8, 36, 18, 2].

The focus of this work is on the evaluation of an adaptive hypermedia course. Thus, we are specially interested in Adaptive Hypermedia (AH) systems. An AH system is also called adaptive web-based information systems (AHS for short), which means that it is developed as hypertext or hypermedia systems.

An adaptive system consists of many layers, and each of them is responsible for a specific step in the adaptation process. The three main parts of these systems are: the *domain model*, *user model* and *adaptation model*. In Knutov's thesis [21] an elaborate comparison is made between existing AHS and applications with respect to these three aspects. In Section 2.1 we briefly summarize the findings from that comparative study, and then also list other "parts" found in different types of adaptive systems.

The developers of adaptive systems are concerned with the effectivity, acceptability, efficiency and usability of the systems. Indeed, these points are a concern in all software products where the main goal is to assist users in their tasks. However an adaptive system has to go beyond that, since the system is created to understand the needs, the characteristics and the goals of the users. Considering that the adaptation takes into account the user's actions and tasks in the system, the adaptive system and application need to be evaluated constantly to confirm that creating and using the adaptation is efficient and effective for both end-users and authors. In Section 2.2 we present the different types of evaluation of an adaptive system and some examples.

2.1 Adaptive Systems Comparison

2.1.1 Domain Model Comparison

Each adaptive application must be based on a Domain Model (DM), describing how the conceptual representation of the application domain is structured. It usually consists of *concepts* and *concept relationships*. A concept represents an abstract information item from the application domain. In all systems we investigated the concepts form a hierarchy. The leaves of the hierarchy are atomic or primitive concepts and all other nodes are composite concepts that have sub-concepts. For example in Interbook [8] a textbook is structured as a hierarchy of chapters and sections with atomic presentations, tests or examples. The pages (and sections) are connected to a structure of concepts, indicating for each page what the required (prerequisite) knowledge is for the page and which outcome concepts the page teaches something about. KBS Hyperbook [37] uses a knowledge base that consists of so-called “Knowledge Items” which are essentially concepts. Each document from the document space is indexed by some concepts from the knowledge base that describes the content representation and hierarchical structure. In APeLS [36] the concepts are encapsulated into a “Narrative” structure where each narrative can be hierarchically split into sub-narratives.

Each system proposes its own way to encapsulate content information: for example like a Pagelet (in APeLS), which contains content and a content model, or it may be an Information Unit just encapsulating content information as in KBS-Hyperbook. These Information Units are indexed to map the Knowledge Items structure. In the AHAM model [16] and in the AHA! system [18] content representation is based on pages which consist of fragments that can be conditionally included by the AH system and which represent the lowest level in the concept hierarchy. (The AHAM model considers fragments to be static but in AHA! the content of fragments can be adaptive.)

A *concept relationship* is a meaningful connection between concepts. In AHAM it is represented as an object with a unique identifier, attribute-value pairs, presentation specification and a sequence of (two or more) specifiers to indicate anchors and a possible “direction” of the connection. Each concept relationship has a type (e.g. direct link, inhibitor, ‘part of’ or prerequisite). Such a DM structure

representation captures the types of relationships that can be encountered in most AH systems. For example, in KBS-Hyperbook one may see the dependency graph of all the KI's (knowledge items), in AHA! there are binary relationships of arbitrary types, and in APeLS there is a relationships map in a Narrative Model, by which adaptive logic is represented. In the GRAPPLE authoring tool [38] a distinction is made between concept relationships that have a meaning in the *subject domain* and relationships that have a *pedagogical meaning*. DM contains subject domain relationships (like *kind-of*, *same-as*, *special-type-of*) and the Conceptual Adaptation Model contains pedagogical relationships (like *prerequisites*). Applications in very different domains make use of very differently named relationships. The CHIP art recommender [33] uses semantic databases (represented in RDF) to connect artworks with “creators”, “techniques”, “subjects”, etc.

2.1.2 User Model Comparison

The User Model (UM) has to be created and kept up-to-date to represent user knowledge, interests, preferences, goals and objectives, action history, type, style and other relevant properties that might be useful for adaptation. Some systems also look at the environment in which an application is used, device properties, work context, etc.

UM properties are usually separated into *domain dependent* and *domain independent* properties. The user typically has domain-independent properties like identity, name and password, all with simple (atomic) values, but UM may also have more complex properties such as a collection of groups the user belongs to, preferences, a number of learning styles, work environment, and so on. The domain-dependent properties of a UM (for a given user) typically consist of some entities, objects or concepts, for which we store a number of attribute-value pairs. For each entity there may in principle be different attributes, but in practice most entities will have the same attributes.

As domain dependent properties we see that most entities in a UM represent concepts from DM, forming an *overlay* over DM, mapping the user's domain-specific characteristics like knowledge of concepts over the domain knowledge space. There may be more domain dependent properties, such as test results and learning objectives which can be problem solving tasks or short term objectives. Typically

these need to be represented in a DM as well in order to be used for adaptation. Thus, even for properties like learning goals the UM will be an overlay of DM; however, not all domain dependent properties should necessarily belong to an overlay. There can be aggregation properties like an “average knowledge” or auxiliary like “has seen an introduction page”, which are difficult to express in a UM as an overlay.

KBS-Hyperbook represents knowledge through a *knowledge vector*. The values represent a “confidence” or “probability” of the user’s knowledge for each concept (from DM). In Interbook (and AHA! and APeLS) the meaning of a knowledge value is *how much* the user knows about the concept. Such values can be more easily aggregated into knowledge values for higher level, composite concepts.

Despite the differences there is a striking commonality between the systems (not including AHA! or GALE): UM has a fixed structure, with predefined attributes per concept, targeting one specific application area, which in these cases is learning (education). AHA!, and inspired by that also GALE, offers no predefined UM structure. Concepts in AHA! and GALE can have arbitrary attributes. A designer can choose which attributes to define and use them depending on the type of application (s)he wishes to create.

2.1.3 Adaptation Model and Adaptation Engine Comparison

Adaptive systems adapt the presentation, the information content and the navigation structure to the user’s level of knowledge, interest, navigational style, goals, objectives, etc. To this end an Adaptation Model (AM) has to be provided, indicating how concept relations in a DM affect user navigation and UM properties updates (for instance whether the system should guide the user towards or away from information about certain concepts). AM actually always consists of two different aspects: rules to translate user activities into UM updates and rules to adapt the presentation to the UM state.

For computing UM updates the most popular technique is to use “forward reasoning”, meaning that an event leads to a conditional action. This is expressed by means of *event-condition-action rules*. The updates lead to more conditional actions, etc. This most closely resembled *triggers* in database systems and to some extent is also comparable to “forward chaining” in rule-based reasoning systems.

Systems using forward reasoning include Interbook, AHA! (or GALE) and to some extent also APeLS. Through forward reasoning one can calculate high-level UM properties, and have their values ready immediately when needed. KBS-Hyperbook uses deduction rules which allows for “backward reasoning”, trying to deduce UM values from events that have happened or from other UM values, somewhat like how rule-based reasoning systems may use “backward chaining” to find evidence for a proposition. Other systems also use backward reasoning but only for rules that determine the adaptation of the presentation. An example of information that is typically calculated (backwards) when a link to a concept needs to be presented is the *suitability* of the link destination for the user. In designing GALE the developers tried to offer a truly generic adaptation engine, which allows to define both types of reasoning for all aspects, i.e. for updating UM as well as for adapting the presentation. Instead of “forward” and “backward” reasoning some other models like the General Ontological Model for Adaptive Web Environments GOMAWE [39] and the Generic Adaptivity Model [40] use the terms “push reasoning” versus “pull reasoning”.

Just like for the structures in UM most systems only have *predefined* types of adaptation rules, specialized for their application. Interbook for instance has a rule that sets a specific amount of knowledge (1.0) to a suitable concept that is studied, and adds a small amount (0.1) to an unsuitable concept, each time that that concept is accessed. AHA! was the first system to allow authors to define arbitrary rules, thus enabling the creation of very different types of application using the same base system. As an illustration of this flexibility a student created an almost perfect simulation of Interbook in AHA! [41]. GALE was developed in such a way to extend this flexibility and extensibility further by also not prescribing the language used to define the rules.

2.1.4 Other “models” in Adaptive Systems

The GAF reference model research [21] has identified other parts of adaptive systems and frameworks that have a specific role. For instance, APeLS explicitly models learning goals or required competencies. Interbook can use concepts themselves as a learning goal to then generate a “Teach Me” page that offers direct guidance. The GAF model includes a Goal Model to represent this functionality. GALE follows the

Interbook approach of not modeling goals as a separate type of object but simply as a DM concept.

Another element in GAF is that it distinguishes between the User Model and a Group Model. Recommender systems may form user groups, consisting of people with a similar interest, and select information based on the commonalities between the group members. In general users may belong to different groups at the same time. The aforementioned adaptive systems, Interbook, KBS-Hyperbook, APeLS, and also AHA!, used mainly in education, do not accommodate a group model. GALE does support group modeling but this feature is not currently enabled so as to guarantee privacy of user data.

GAF also distinguishes between an *Application Model* (with rules that are essential for the application but not for personalization), the *Adaptation Model* (with rules for updating UM) and the *Presentation Model* (with rules for creating a personalized presentation based on UM). This follows an earlier distinction between Adaptation Model, Presentation Model and Navigation model in the General Meta-model for AH system described in [42]. Although these models deal with conceptually different aspects of an adaptive application the same rule language can be used to express these three parts. In GALE we treat all these parts as being the Adaptation Model.

2.2 Evaluation of Adaptive Systems

One of the first works about the evaluation of adaptive system dates back to 1990's. Totterdell and Boyle [26] present a few metrics to be related with different components of the logical model of an adaptive system interface. In 2001 Chin [43] said that only one third of the publications of the journal User Modeling and User-Adapted Interaction (UMUAI) for the past 9 years contained at least a small part of the evaluation of a system. However, it is important to highlight that the number of works from 1998 to 2001 with some evaluation is twice that of the five years before 1998. This is an indication on how importance of the evaluation has been increasing in the area of adaptive systems. Paramythis et al. [44] say that between 2007 and 2009 all articles published in the UMUAI, with the exception of surveys and introduction to special topics, have at least a small part of evaluation

of an adaptive system. The focus of this thesis is on the evaluation of an adaptive course created and served by AHA! and GALE, based on empirical evidence and questionnaires. We also propose a continuous-time layered evaluation variant of the layered evaluation described in Section 2.2.2.

There are different techniques and methods to evaluate an adaptive system. In the next subsections we describe a few of them.

2.2.1 Comparative Evaluation

The first attempts in the evaluation of adaptive systems compared adaptive systems with non-adaptive versions of the same systems. These works, that we call *comparative evaluation*, became very popular in the 90's and in the beginning of 2000 [24, 45, 46, 47, 28, 48, 25]. It is natural that this kind of evaluation became popular in the first works of adaptive system evaluation, since the systems were created to assist the users to find their goals and solutions. Typically, a comparative evaluation is made with two groups of users navigating through the system: one group using a version of the system with adaptation, and another group using a version of the system without adaptation. At the end of the navigation, the researchers compare different things, such as the navigational paths, test results and questionnaires. The main goal of this kind of evaluation is to analyze whether the users have substantially different results (less effort, better grades, less navigation etc) between the groups. Therefore a comparative evaluation should bring some insights about the effectiveness of these systems. However, in 1990 Totterdell and Boyle [26] pointed out the potential problems that this kind of evaluation could have. Paramythis et al. [44] summarized Totterdell and Boyle's work and highlighted the following problems:

- Selection of non-adaptive controls. An adaptive system has different states for different users. The authors call *state* the current goals, navigation paths, learning style and all the characteristic that the system can store from the user (the values stored in the user model) and the adaptation that the system provides at that moment. Therefore, the biggest problem is to select one of these states to compare to the non-adaptive version. The question is which of these states is appropriate to compare both systems? The appropriate state

should be selected from the best current practice. In this case, what should be the “best” state, since the adaptive system should have a very large space of potential system states. This is a very hard task.

- Selection of equilibrium points. All adaptive systems require a minimum “time” and navigation to “get to know” the user characteristics, goals etc, before the system can start to adapt the content to the user. The time required depends on the system. Consequently, the evaluation must be done after the user model has been updated, and after any other effect that influences the adaptation of the system has already taken place. Another important point to be considered is to identify the exact moments when the system reaches new points of equilibrium. The point of equilibrium is considered by the authors as the interval of time wherein the adaptation does not work while the user is acting (navigating, performing tests and tasks etc) through the system.
- Dynamics of adaptive behavior. An adaptive system can adapt itself in different ways for the same user, but this adaptation can sometimes not be beneficial for another user. In this case, the evaluator must show that that state is beneficial to that user, but the system has different “optimal” states that need to be found. In this kind of evaluation, all these states should ideally be taken into consideration.

The problems observed in the comparative evaluation motivated new proposals of evaluation research, such as the layered evaluation described in Section 2.2.2. For example, the works on comparative evaluations suggest that it does not produce generalizable results, i. e., they only work for the system or application that is being evaluated. Another point is that the results do not present sufficient data to allow the researchers to detail the system behavior effects, as pointed out by [27]. Brusilovsky et al. also observed that it is also important to remark that comparative evaluations do not evaluate an exact aspect of adaptation, hence it becomes very hard to point out the causes of the “success” or “failure” of the adaptation. Specifically, it is very hard to identify in which conditions and why one aspect of adaptation can be applied to reach the goal.

2.2.2 Layered Evaluation

It is common to find in the literature two kinds of evaluation of adaptive systems: the traditional (largely called “evaluation as a whole”) and the layered evaluation. The traditional evaluation consists of applying the methods over the whole system, without distinguishing between parts or layers that could be evaluated separately. Consequently, the system is treated as one block only [28, 47, 49]. For this reason, it is not possible to identify in which aspect of adaptation the possible problem is (if there is any). Brusilovsky [1] emphasizes that the evaluation “as a whole” means that the system needs to be already developed completely. Thus, the required changes that could be deduced from the results of the evaluation may not be easy to implement anymore, leading to an extra effort of development. Considering only the system layers, the traditional evaluation does not give feedback for each of these layers, and then the project pattern can not be reused in other systems. On the other hand, it is possible for the layered evaluation to identify problems and pitfalls for a specific layer of the system (the layers can be evaluated in union or concurrently). Consequently, the authors and developers of the AHS are able to focus on the resolution of the problem for that specific layer. The layered evaluation became popular in the past 12 years [50, 27, 51, 1, 52, 44].

Brusilovsky et al. [27] present the benefits of a layered evaluation in an adaptive system by revisiting a traditional and comparative earlier evaluation [28]. In [28], Brusilovsky did not have important results about the course being evaluated. For these reason, the authors in [27] decided to revisit this work using the layered evaluation to have better results. In [28], Brusilovsky realized that the adaptation did not assist the student in their learning process, because the students who used the non-adaptive system version got better results in the course than that ones who used the adaptive version. Brusilovsky et al. [27] present important results using the layered evaluation. For example, the authors observed that the students spent significantly less time on pages with status “nothing new” (a page that has neither unknown outcomes nor unknown prerequisites) than in pages with status “ready and recommended” (a page that has no unknown prerequisites and at least one unknown outcome concept). In this example the results show that the students considered

the status of the page they were reading, and the system worked as idealized by the system's author.

2.2.3 Empirical Evaluation

A good way to evaluate an adaptive system using a traditional or layered technique is through empirical research. An empirical evaluation observes participants performing tasks and tests as described by Chin [43] as “the appraisal of a theory by observation in experiments.” It is out of the scope of this work to present the plethora of works on empirical evaluation, and here we restrict our discussion to empirical evaluation of AHS. For the AHS the most important thing is the user, since the AHS need the characteristics, knowledge and goals of each user. For this reason, the participants of empirical evaluation can give important feedback about the system if the experiment are well formed and controlled. For example, an empirical evaluation considers the frequency with which a user accesses the system and the number of errors made by a user to complete a task. Chin contributed with his paper by presenting the common errors and problems found when researchers are creating their experiments. A common error is that researchers do not consider the knowledge required for the task, the reading skill or the visual perception. All these aspects influence the measurements and the results. There are a few other problems that are present in the experiments, for example, the nonexistence of guidelines and documentation of the experiment, inconclusive results, an insufficient number of users, leading to results without statistically relevant meaning and experiments with insufficient and adequate control. All these issues lead to bigger problems, such as the impossibility to repeat the experiment and the difficulty to obtain similar results with a different group of users.

To avoid problems in experiments, Chin suggests guidelines to assist researchers elaborating an experiment. We highlight the following suggestions: randomly assign enough participants to groups, be prepared to discard participant data if the participant requires interaction with the experimenter during the experiment and run a pilot study before the main study. The author also suggests basic measures to appear in the empirical experiments, such as the quantity, the source and the users' prior knowledge, the method and technique of the analysis and the raw data. These

empirical studies were also observed in different areas of research, such as software engineering [53], medicine [54, 55] and psychiatry [56].

2.2.4 Evaluation Frameworks

From 2000 onwards the researchers and developers of AH systems started thinking of the development of frameworks, guidelines and patterns for the evaluation of AH systems [50, 52, 57, 58, 59, 44]. The main goal of these works was to eliminate the problems in the evaluation methods or the results presentation.

Gena and Weibelzahl [58] present an approach to evaluate AH systems on the Web. Their approach is based on a user-centered evaluation (UCE). They focus on the evaluation of AH system based on each phase of development: requirement phase, preliminary evaluation phase, and final evaluation phase. For each phase they present evaluation techniques and examples. At the end, the authors suggest solutions and work to be done while applying the techniques to avoid problems and pitfalls.

In 2008 Van Velsen et al.[59] present a systematic review of UCE for adaptive and adaptable systems. It is important to remark that the UCE are the methods used in the evaluation, they are not the whole evaluation as we stated before when we discussed traditional and layered evaluation. For De Jong and Schellens [60] the UCE is used to reach three main goals: verify the quality of the application, detect problems and support the decisions of implementation and design.

For Van Velsen et al. an adaptive system “tailors its output, using implicit inferences based on interaction with the user” while adaptable systems “use explicitly provided input to personalize output.” The authors divided the analyzed works in four categories: concerning attitude and experience, concerning actual use, concerning system adoption and concerning system output. The concern about “actual use” is presented in more than 50% of the analyzed works. The three most used variables to evaluate the concept were usability, perceived usefulness and appropriateness of the adaptation. Besides the identification of the methods used in an UCE, the authors presented a model to be used as a framework to guide the evaluation process or the presentation of the results produced by an UCE. The most commonly used method was the questionnaire. For the authors, the questionnaire was not well formulated and could not be used to evaluate an adaptive system,

which means that the quality of the evaluation was not significant and should not be used for further analysis. Some of the problems found by the authors are the nonexistence of documentation about the procedure and methods used in the evaluation (making it difficult to replicate the same experiment), experiments with the developers' opinions (little criticism about the system) and only little data about the usability of the system (which means the way the users see the system).

Paramythis et al. [44] in 2010 propose a framework to guide the layered evaluation of interactive adaptive systems (IAS). The authors consider an IAS as systems that have an interactive front-end and are capable of self-adaptation (applied to, or experienced through, the aforementioned interactive front-end), which means that it is not restricted to AH systems. The proposed framework is the union of previous approaches presented in the literature (mainly the ones which is presented in [26, 27, 61, 51, 50, 52]), and it proposes the decomposition of an IAS in five layers: collection of input data, interpretation of the collected data, modeling of the current state of the “world”, deciding upon adaptation, and applying (or instantiating) adaptation. The authors say that a layered evaluation does not have to evaluate all the proposed layers, it depends on the nature of the system and the layers relevance. The authors also summarize and presents the common methods and techniques used in the layered evaluation of IAS, discussing the application of these methods and techniques clearly and presenting the best practice to a layered evaluation. The authors focused on the methods and techniques used in formative evaluations. Formative evaluation aims to identify shortcomings or errors in a system in order to further improve it and to guide the system design and development. In contrast, summative evaluation aims to determine the value or impact of a system. Therefore, the authors present the variables used in the evaluation and relate them to each layer of the system, showing practical examples of that.

Paramythis et al. [44] mention that “Formative evaluations may be more time- and labour- intensive compared to most forms of summative evaluation due to relying more on qualitative methods.” Indeed, the observational methods suggested by the authors for testing the system with real users take considerable amount of time and, in most of the case, it takes more than one session test. Another point is that the *thinking-aloud* method and the *co-discovery* method, for example, may interfere

with the participants, sometimes slowing them down or, even, making them behave in different ways [43].

We consider the methods and techniques proposed in [44] to elaborate our evaluation method, which we present in Chapters 3, 5, 4 and 7. For example, we do use users to evaluate the adaptive course presented in this thesis, but in that case the evaluation concerns the impact of the system on the student's behavior instead of the system's errors (correctness of implementation).

3 Preliminary Evaluation of an Adaptive Course

The *Adaptive Hypermedia for All!* (AHA!) development lead to GALE [22, 2]. As a consequence, GALE can be considered the AHA! version 5.0. The AHA! development started in 1996 [9], and it is considered pioneering work in the area. It is important to remark that all the applications provided by AHA! can still be served by GALE. In this chapter, we consider the evaluation of an adaptive course provided by AHA! version 3.0.

An adaptive course is an educational application built in and provided by AHA! (and also GALE). The main goal of the evaluation is to study the combined effect of link structure and annotation/hiding on the navigation patterns of users. This is a preliminary evaluation of the adaptive course, and our main goal is to identify and document problems in order to improve the quality of the evaluation and apply a modified version of this evaluation in later courses for further analysis, as shown in Chapter 4.

We first introduce, in Section 3.1, the adaptive course by itself and its adaptation rules, and then present the dataset used in the evaluation in Section 3.2. In Section 3.3 we present the methodology and the results obtained in this evaluation. At the end, in Section 3.5, we present the conclusion of this evaluation.

3.1 Description of the Adaptive Course

The AHA! system is heavily inspired by the AHAM reference model [16, 19]. Consequently, AHA! do not distinguish between an *Application Model* (with rules that are essential for the application but not for personalization), the *Adaptation Model* (with rules for updating UM) and the *Presentation Model* (with rules for creating a personalized presentation based on UM) as proposed in [42]. In AHA! (and also GALE) all these three parts is considered the *Adaptation Model*.

An adaptive course in AHA! (and also in GALE) consists of concepts that are connected to pages (or concepts). The pages contain information (text, images, tests, and so on) and links to concepts. Links must refer to concepts instead of pages, and the link destination page may be adaptively selected by the adaptation engine. The tests are multiple choice forms, where the questions and answers can be randomized by the system (we randomize the order of the question and the answers,

but we do not change the correctness of the answers and the answers are randomized for its own question).

Knutov et al [15] present a taxonomy of all the existing adaptive techniques until 2009. One of these techniques is called Link Hiding/Annotation and it is part of the navigational techniques used in AHA! (Brusilovsky [13, 14] uses the term *adaptive navigation support* in his taxonomy with a slightly difference in the techniques names, he called *adaptive hiding/annotation of links*). The link hiding/annotation consists of the use of different colors to present the links on the pages, perhaps with additional icons. In general, the author of the application defines rules to determine the conditions under which a presentation class is associated with a link. One of these presentation classes defines a link with the same color of the normal text, and it is not underlined. Consequently, the link is “hidden” in the text. The link colors suggest which links are relevant to each student. This approach is for instance discussed in [62]. In spite of other techniques provided by the AHA! system, the focus of this work is the link hiding/annotation technique, which is part of the adaptive navigation technique shown in the taxonomy presented by [15]. For AHA! (and GALE) courses the default presentation classes are called *bad*, *good*, and *neutral*, and they have the following meaning and presentation style:

- The *bad* links point to non-recommended concepts, which means that according to the rules defined by the author, the student is expected to study something—do some reading or perform some tests—before accessing these concepts. *Bad* links are colored in black and are not underlined, which implies that they are indistinguishable from the textual information of the page. So, bad links are hidden within the text, though they are fully operational and can be clicked on at any time. An important remark about the term “bad” used here is that it refers only to the link class, there is nothing really bad following these links.
- The *good* links point to a recommended concept that the student has not yet visited after it became recommended. *Good* links are colored in blue.
- The neutral links point to a recommended concept that the student has already visited after it became recommended. Neutral links are colored in purple.

Because of the choice of adaptation and colors the hypertext course looks like a standard (non-adaptive) website with pages and blue and purple links. Figure 1 shows a screen-shot of the *welcome* page from the hypertext course. The screen-shot shows clearly the *adaptive navigation support* technique, where *good* links are shown in blue text color *Introduction*, *Definition of hypertext and hypermedia* and *history*, *neutral* links are shown in purple text color: *back to course topics* and *review the “readme” page*, and *bad* links are shown in black text color, which is not visible but all the *bad* links are shown in the list of topics starting in *The architecture of hypertext systems* and ending at *The Future of Hypertext and Hypermedia*. For more details about how AHA! performs adaptation refer to [30].

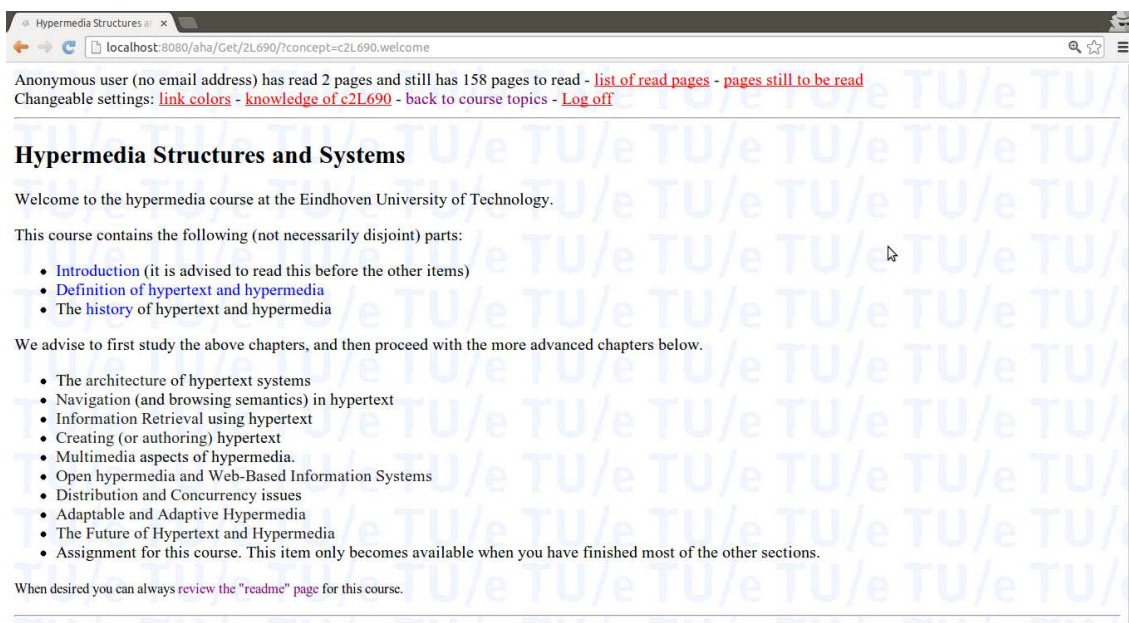


Figure 1: Screen-shot of the *welcome* page from the hypertext course, showing also *bad*, *good* and *neutral* links in the main view.

The presentation classes are used to guide the students in an AHA! adaptive course, since it shows recommended, not-recommended and neutral links. Our focus is in the students’ navigation and in the course structure of links. Precisely, we have two goals: 1) identify which pages and links influence the choices of the students and contrast this with the rules employed by the adaptation engine; and 2) verify whether the adaptation rules, created by the author of the course, influence the navigation of the students. Summarized results can be found in [49].

The course is called “Hypermedia Structures and Systems” and it has been offered by Eindhoven University of Technology since 1994. The course we evaluated took place during 2006 and 2007 and our analysis was made in 2009. The course was offered only through distance learning and had students from Dutch and Flemish universities. The course is composed of 171 concepts, including: 11 multiple choice tests, 1 final test and the Intro concept (presented only on the very first access). From all the concepts, 12 are called the “main” concepts. We consider main concepts to be the concepts that introduce a major topic in the course; the other concepts are sub-concepts of some main concept.

The main concepts are listed in (and linked from) the course’s welcome page (Figure 1 show a screenshot of this page). The welcome page is linked to a concept and it is presented every time a student logs in. Later, when the course was moved from AHA! to GALE students could go directly to a different concept by using the URL of that concept. This would happen for instance when a student tries to follow a link when his session is expired. In that case, after the student logs in (again) the system redirects him to the concept he tried to access. The 2006/2007 versions of the course served by AHA! always returned to the welcome page after logging in.

The 12 main concepts have one multiple choice test each, except for the *future* concept. Initially, on the welcome page, the system presents 3 out of the 12 main concepts as good links (recommended). The first 3 recommended main concepts are: introduction, definition and history. In the list of the recommended concepts there are 3 tests related to the 3 (already) recommended main concepts. After the student performs these first 3 tests, 90 out of 100 non-recommended links become recommended.

The student cannot repeat the tests about the main topics. The adaptation only depends on taking the test, not on the score. For the final test, the system lets the student repeat it until he scores more than 95%. The course continues with an assignment which only becomes accessible to students after they score more than 95% on the final test.

The first significant research about the structure of a non-adaptive hypertext was done by Botafogo et al [63]. The link structure analysis of a static hypertext document is feasible because its structure is static and do not change by the users’ navigation, while in an adaptive hypertext the available links may change while the

users are navigating. This leads to an initial question of our research: how do we consider the link structure of the AHA! courses, since they are adaptive and the links may change for each student?

Our research about link structure analysis focus on the *hubs* and *authorities* defined by Kleinberg [29]. In short, hubs are pages with high number of outgoing links, while authorities are pages with high number of incoming links. Regarding that we want to identify pages and links which can influence the student's choice, we also want to check if hubs and authorities can contribute to these choices.

3.2 Dataset

The data considered for this evaluation consists of the access log of the “Hypermedia Structures and Systems” course offered by Eindhoven University of Technology during 2006 and 2007. The access log contains 62993 entries of 127 students, but we disregarded 51 users with a total of 19886 entries for which the information regarding the test they performed was missing. We also disregarded 9257 entries for which the “referrer” information necessary to identify the link used to reach a concept or page was missing. There are several possible causes for the “referrer” problem such as the Web browser (e.g. Opera¹) which may reuse pages from its local cache or the students who may open a page in a new browser tab or window.

The access log is composed by the following fields:

- Student Id - an identification of the logged user.²
- Date - the date the student accessed a page, including hour, minutes and seconds.
- Concept - the concept requested.
- Activity - indicate if it is a test, assignment or concept access.
- Session - the session number provided by the server.

¹AHA! assured that the most common Web browsers, such as Firefox and Internet Explorer, did not use their local cache, but Opera ignored these settings and continued to use its cache.

²The system accepts anonymous users, but in this research the few anonymous users that existed were disregarded because their test entries were missing.

- Score - the score received by the student on a test or assignment.

When we started our evaluation research we realized that the logs did not contain all the data we would like to analyze. For example, we were unable to analyze the link presentation class of each log entry. For that reason, we replayed the logs to get all information we need. The task of replaying log was a challenge. The log entries were “hiding” important information about, for example, the student’s session or the referrer (the page where the student came from). These challenges were tackled and can be found in detail in the next section.

3.3 A Quantitative Evaluation Approach: Empirical Evidence

To tackle the “replaying log” problem, we developed software to read the log entries and process them. The log entries are first ordered by *student ID*. The log is replayed on a different server to ensure that the user models are properly initialized and that the replaying has no effect on the production server. The software reads each log entry and distinguishes whether it is a request for a concept or a test (answer). For each concept entry, the software processes the following steps: 1. Send an HTTP request to the server telling AHA! to get a concept and update the user model; 2. Process the HTTP answer getting all the outgoing links, saving the presentation state of each link into what we call *processed database*; for each test entry, the software sends an execution code to the server updating the user model with the test score.

The first step is important to get information of the concept, including the content and the links as they are at that time for a specific user. These data enable us to process the links attributes, like the link class (bad, good or neutral), and infer the referrer of the request. That processing is made by a new developed application. This application navigates through the log entries, and for each log entry, it gets the concept name (we refer to it as ConceptA), date and session. For the same user and session, it looks for the closest earlier entry (we refer to it as EarlierConceptB) into the processed database, gets the link class of the outgoing link pointing to ConceptA, and save EarlierConceptB as the referrer of ConceptA request into a new database called *results database*.

The processing has other rules to be followed:

1. If the session of EarlierConceptB is not the same as the session of ConceptA, it means that ConceptA is the first accessed concept in the session. In consequence of that, the referrer and the link class are *undefined*;
2. If EarlierConceptB does not have an outgoing link pointing to ConceptA, we treat the link class as *not defined*;
3. Consecutive accesses to the same concept is considered “reloading”, that is a student is reloading the page;

The *results database* is processed to get the following statistical measurement:

1. Number of accesses per link class;
2. Number of accesses per link class per user;
3. Number of accesses per link class per concept;
4. Number of outgoing links for each concept, known as out-degree of a concept;
5. Number of accesses for each concept;
6. The *Empirical Hub Coefficient (EHC)* of a page X, which is the ratio between the number of times that students clicked on a link of X to go to a different concept and the number of times that students accessed X. (See next section for a more detailed explanation.)
7. The time spent by students per concept in the very first access, which consists of the very first access of a student to the concept provided he followed a link or pressed the back button;
8. The time spent by students per concept on the first time they followed a link.
9. The time spent by students per concept on the very first time with click, which is the intersection of the previous two.

The first finding we want to point out is that the log file does not contain all information needed to create the results database directly. For this reason, we replayed the log to get the needed information. For example, AHA! does not log the referrer and the link class used to navigate through the course. In later research (see Chapter 6) using GALE we changed the logging so that all needed information would be available without the need to replay all sessions.

Another piece of missing information that is useful for the analysis is the system view from within which the student followed a link. AHA! (and GALE too) offers different ways to present content along with additional “views”. This can be done using CSS and DIV HTML tags for instance³, dividing the page into blocks of views. In the analysis for this chapter we do not have the view problem, because the course we are analyzing has no views with links to concepts, it has only 3 views: *content view*, with the main content of the concept; *header view*, with a settings menu and greeting message; and, *footer view*, with a copyright notice and a small photo of the professor responsible for the course. In later versions of the same course students could setup (or automatically receive) a *concepts menu view*, that is a tree menu to the concepts of the course. Links could then be followed from the menu view as well as from within the content view.

We start our analysis by measuring the number of access per link class. We notice that even though links associated with the presentation class *bad* were hidden within the text, the number of accesses via bad links was substantial: roughly 6.58% (2,293) of all the links followed were *bad* links, while 42.29% (14,738) were *good* links, and 51.13% (17,031) were *neutral* links. The issue of users following non-recommended links has been observed in other systems but was never studied closely. Our goal in the sequel is to analyze this issue further in order to discover why students go against the recommendation of the author of the course and access *bad* links.

Consider Figure 2, which shows the presentation classes used to access the 12 main concepts of the course. We consider main concepts to be the concepts that introduce a topic in the course; the other concepts are subconcepts of some main concept. The main concepts are listed in (and linked from) the course’s welcome page. Since this page provides a bullet list of topics the link hiding technique does

³DIV tags define a division or a section in a HTML page. In most of the cases, DIV tags are used to group block-elements to format them with the same CSS style.

not really succeed in hiding links here, but it does provide annotation by not showing blue links for the non-recommended links.




Figure 2: Total number of accesses performed by the students to each main concept of the course. The colors represent the presentation class of the link used to access the concept.

Note in Figure 2 that the concepts definition, history, and intro are never accessed through bad links because these links are always recommended. These concepts represent the first three main concepts or chapters of the course. Each chapter contains a multiple choice test. The “advanced” topics only become recommended after the student has completed the tests of the introductory chapters. (Only taking the tests is considered, not the score that is obtained.)

The eight “advanced” concepts shown in Figure 2 have sometimes been accessed through a bad link. The Adaptive Hypermedia concept has 42 accesses via bad links and 174 accesses via good or neutral links, giving roughly 20% of its accesses via bad links. We now focus on the bad accesses to this Adaptive Hypermedia concept. One might think that most bad accesses to any concept would come from the “welcome” page where the non-recommended links are still visible in a bullet

list. However, the collected data shows that less than 12% of the bad accesses to “Adaptive Hypermedia” came from the welcome page; hence, most bad accesses came from pages in which the link to “Adaptive Hypermedia” is effectively invisible (colored black and appearing in the main text). This suggests a behavior that we refer to as *curious browsing*: through the welcome page the students are hinted about the existence of the Adaptive Hypermedia concept, but they are not immediately drawn to it. But throughout the text, on different pages, the concept is mentioned again and again, and at some point the student’s curiosity prevails and she finds out that the term is actually a hidden link and clicks on it. In addition to “being mentioned” the term Adaptive Hypermedia is named as a technology used in the hypertext course itself and students may wish to find out what that actually does.




Figure 3: Number of students that accessed a concept via bad links. We show the five concepts with the largest number of student accesses via bad links.

We dwell on the idea of curious browsing a little further, and show in Figure 3 that 24 students accessed the Adaptive Hypermedia concept via bad links out of a total number of 42 accesses, which shows that some students accessed this concept more than once via bad links. We also point out that 76 students appear in the

access log, so roughly 30% of the students accessed Adaptive Hypermedia via bad links.

The concepts *hierarchies*, *distancematrix*, and *usermodel* can be reached from the concept Adaptive Hypermedia; so Figure 3 shows that students who follow a bad link to read about Adaptive Hypermedia do not stop at the first page but look for more hidden links to detailed information from that concept. Aside from Adaptive Hypermedia, the other concepts had roughly an average number of 1.5 accesses via bad links per student, which suggests that after revisiting the page, the students refrained from following the bad link again and waited until the link became good. This gives evidence that the students generally abided by the adaptation rules and that the adaptive engine of AHA! was to a large extent successful in guiding the students' navigation. In order to reinforce our suggestion that curious browsing for "Adaptive Hypermedia" was caused by the topic being mentioned in several places, we note that the in-degree (number of links pointing to the concept regardless of the link's presentation class) of the "Adaptive Hypermedia" concept is 14, which is the fifth largest in-degree in the course.

The observations in this section show a connection between the link adaptation and link structure: the adaptive guidance offered by the system is generally followed by the students, except for concepts that have many incoming links. Kleinberg created an important algorithm called HITS that uses ideas from random walks and measures each page according to two classes, called *hubs* and *authorities* [29]. Intuitively, good authorities are considered the interesting Web pages (i.e., Web pages that people want to refer to and that as a result have many incoming links and probably also get many visits), while good hubs are Web pages that have not just many links but many links to authorities. (The expectation is then also that hubs get many visits because they are gateways to interesting pages.) In the next section we study link structure further to find out how adaptation and link structure influence each other.

3.4 HUBS and Informative Pages

Earlier in this chapter, we thus noted that authorities may counteract the effect of link adaptation: authorities are visited even when the system recommends against

them. Now let us consider the (possible) role of hubs in adaptive courses. Hubs are pages containing a large number of links (i. e., they have a large out-degree). However, in the adaptive course that uses link hiding the number of *visible* links can change. We need a new definition for hubs that takes the difference between recommended and non-recommended (invisible) links into account.




Figure 4: Concepts whose links were most frequently used to access another concept.

Figure 4 shows a rank of the concepts whose links were most frequently used to access other concepts. We remark that twelve of the concepts presented in Figure 4 have the largest out-degree in the course. This shows how a concept is used as a hub, since some concepts not only have large out-degree, but also contain links that were indeed used by the students to access another concept.

3.4.1 The Empirical Hub Coefficient

To investigate the idea of hubs for an adaptive course, let's define the *Empirical hub coefficient* (EHC) of a page X , which is the ratio between the number of times that students clicked on a link of X to go to a different concept and the number of times that students accessed X . Intuitively, pages with large EHC are the ones used as

hubs. Note that the EHC is a number between 0 and 1, where 1 means that each time a student visited page X she clicked on a link to another concept and 0 meaning that nobody ever clicked on a link in X that leads to a different concept (most likely because X has no links).

EHC \ Out-degree	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
0 - 0.05	19	3																
0.05 - 0.1	11	4	2		1													
0.1 - 0.15	1	8	1	2	2													
0.15 - 0.2		1	1		1	1												
0.2 - 0.25		5	5		2					1								
0.25 - 0.3		1	3		1	1		1										
0.3 - 0.35		1	2		1	1	1		1	1								
0.35 - 0.4		2	1	3		1		1										
0.4 - 0.45		3	3	1	1			1										
0.45 - 0.5			2	1	1	2	1	1	2			1						
0.5 - 0.55	1		1			2	1		1	1								
0.55 - 0.6			1	1	1	2		1	1		2							
0.6 - 0.65				1	3		1	1					1					
0.65 - 0.7				2	1	1		2	1						2			
0.7 - 0.75						1	1										1	
0.75 - 0.8										1								1
0.8 - 0.85																		
0.85 - 0.9																		
0.9 - 0.95																		
0.95 - 1																		

Table 1: Number of pages with EHC according to the row and out-degree represented by the column.

Table 1 contrasts the out-degree with the EHC of the pages of the course. The rows represent the EHC split into intervals of 0.05, and the columns represent the out-degree. Entry (i, j) of the table then corresponds to the number of pages having EHC falling into the interval represented by row i and out-degree represented by column j .

In the hypertext course, links to a given concept could be found in many other concepts. The course is quite densely linked; links to a single concept can even be found in the subconcepts of different *main concepts* or *chapters*. This tends to decrease the EHC of the concepts since students are unlikely to revisit concepts that they have already visited while studying another chapter. On the other hand, Table 1 shows that there are many pages with out-degree between 3 and 9 and with EHC larger than 0.5. This may be another consequence of the absence of a centralized menu. In this course, when finding a link, students seemed to prefer to follow it immediately instead of waiting for the link to appear in subsequent pages.

Such behavior (“let’s visit the page now because we risk not finding it again later”) contributes to having pages with large EHC.

	Out-Degree	Concept	Total Number Of Users	First Access		First Click		First Access with Click	
				Students	Time	Students	Time	Students	Time
1	18	history	76	59	86	68	73	57	79
2	17	authoring	76	60	89	68	88	57	98
3	15	intro	76	50	182	73	116	49	174
4	14	dexter	75	57	135	64	148	56	145
5	14	welcome	76	61	71	75	62	61	71
6	12	openhypermedia	74	52	99	57	74	46	87
7	11	definition	76	49	159	70	145	49	171
8	11	www	76	62	159	59	178	51	184
9	10	link	76	60	106	60	128	51	119
10	10	microcosm	66	48	78	19	53	17	48
11	10	navigaids	72	58	24	64	26	54	25
12	10	scripting	68	50	89	14	81	14	81
13	9	dynamicview	67	45	109	33	119	29	123
14	9	hyperties	75	59	130	55	115	49	129
15	9	intermedia	74	61	93	57	92	51	94
16	9	navigation	76	54	86	69	80	52	93
17	9	sculptural	66	46	113	40	103	37	114
18	9	wwwserver	68	54	91	48	85	43	87

Table 2: First access information for some concepts. The table shows the out-degree, the number of students who accessed each concept, and three types of first access: *first access*, which consists of the very first access of a student to the concept provided she followed a link or pressed the back button of the browser after visiting the concept, *first click*, which is the first time a student followed a link from the concept, and *first access with click*, which is the intersection of the previous two. For each type of access, the table shows the number of students and the average time in seconds they spent on the page.

A standard statistical measure that can be used to see the correlation between the out-degree of a page and its EHC is the correlation coefficient [64]. For the data shown in Table 1, the correlation coefficient is 0.71, which means a large out-degree implies large EHC.

3.4.2 Informative Pages

As Table 1 shows, this course has many pages with a large out-degree, which can be a problem since a myriad of links on a page may generate confusion about which link the student should follow. On the other hand, pages with large out-degree are beneficial for decreasing the depth of the link structure of the course, i. e., the

minimum number of links that must be followed between any two pages that are furthest apart in terms of number of links.⁴

The hypertext course has 18 concepts with out-degree larger than 8, which is a quite large number. We examine this issue further with Table 2, which shows concepts ordered by out-degree. For each concept, it shows the total number of students that accessed that concept and three types of first access: first access, first click, and first access with click. For each of the three types of access, Table 2 shows the number of students that had an access of that type to the concept, and the average time in seconds that the students spent in the Web page during such access.

Table 2 reveals an important characteristic that motivates us to define *informative pages*. We classify a page as an *informative page* if it contains information that can explicitly guide the student on the decision about which link she should follow. Therefore, the presence of *informative pages* mitigates the problem inherent to pages with large out-degree. An example of an *informative page* is the *history* concept. This concept contains textual information regarding the order according to which the student should follow its links. Note that for the *first access with click*, students spent 79 seconds on the page on average, which gives evidence that the student indeed read the textual information in the page before following a link. On the other hand, consider the concepts *welcome* and *navigaids*. The concept *welcome* is special because on the first visit (when just starting the course) it actually shows a “readme” page which has a lot of information and just one link, to itself. After that it becomes a page with almost no information and a large number of links. The “readme” is studied carefully, requiring 71 seconds on average. The concept *navigaids* is a typical example of a non-informative page: it has very little text, and a large number of links. There is not enough text to really guide the student into choosing a particular link. The reading time of the first access is only 25 seconds.

Now, consider the concepts *microcosm* and *scripting*. For *first access*, note that 48 students accessed *microcosm* and 50 students accessed *scripting*. However, for *first click*, only 19 and 14 students accessed *microcosm* and *scripting*, respectively. This means that most students that accessed these concepts never followed their links; they most likely clicked on the back button of the browser. We see this kind

⁴In the terminology of [63] this would imply a high *compactness* value.

of concepts as (de-facto) dead ends. (They either have no links or only links the user already followed.)

Table 2 corroborates the idea that *informative pages* play an important role guiding the students through the maze of links that Web pages with large out-degree can create. We observe that students spent a large amount of time in their first access to the concepts, suggesting that they stopped to read the information provided in the page. *Informative pages* combine the functionality of a *hub* with a normal course page, whereas other pages that are just *hubs* but have a short reading time offer a lot of choice but no guidance. Such pages are to be avoided. The author should add information to the pages that can guide the student through the many links they offer.

3.5 Discussion and Conclusion

Our analysis was made through an empirical evaluation that had some missing data. As a consequence of our findings, and looking for the reasons the students follow bad links, we considered the possibility of asking the student why they are following bad links. Some of our conclusions are educated guesses based on inferences that might not be correct. In later research using GALE we implemented this through a form that would come up after a student followed a number of bad links and ask for the student's reason for following bad links. We present the results of this new run in Chapter 4.

Our analysis is important for the authors of adaptive courses. Authors should not just consider the “static” link structure of the courses they create but also consider the way their students really navigate through the course, thanks to or in spite of offering adaptive navigation support.

It is also interesting for authors to know how the students are answering their tests: knowing which questions were presented to each students and what their answers were. This could not be done with the AHA! version used in the experiments described in this chapter. We did implement two plugins for GALE that can present most of the measures and results used in this chapter and also this additional information about the performance on tests. In Chapter 8 we show the Analysis plugin and Test Analysis plugin in detail.

4 Second Evaluation of an Adaptive Course

This chapter presents and discusses in detail the results obtained through an evaluation of an adaptive course. This course is served by GALE, which is an evolution of AHA! system, and is the platform used for the adaptive course evaluated in the first evaluation in Chapter 3. In Chapter 3 we presented an earlier analysis about the influence of adaptation and link structure on the students' behavior (an overview can also be found in [49]). The earlier analysis was based solely on data because we analyzed the logs after the end of the course, therefore we were unable to ask the students' opinions. Getting meaningful information from the logs posed a challenge because the logs did not contain all data we wished to analyze. Fortunately it was possible (for almost all the access log entries) to retrieve or discover the desired data by replaying the logs for each student. In this chapter, we tackle these challenge by developing a few tools to interact with the students and by saving more data about the adaptation and the student's navigation in the logs. These implementations allowed us to analyze the logs without replaying them. We present these implementations in more detail in Chapter 8.

We start this chapter with a brief description of the adaptive course in Section 4.1. A brief overview of this chapter can also be found in [65]. The evaluation is based on quantitative and qualitative methods. We use data obtained from the access log, test logs, questionnaire, and exam grades from bachelor students at the Eindhoven University of Technology. In Section 4.2 we describe the dataset used in this evaluation, and in Section 4.3 we discuss the results and the conclusions obtained by our evaluation. For example, our study suggests that link hiding and link annotation as used in this adaptive course is a fairly unobtrusive adaptation technique. In Section 4.4 we revisit our study about hubs and informative pages presented in Chapter 3.

To understand the way users navigate through an adaptive course, our goal is to identify which pages and links influence the choices of the students and to contrast this with the test logs, questionnaire answers and exam grades as well as the adaptation rules employed by the adaptation engine. Another goal is to verify how the adaptation rules created by the author of the course influence the navigation by

the students. In Section 4.5 we discuss the results obtained and presented in this chapter.

4.1 Description of the Adaptive Course

GALE is an extensible generic and general purpose adaptive hypermedia engine. Chapter 6 presents the GALE system in detail. In this section we present the adaptive course used for the evaluation in Section 4.3.

An adaptive course in GALE consists of concepts that are connected to pages. The pages contain information (or tests, exercises, assignments or anything else that can be represented by Web pages) and links to concepts. Links in GALE always refer to concepts instead of pages, and the link destination page may be adaptively selected by the adaptation engine. A page may also contain links to external pages on the Web, but we shall not consider these here. One approach used in GALE to adapt a course is to use different colors for the links of the pages. The author of the course defines rules to determine the conditions under which a presentation class is associated with a link. This approach was used in the “Hypermedia Structures and Systems” course (to simplify we call *Hypermedia* course) presented in Chapter 3. Indeed, the “Design-Based Learning Hypermedia” course (to simplify we call *DBL* course) presented in this chapter is an update of the *Hypermedia* course.

The DBL course was offered to bachelor students at the Eindhoven University of Technology. It is an adaptive course that introduces the basics concepts of hypermedia and the Web. This course has existed for a long time but the content was updated several times and the course as a whole was moved from AHA! to GALE. This has also allowed us to use the extensibility of GALE to perform more extensive logging, as needed for our evaluation.

All the concepts were reviewed by the author during the course update. The author also added five new concepts: CSS, Empty Fragment, FOHM, Mundaneum and Ties/Hyperties. The big differences between the two courses are in the tests and in the logs and in the (possible) presentation of a navigation menu. In the earlier AHA! version each test was associated with a concept and the result of the test was a “knowledge” score for that concept. In the new *DBL course* each question of a test is linked to a concept, allowing for more detailed feedback. In the old course

students were required to repeat the final test until they scored 95%, at this new course the percentage was decreased to 90%. The logs store new data, besides the date and time, the user, the concept and the score (if it is a test):

- view - indicates from which view the user followed the link (the navigation menu or the content view showing the course page);
- link class - indicates the link class of the followed link (good, neutral or bad);
- referrer - indicates the concept the student came from.

These data are logged to solve the problems we had in our earlier experiment, for more details about the problems see Chapter 3. Briefly, the *view* is logged to know whether the student follows links from the main content part or from the menu of the course. We may observe a different effect of the adaptation or different overall user behavior by knowing whether students mainly use the navigation menu to navigate through the course or whether they mainly use the links embedded in the pages (as in the old course). The *link class* and the *referrer* are logged to enable analysis of link sources and of users clicking on non-recommended links without us having to replaying the logs to retrieve that information.

In our evaluation, an important characteristic about GALE courses is the way GALE present the colors of the links, which we call presentation classes. These classes are presented in Chapter 3, but it is very important to our evaluation that we recall them briefly here. If this thesis were adaptive, not in paper, we, based on your previous knowledge, would adapt (enable or disable) the next paragraph. Here, we intentionally show the next paragraph with the definition of the presentation classes in GALE. These classes are named “bad”, “good”, and “neutral”, and have the following meaning and presentation style:

1. The *bad* links point to non-recommended concepts, which means that according to the rules defined by the author, the student is expected to study something else—do some reading or perform some tests—before accessing these concepts. *Bad* links are colored in black and are not underlined, which implies that they are indistinguishable from the textual information of the page. So, *bad* links are hidden within the text, though they are fully

operational and can be clicked on at any time. An important remark about the term “bad” used here is that it refers only to the link class, there is nothing really bad following these links.

2. The *good* links point to a recommended concept that the student has not yet visited after it became recommended. *Good* links are colored in blue.
3. The *neutral* links point to a recommended concept that the student has already visited after it became recommended. *Neutral* links are colored in purple.

Because of the choice of adaptation and colors the hypertext course looks like a standard (non-adaptive) website with pages and blue and purple links. In Figure 5 we show an example screen-shot of the DBL course served by GALE (and called upon from the Sakai learning management system).

At the beginning of the course (upon the very first access, and only on that access), the student gets a page with some explanation about the course and adaptation that is used. An important message to the student in this *intro* page is that in (almost) all the cases when the student request a page, the adaptation process is done, even when the student clicks on the *back* button of the browser. The message in this *intro* concept presents the following sentences:

The course pages are adapted each time you request them from the server. This should also happen when you use the “back” button of the browser, but some browsers (or browser configurations) may refuse to do so in which case you don’t see the adaptation. There is nothing we can do about this.

...

*Links in this course are not underlined (no, there is nothing wrong with your browser). Links appear mainly in three colors: **blue**, **purple** and **black**. Since the main text is also black there may be links you don’t see. These links are black because the system thinks these links are not (sufficiently) appropriate for you at the time you view the page containing them. Don’t be alarmed, when you revisit the page later these links may become blue (or purple). The link colors are determined by the server, not your browser. You should not configure your browser to override the colors requested by the server.*

Adaptive Course Text

Setup

Paul De Bra

Hypermedia and the Web
 Introduction
 Definition of Hypertext
 Nodes
 Links
 Link Anchors
 Databases
 Cross-References
 Stricter Definition
 Other Definitions
 Network Structure
 Bi-directional Links
 Test about Definition
 History of Hypertext
 Architecture of Hypertext
 Navigation in Hypertext
 Information Retrieval
 Authoring of Hypertext
 Multimedia Aspects
 Open Hypermedia and WIS
 Distribution / Concurrency
 Adaptive Hypermedia
 Future of Hypertext

small hyperdocument, having only five nodes and seven links. This figure also shows that links are tied to a specific point (or word or region) within a node, called an **anchor**. Some of the links represent a hierarchical structure, while some are **cross-references**.

A simple way to distinguish the hierarchical from the cross-reference links is to consider the links on the shortest paths from the root node ("A" in the example) as hierarchical links, and all the others as cross-reference links. However, such a structure can also be created explicitly, by using link types. And in **HTML** you can use **style sheets** to define a different presentation style for these different link types.

In paper documents there are a few limited forms or types of links as well. The *index* is a source of links, but it is not possible to go directly from a word in the book to one of the pages indicated in the index, without first jumping to the index and then to the desired page. Examples of direct links are *references to the bibliography*, and, more importantly, *footnotes*. Hypertext is sometimes called the

Figure 5: Screenshot of a page from the hypertext course, showing also the navigation menu (left) and a header that gives access to settings and progress information.

This message in the intro alerts the students about the importance of the link adaptation. The author suggests that students should not follow black links. The black links, as the course presents, “...are not (sufficiently) appropriate for you...”. The author’s intention is to tell the student to follow the structure of the course in a way that at some point they will have all the links colored in blue or purple.

In earlier analysis [49], we noticed that students followed roughly 6% of the links via black links, but we could not ask them why they did that, because we analyzed the logs long after the end of the course. For our new evaluation we implemented a plugin in GALE (you can find more detail about plugin implementations in GALE in [22, 2]) that presents a multiple choice question whenever the student has followed 5 non-recommended links. We called this plugin *quiz*.

The students all started the course at the same time and they were aware that after three weeks there would be an exam about the course’s content. Consequently, they had 21 days to learn about hypertext by studying the on-line course text. They were also aware that they would be allowed to consult and navigate through the course text during the exam. (In previous years students were free to take the course at their own pace and complete it with an assignment for which there was no deadline. This led to students postponing the work, concentrating on other courses with fixed deadlines and exams first.)

The course is composed of 176 concepts, including: 11 multiple choice tests, 1 final test and the Intro concept (presented only on the very first access). It is important to remark that this course has 5 concepts more than in the course presented in Chapter 3. The content of the others 171 concepts was updated, including the multiple choice tests.

From all the concepts, 12 of them are called the main concepts. We consider main concepts to be the concepts that introduce a major topic in the course; the other concepts are sub-concepts of some main concept. Thinking of the course as a book, the main concepts would be the chapters with an introduction and the others would be the book’s sections and pages. If we consider the adaptive course as a pure hypertext course, with no adaptation, the big difference between it and a book is that the student could find the same section in different chapters. Another difference is that the students do not follow a fixed page order, they follow embedded links in

concepts in any way they like. In the adaptive case, the embedded links and their presentation class suggest how students should navigate through the course.

The main concepts are listed in (and linked from) the course's welcome page. The welcome page is linked to a concept and it is presented every time a student logs in, unless the student tries to access a concept with a direct URL but he is not logged in (for instance when his session is expired and he clicks on a link to a concept). In that case, after the student logs in the system redirects him to the concept he tried to access.

The students can set up the course to present a menu with the main concepts and their sub-concepts. The menu links are also adapted like the links in the pages.

The 12 main concepts have one multiple test each, except for the *future* concept. Initially, on the welcome page, the system presents 3 out of the 12 main concepts as good links (recommended). The first 3 recommended main concepts are: introduction, definition and history. At the beginning of the course, there are 74 recommended concepts (good links), 1 visited concept (the Intro concept) and 101 non-recommended concepts. In the list of the 74 recommended concepts there are 3 tests related to the 3 (already) recommended main concepts. After the student performs these first 3 tests, 91 out of 101 non-recommended links become recommended.

The welcome page also shows a sentence saying that the student has to perform all of the tests in order to be able to perform the final test. Indeed, to access the final test the student has to perform all the 11 tests in the course. It is important to remark that the final test is not just presented as a non-recommended link, it is blocked until the student performs the other tests. At the final test, the system lets the student repeat it until he scores more than 90% (instead of 95% from the earlier course).

The next section presents the evaluation data and analysis.

4.2 Dataset

The data considered for this chapter consists of the access log of 46 bachelor students of the adaptive course offered in 2013 and a questionnaire answered by 28 of the same group of students. These were first year students in the programs of Web

Science, Software Science and Psychology and Technology. The on-line course text was intended to prepare the students for the exam and for a later (group) assignment. The questionnaire was made available online to all the students, shortly before the exam (and left open until a few days after the exam). The students were not required to complete the questionnaire and the system did not log their identity. The anonymous nature of the questionnaire should give students more “freedom” to answer the questions honestly. The course’s access log contains 17,318 entries, where 609 entries are answers for the tests and 1,001 entries are accesses to the test concepts (the questions). So clearly students sometimes requested a (multiple-choice) test, realized they did not yet know the answers and went back to study some course pages before trying to complete the test. Note that the tests are generated randomly from a larger collection of questions, so when a student views a test but does not answer the questions the next visit to the test may show different questions. This is also true for the final test that can be repeated: each time the student will get (some) different questions.

4.3 Qualitative and Quantitative Evaluation Approach: Empirical Evidence and Questionnaires

We start our analysis by relating the test log with the exam grades (or to simplify, just grade). Students were told that the tests that are part of the adaptive course test are only there in preparation for the exam and do not contribute to the course grade. The grades that were obtained have an average of 5.2 out of a maximum of 10. (Grades were integers.) The highest grade was 7 and the lowest is 0 (only 1 student); the second lowest grade was 2 (only 2 students). 20 students obtained a grade above the average (i.e., a 6 or 7).

To understand Table 3 with the attempts and scores for the final test it is important to understand the operation of the course: only students who completed the tests embedded in each of the main concepts (chapters) were allowed (and had access) to the final test. Out of the 46 students only 12 performed all preliminary tests and gained access to the final test. Of these 12 only 8 students actually attempted to complete the final test. Students were told they should attempt to score at least 90% on that final test. The test could be repeated as often as desired

(until the score of 90% was reached). The preliminary tests could be done only once and their score was not used for deciding on access to the final test and did not influence the adaptation. The final test was tried 319 times (by only 8 students). It is interesting that 5 out of 8 students who tried the final test got an exam grade above the average. We also note that all the students who got a score higher than 90% in the final test got a grade of 6.0 or 7.0. These findings suggest that if a student followed the course and tried to get a high score in the final test, he would be better prepared for the final exam than the “average” student. Although this is just anecdotal evidence (with few students taking and passing the final test) it is good to know because before the students started the on-line course they were told that the final test would be a good preparation for the exam. The anecdotal evidence “proves” us right.

Table 3: Number of time a student answered the final test, his highest score and the exam’s grade

Student	# Tries	Highest Score	Exam Grade
Std1	1	27	7.0
Std2	1	27	7.0
Std3	3	40	3.0
Std4	34	100	7.0
Std5	47	47	3.0
Std6	49	93	7.0
Std7	51	60	4.0
Std8	133	93	6.0

The first three tests associated with the three recommended main concepts (introduction, definition and history) were performed by 40 students. Performing these tests turned the “advanced” concepts into recommended concepts, so that was a clear motivation for taking these tests. These 40 students thus clearly saw the link color changes after completing the first three tests. The questionnaire has the question: *Regarding LINK ANNOTATION, where the system presents links in different colors, perhaps with additional icons, which of the following statements*

applies? For this question the answers was (in parentheses there is the number of students who chose that answer):

- I was not aware of link annotation. (1)
- I only notice that visited links became purple. (6)
- The link annotation was clearly intended to offer guidance through the course. (10)
- The link annotation was manly intended to avoid to going to some pages. (2)
- No answer. (9)

The answers show that 12 out of 19 students were influenced by the link annotation, since they recognized the importance of the link annotation to guide them through the course or to avoid them following links.

Even though students who performed the first 3 tests then got many more recommended links, they did perform some navigation through bad links before the advanced topics became recommended. The logs contain a total of 471 visits via bad links by 36 different students. To understand why the students would visit non-recommended concepts, the system presented a question to the students who accessed more than 5 concepts via bad links: *Why are you following black links (not suitable yet)?* The possible answers were:

1. I am curious to see what happens if I click on a black link.
2. I would like to explore the course a little bit before learning.
3. The system presented a lot of black links and I would like to know what it means.
4. I read many times about one concept and it is not suitable yet, but I would like to learn about it.
5. I do not know what a black link means.
6. Other (please specify).

15 students visited less than 5 concepts via bad links and did not get this question. We only wanted to question the “repeat offenders”.

Table 4 presents the summary of the students’ answers. It is interesting to note that 13 out of 18 answers were given on the day of the exam, including 3 of them during the exam, showing that students studied the course text just in preparation for the exam, not really considering that the most important aspect of the course was that it prepared them for the later assignment. One student, called *Std1*, chose the *other* answer during the exam and he wrote: “Trying to find an answer to a question”. Std1 had already answered this question a few hours before, and at that time he chose the answer 2. This student reset his profile after the first “exploration phase”, which means that the system deleted everything about what he did from his user model, including the performed tests and visited status of concepts. (The log however remained unaffected by a reset.) We can also estimate the time students spent reading pages, by considering the time difference between two log entries for that student. Std1 showed an average reading time of roughly 9 seconds per page. Clearly in the exploration he did not spend enough time on the pages to actually read and study the text. The system adaptation influenced Std1’s steps in the course in spite of the fact that he was not concerned with learning the material and was only exploring the course. The adaptation in the course only depends on access, not on reading time. Because of the profile reset this student received the question about navigating through black links twice. And it turns out he gave a different answer the second time.

Table 4: Summary of the answers about why students follow bad links.

Answer	Before Exam Day	Exam Day	Total
1	2	2	4
2	2	6	8
3	1	0	1
4	0	2	2
5	0	1	1
6	0	2	2
Total	5	13	18

Table 4 also summarizes how we could group students who do not follow the system recommendation:

- students who have curiosity about what would happen if they do not follow the adaptive structure. It is represented by students who answered 1, 3 and 4;
- students exploring the course: students who want to explore the course before they start learning. It is represented by students who answered 2;
- late students: students who start studying the course when it is (too) late. Three of them were still following bad links during the exam.

As it appears that many students started studying for the exam quite late it is interesting to look specifically at the log for the day of the exam.

On that day the bad links access log shows 278 accesses (including during the exam), which is roughly 60% of the all bad link accesses. During the exam there are 168 accesses via bad links, representing roughly 36% of all bad links accesses. These 168 accesses are concentrated in 9 different students, where 3 of them accessed only 1 bad link each and the others have an average of 27.5 (4 students above the average). Teachers often state without proof that some students go to an open book exam without any preparation, thinking they can look everything up as needed during the exam. Our log proves that some students indeed did not prepare for the exam and tried to search for the answers during the exam, and thus often going through bad links. Before the exam day the bad link accesses are mostly concentrated in the log for 5 students who performed 120 of these accesses, representing roughly 25% of all bad links accesses. These students are in the category of student curiosity or student exploring the course according to their answer on the question about why they followed bad links.

In the questionnaire, the students answered the question: *Regarding LINK HIDING, where the main text is black, the black link is hidden in the main text but you can still click on it. Have you tried clicking on these links?* For this question, we have similar results to the one presented in Table 4: 67% of the students who admitted to have followed black links talked about the curiosity; 25% admitted to explore the course before learning; and the 8% said that they did not remember if they followed black links.

Another point to corroborate the idea of the students' curiosity is the fact that the first two most accessed concepts via bad links are the first two of the main non-recommended concepts appearing in the list of topics on the welcome page. (These are "The architecture of hypertext systems" (*architecture* concept) and "Navigation (and browsing semantics) in hypertext" (*navigation* concept) as shown in Figure 6.) The referrer of a visited concept indicates the concept the student came from. So we could also analyze whether other non-recommended concepts were visited directly from these concepts or whether they were accessed through cross-reference links from other main concepts. So we looked at the all pages from which bad links were followed. Again, the navigation and the architecture concepts appear in the top four rank, losing the first and second positions only to the welcome concept that lists all the main topics and the *TO DO* list (a page containing a list of all the concepts that were not visited by the student before.), two links that are always recommended. This indicates that when a student followed a bad link to a main topic he was not considered ready for, he did not stop at the first page of the non-recommended topic but kept following bad links on that topic. When asked about this behavior the answer to the bad link question was "by curiosity".

The use of link colors in the adaptive course text is described as *link hiding* because black link anchors appearing in a paragraph of black text in fact causes the link anchor to be hidden. However, the black links also appear more clearly visible in three places: on the welcome page that shows a list of links to the main concepts (Figure 6 show a screenshot of this page), in the optional navigation menu of the course (also shown in Figure 6), also showing bad links and in the *TO DO* list (of all unvisited concepts). In these places the differences in color should be referred to more correctly as *link annotation* instead of *link hiding*. For example, Figure 6 illustrates the presence of bad links. In the menu on the left, the items from "Architecture of Hypertext" to "Future of Hypertext" are web links in black color. Similarly, in the welcome concept (right) shown in Figure 6, the list of topics from "The architecture of hypertext systems" to "Assignment for this course" are web links in black color that can be clicked by the student. We can thus distinguish between bad links followed through annotated links versus bad links followed through hidden links. It turns out that 87% of all bad link accesses are through the three mentioned places that use *link annotation*. When we use the same three sources

(welcome, menu and TO DO list) to check the accesses via good links, the percentage falls down to 57%. The accesses that remain thus indicate that the students follow many links from the course pages but few bad links in the course pages. Link hiding is thus effective in keeping students away from non-recommended topics but link annotation is not. Or in other words, students look for good links when they are navigating through the course pages but consider all the links that appear in menu-like structures.

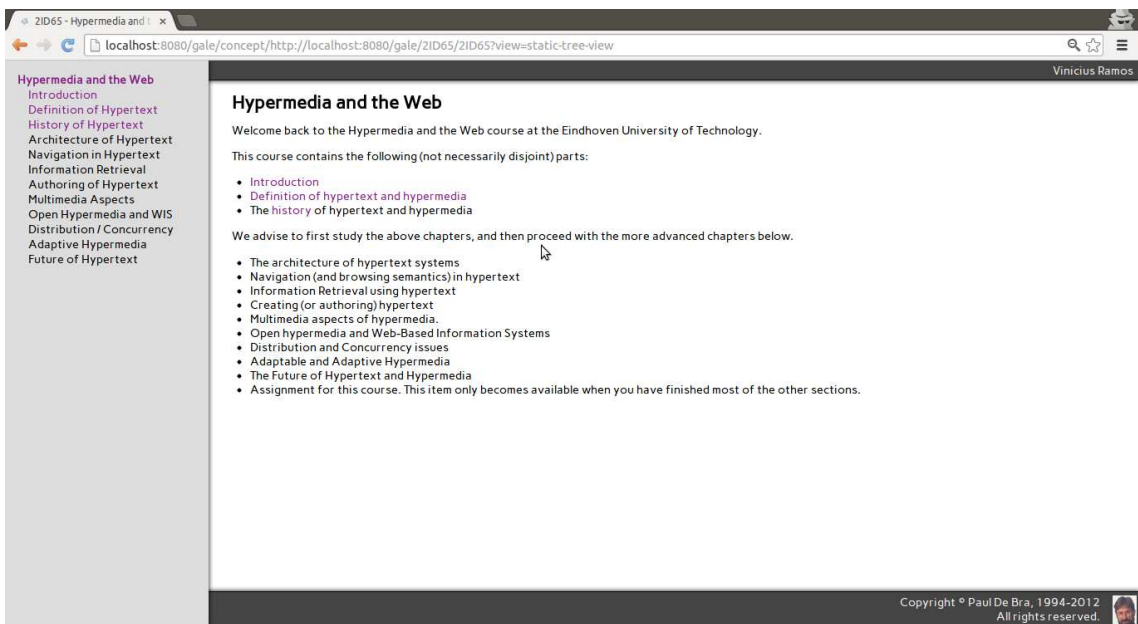


Figure 6: Screenshot of the *welcome* page from the hypertext course, showing also *bad* and *neutral* links in the menu (left) and in the main view (right).

We asked the students: *Regarding adaptation, what was your user experience?* The students answers seem to contradict the logs: 8 out of 28 students (28%) marked the option: “I hardly noticed it”, while 9 others marked the option “I found it patronizing”, and only one “found it helpful”. 10 students did not give an answer at all. It is important to recall that roughly 90% of the students completed the first 3 tests of the course. Consequently, they got 90 new recommended concepts, of which 9 are presented in the menu and in the welcome page. For this reason, it is easy to see that the adaptation is working. Yet 28% of the students said they “hardly noticed it”. Apparently link hiding and link annotation as used in this

course (having links become recommended but never making a recommended link become non-recommended) is a fairly unobtrusive adaptation technique.

When students were asked to talk about their feelings about the adaptation in the course, two of them talked about the menu that could be enabled or disabled. They would like the menu to be enabled by default, whereas in the course it was disabled by default. They see an advantage of following the menu while following links, instead of relying on links that appear within the pages. Another student said that *It's not my cup of tea. I see the advantages, but I prefer an old-fashioned paper book.* Two students also suggested the possibility of having a back link/button within each page of the course, to be used instead of the browser's back button, so as to allow them to keep navigating through the pages' content.

In Section 4.5, we discuss the results presented in this chapter and the suggestions made by the students in the questionnaire.

4.4 HUBS and Informative Pages

In Chapter 3 we started an investigation to find out how adaptation and link structure influence each other. If this thesis were an adaptive application we would recommend (in a blue or purple link, considering a GALE application) or not (black link) the reading of Section 3.4 based on your previous knowledge. Another way of adapting the current section in a GALE application is to use the conditional inclusion of fragments, where the system includes or not part of the content depending on the user's knowledge. Since this thesis is written on a paper and is not an adaptive application, we (intentionally) write (or recall, depending on what you have read before) the definitions of *Empirical Hub Coefficient* (EHC) and *Informative Pages* in this section.

4.4.1 The Empirical Hub Coefficient

We define the EHC of a concept X as the ratio between the number of times that students clicked on a link of X to go to a different concept and the number of times that students accessed X. Intuitively, concepts with large EHC are the ones used as hubs. Note that the EHC is a number between 0 and 1, where 1 means that each time a student visited page X he clicked on a link to another concept and 0 meaning

that nobody ever clicked on a link in X that leads to a different concept (most likely because X has no links).

It is important to remark that the adaptive course evaluated in this chapter is an update of the *Hypermedia* course presented in Chapter 3. We refer to the first course evaluated as the *preliminary course*. There is a big difference in the layout of the two courses: a centralized menu (as seen in Figures 5 and 6). The menu presents a list of concepts in a tree hierarchy view, where the main concepts (chapters) are the roots of the tree and the subconcepts are within the chapters.

Table 5 contrasts the out-degree with the EHC of the concepts of the course. The rows represent the EHC split into intervals of 0.05, and the columns represent the out-degree. Entry (i, j) of the table then corresponds to the number of pages having EHC falling into the interval represented by row i and out-degree represented by column j . The out-degree of a page do not consider the menu, since it is a different view in the course. In this case, our EHC definition fits this restriction. Another important remark about the menu is that we do not log the moment the student enabled that. We are able to know from which view the student came from, for example, from within the menu, and in that case we know that he has a menu. Therefore, to calculate the EHC we restricted the clicks within the *content* view. For example, if the *StudentA* followed a link within the content view from *conceptD* to *conceptC* and later he followed a link within the menu view from *conceptC* to *conceptB*, we consider that *conceptC* and *conceptB* were visited but we do not consider that the *StudentA* clicked on a link of *conceptC* to navigate to *conceptB*.

In our preliminary evaluation we stated that the student’s behavior “let’s visit the page now because we risk not finding it again later” contributes to having pages with large EHC. At this course we have a menu, which means that the student can revisit the concept later and do not have to visit that concept in the moment he see a link. In the *preliminary course* the correlation coefficient between the EHC and the out-degree was 0.71 (refer to Section 3.4.1). Thus, we would expect a decrease of the correlation coefficient in this course. Although, we still have a high correlation coefficient that is 0.65, which means a large out-degree implies large EHC.

EHC \ Out-degree	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
0 - 0.05	12	16	2	2	2					1										
0.05 - 0.1	6	3	1	3	1	1				1										
0.1 - 0.15	1	6	3	1	2	2	2	2	1				1							
0.15 - 0.2	1	1	1	2	2	1	1	3												
0.2 - 0.25	1		2		2	2	3													
0.25 - 0.3	2		1	3	4	1		1	2					1						
0.3 - 0.35			1	1	2		1			1									1	
0.35 - 0.4						1						1								
0.4 - 0.45	1			1					1	2										
0.45 - 0.5							1	1	2					1						
0.5 - 0.55				1																
0.55 - 0.6																				
0.6 - 0.65																				
0.65 - 0.7																1				
0.7 - 0.75																				1
0.75 - 0.8													1							
0.8 - 0.85																				
0.85 - 0.9																				
0.9 - 0.95																				
0.95 - 1																				

Table 5: Number of concepts with EHC according to the row and out-degree represented by the column.

4.4.2 Informative Pages

The data presented in Table 6 and Table 5 reveals an important characteristic of concepts with large out-degree. The adaptive course has 19 concepts with out-degree larger than 8, which is quite large. In one hand, a concept with large out-degree may generate confusion about which link the student should follow. On the other hand, pages with large out-degree are beneficial for decreasing the depth of the link structure of the course, i.e., the minimum number of links that must be followed between any two pages that are furthest apart in terms of number of links.

Table 6 shows concepts ordered by out-degree. For each concept, it shows the total number of students that accessed that concept and three types of first access: *first access* (consists of the very first access of a student to the concept provided she followed a link or pressed the back button of the browser after visiting the concept), *first click* (consist of the first time a student followed a link from the concept), and *first access with click* (it is the intersection of the previous two). For each of the three types of access, Table 6 shows the number of students that had an access of that type to the concept, and the average time in seconds that the students spent in the Web page during such access.

We define a page as *informative page* if it contains information that can guide the student on the decision about which link she should follow. It is important for this thesis to revisit the evaluation presented in Chapter 3. Thus, we start our analysis of Table 6 revisiting the *informative pages* we analyzed in this thesis before.

The *history* and the *intro* concepts are good example of *informative pages*. Note that for the *first access with click*, students spent 53 and 128 seconds on the page on average, for *history* and *intro* concepts, respectively. It give us an evidence that the student indeed read the textual information in the page before following a link. On the other hand, *navig-aids* is a typical example of a non-informative pages: it has very little text, and a large number of links. The reading time of the *first access* is only 25 seconds.

The lowest average time for *first access with click* is for the *scripting* concept. As we noticed in our earlier experiment, this concept is a dead end, where all the 17 students that accessed this concept, did not follow a link. We can see in Table 5 that there is one concept with EHC lower than 0.05 with out-degree equals to 10, this is the *scripting* concept. Another dead end is the *microcosm* concept. This concept has also a very low EHC of 0.07. Indeed, only 3 users clicked on a link from within this concept.

The *authoring* and the *openhypermedia* concepts are main chapters of the adaptive course. They were created to be *informative pages*, but, as presented in Table 5 they have EHC 0.33 and 0.13, respectively, which means that only a few students follows links from within these concepts, and most of them without even reading the concepts' content.

We reinforce the idea that *informative pages* combine the functionality of a *hub* with normal course page, whereas other pages that are just *hubs* but have a short reading time offer a lot of choice but no guidance. In this way, authors should add information to the pages that can guide the student through the many links they offer.

4.5 Discussion and Conclusion

The focus of this case study has been to analyze qualitatively and quantitatively whether the adaptation mechanism influences the students' learning and navigation

	Out-Degree	Concept	Total Number Of Users	First Access		First Click		First Access with Click	
				Students	Time	Students	Time	Students	Time
1	20	history	46	28	51	44	48	27	53
2	18	authoring	40	29	36	38	33	28	38
3	16	intro	48	27	128	46	123	27	128
4	14	definition	46	24	81	46	85	22	88
5	14	dexter	34	23	111	30	104	22	116
6	13	welcome	48	19	44	48	35	19	44
7	13	openhypermedia	38	31	25	36	39	29	26
8	12	www	40	21	102	33	99	21	102
9	10	hyperties	35	17	78	24	70	17	78
10	10	microcosm	19	18	97	19	92	17	103
11	10	navigation	45	32	50	43	40	32	50
12	10	scripting	25	17	9	19	10	17	9
13	10	www-servers	35	26	81	30	71	26	81
14	9	about-html	36	18	50	29	57	18	50
15	9	intermedia	40	17	57	31	57	17	57
16	9	link	41	26	70	38	67	26	70
17	9	navig-aids	36	31	25	35	40	31	25
18	9	retrieval	39	30	76	38	68	29	78
19	9	sculptural	18	14	67	16	60	14	67

Table 6: First access information for some concepts. The table shows the out-degree, the number of students who accessed each concept, and three types of first access: *first access*, which consists of the very first access of a student to the concept provided she followed a link or pressed the back button of the browser after visiting the concept, *first click*, which is the first time a student followed a link from the concept, and *first access with click*, which is the intersection of the previous two. For each type of access, the table shows the number of students and the average time in seconds they spent on the page.

behavior. We approached these topics in two ways: verifying the access log of the course with the exam grades and analyzing a questionnaire answered by the students.

The main analysis of the access log relates to the visits of concepts via the so-called *bad* links (a link that points to a non-recommended concept) and whether the presence of these links affected the student’s behavior. We found that students who start studying late in the course period tend to follow more bad links than students who use the entire period. We also noticed in the log and questionnaire analysis that students often follow bad links because they are curious about the linked non-recommended concept or because they are exploring the course to get to know all the material available before they start studying the topic more in depth.

One remark is that when students navigate through bad links 87% of the clicks came from the welcome page, the menu view links or the TO DO list of concepts. On the one hand this suggests that they are curious or exploring the course before

they start learning as confirmed by them in the questionnaire and by answering a specific question about following bad links during their navigation. On the other hand it also shows that the *link annotation* used on these three pages has less effect on the students' navigation than the *link hiding* that is used for links appearing in the running text of the pages.

Considering that 51.5% of the log entries correspond to the last day of learning and 26.5% of the whole log belongs to the day of the exam (including during the exam), it appears that the students were much more concerned with quickly preparing themselves for the exam and were not too concerned with a comfortable learning experience for which the adaptation is intended. The log entries accumulated until two days before the exam represent 23% of the whole log, and the bad links entries on these days represent only 15% of the total number of bad links followed. Students who started learning in the beginning of the course period followed the course's link structure and adaptation better than of the students who started learning a few days (or even hours) before the exam. Clearly, when rushing towards the exam students are not open for advice that tells them they are not yet ready for certain topics and should study something else first. They want to take it all in, as quickly as possible and in any order.

It is important to notice that there are more entries during the exam (4,495 entries) than during the whole period before the day before the exam (3,976 entries). This suggests that, independently of the exam (which the adaptive course was supposed to help the student to study for that), the students need to be stimulated to start studying with the system. The absence of classes or other contact with a tutor during the course period leads to students postponing their study activity until right before the exam.

The questionnaire gave us a good insight about the needs of the students, who gave us a few suggestions to improve the quality of the course and its navigation. For example, students suggested to have a navigation menu presented at all times (by default). We made this optional because the course was on the topic of hypermedia and we wanted to stimulate navigation by following links in the pages. But for other course topics it is certainly worthwhile to offer a navigation menu by default. Some students remarked also that the course structure (i.e., the navigation support including a menu) should be more similar to the way they are used to study: it

should be more similar to a book. The new paradigm of navigating through links in pages makes the navigation, as a student wrote, “quite difficult and annoying.”

Students also want to know how far along they are in the course. Such a “count down” counter exists in the course but it is not permanently displayed. Instead of a menu a progress bar could also be used. As the deadline of the exam was approaching, it became clear that students lose their patience in following the advice to study some concepts before some other concepts. They seem to want to click through the whole course and the whole material quickly.

We have also carried out a structural analysis of the course, comparing the definition of *hubs* to new empirical measures, called *empirical hub coefficient* and *informative pages*. Regarding the fact that the *authoring* and the *openhypertext* concepts have low *empirical hub coefficient* and, consequently, both concepts have low time reading. Most of the *informative pages* were used as they were intended for: it has many links, but the information can guide the students through the navigation over these links.

5 Continuous-Time Layered Evaluation in an Adaptive Hypermedia System

5.1 Introduction

In this chapter we present the continuous-time layered evaluation technique, it was proposed in [66]. We start explaining the layered evaluation technique, which was proposed to break the paradigm of evaluating an adaptive system as a whole [67, 68, 69, 1]. In [1], it is proposed to decompose an AH system in two distinct layers, called *user modelling* (UM) and *adaptation decision making* (ADM). The authors' proposal is based on an adaptive educational system. They define the UM layer as the evaluation of the user model of the students. For example, the UM layer is responsible for analyzing whether a student has not visited a certain important webpage, has not understood a specific material, or has not succeeded in finding the information she needs. They define the ADM layer as the evaluation of specific aspects of interaction. It uses the UM layer to adapt the AH system to the needs of the student. For example, in the hypermedia courses evaluated in this thesis, the logic of the ADM is captured into a set of *adaptation rules*. The adaptation rules determine which adaptation aspects should be selected according to the results of the UM process; i.e., these rules are responsible for the adaptive presentation, including hiding and annotation of links. The UM layer, for example, detects that a student has not succeeded completing a task, a pop-up message can be triggered by the ADM layer with additional information and links that can help the student finish the task. Figure 7 shows the layer decomposition made by the authors.

The other pillar of our layered evaluation technique is the continuous empirical evaluation introduced by Ortigosa and Carro [69]. Their goal is to evaluate adaptive courses in order to identify possible fails in order to improve or, at least, suggest possible actions to be done in these courses. They discussed situations in which an empirical analysis can help detect possible design problems leading the students to low performance (students fail to do the exercises related to the topic being studied), disorientation (students browse the course with no logical order), lack of motivation (students do not browse the course frequently), and dissatisfaction (students drop off

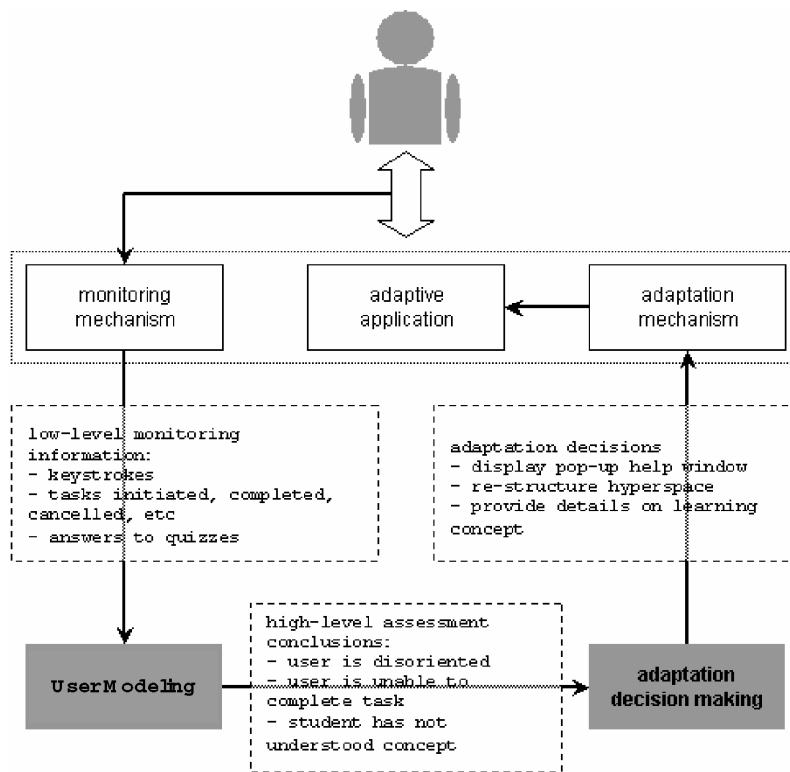


Figure 7: Adaptation decomposed (presented in [1])

the course at the beginning). Ortigosa and Carro also give some insights about which aspects of the course design could be the cause of these problems. For example, they mention that if students obtain a low performance in some activities of the course, then it may be the case that some information is not clearly explained in some webpage or the topics of the course are not well organized. Another point raised by Ortigosa and Carro was that there may exist unidentified dependencies between topics. Usually, this type of problems are very hard to identify, but the authors remark that students may need to know some topics (called prerequisite) before studying others. Then, it would be natural to request that a previous visit to a prerequisite or a minimum score in the tests associated to them should be required in order to access a certain page. Then, if there are unidentified dependencies in a course, this could become evident if students who visited the prerequisite get better results in the current topic than those who did not.

We not only consider the aspects pointed out by Ortigosa and Carro, but we also propose new aspects for the evaluation process. In Chapter 3, we present our findings about the evaluation of an adaptive course created and served by AHA!. For that course, we observed that students constantly return to the same *concept*, a phenomenon that we call repetitive browsing, and that students visit advanced pages without having visited more fundamental pages or doing required activities, which we call curious browsing. These two aspects represent anomalies in the way one would want a student to visit the pages of the course [49]. We also investigate how much time students spend on each concept. This turns out to be interesting empirical information regarding how effective the content of the pages is in guiding the students through the links of the course, especially on pages that contain a very large number of links, which could in principle lead the student not to know which link to follow [49].

5.2 Methodology

The combination of the layered evaluation approach by Brusilovsky et al. [1] and the continuous empirical evaluation by Ortigosa and Carro [69] forms the basis of our methodology. The methodology consists on applying the proposed techniques described in this section to the adaptive courses served by AHA! (or GALE) in order to detect situations where the course does not satisfy the student's needs or to identify possible problems in the course's design (the structure or the adaptation rules) to improve it, and depending on the faced problem, suggest updates to the authors. The proposed evaluation technique uses the idea of the layered evaluation, focusing on UM and ADM layers, to evaluate a course continuously. For example, the *hypermedia course* evaluated in this thesis and presented in Chapters 3 and 4 is under *continuous* update and it has been offered and revisited every year since 1994 [9]. Different groups of (bachelor and/or master) students navigate through this course every year in different contexts and with different goals. Consequently, we have navigational paths and performed test logs for each group of students with their different goals. For these reason, a continuous evaluation approach should be applied in this course so that students of later courses benefit from the improvement made after previous evaluations. The goal of the continuous evaluation is to reveal

possible problems in the course or in the way students are using the course while the course design can still be modified. Therefore, the evaluation can be used to improve the course design in real time and provide a better learning experience to the students.

The proposed layered evaluation will encompass the following aspects: the students navigational paths (i. e., sequence of pages) that students follow during the learning process, the structure of the course, the improvement on the students' knowledge, the students' possible disorientation, and their opinion expressed in response to surveys.

1. Navigation paths analysis

The navigation paths are used to represent the sequence of concepts visited in a course. Each identifier in the sequence represents a concept in the course. The main idea of analyzing the navigation path is to check if students completing a test successfully visited some page before finishing the test that students failing the test did not visit. At this point we want to find patterns among all the students who failed the test. After getting a pattern, we shall contrast it to students who pass the test. For illustration purpose, we refer to Figure 8 (adapted from [69]). In (1), we show the navigation paths for three students who failed the test **S**. In (2), we show the pattern extracted from these three students (i. e., which concepts these students visited in common). Figure 9 presents, in (3), three students who passed the test **S**. In (4), we show the pattern extracted from the students who passed the test. In order to determine why students failed the test **S**, one needs to compare the patterns for the students failing the test (Figure 8 (2)) and for the students passing the test (Figure 9 (4)). In the examples of Figures 8 and 9, one can observe that the path “2,*,5,*,7” appears only in Figure 9(4). This suggests that students who have completed the test successfully visited the concept 5 before performing the test **S**, while students who failed the test have not visited concept 5. This gives evidence that visiting concept 5 before performing the test **S** is important to the students. In one hand, the evaluation mechanism can suggest to the author of the course to add a *prerequisite* so that students can only perform test **S** after visiting concept 5. Alternatively, in order to corroborate the hypothesis that concept 5 is important for test **S**, the author of the course may add a




Figure 8: Navigation Path sequences from students who fail in the test S. (1) represents the students' paths. (2) represents the common pattern extracted from (1), where * stands for visits to zero or more concepts.

review with the content of concept 5 before the test S. On the other hand, the evidence should not be statistical significant, and the system should send an warning message to the author. We would also suggest a different evaluation approach, such as observational studies or controlled experiments, to verify the evidence.

2. Average time spent

The evaluation of an AH system for distance learning is always concerned about usability, student performance, or students' knowledge improvement. One point that should be evaluated is the average time spent by students on a concept. We can distinguish the three presentation classes in *AHA!*: *bad*, *good* and *neutral* links, which represent the status of a concept. We associate these presentation classes to the concept status. We consider the concept status: non-recommended, recommended but not visited, and recommended and visited. These concept statuses are associated with the link status *bad*, *good* and *neutral*, respectively. For example, a *good* link represents a concept that has not yet been visited (and consequently has not yet been learned by the student). We expect that students spend a large amount of time reading




Figure 9: Navigation Paths sequence from students that succeed in the test S. (3) represents the students' paths. (4) represents the common pattern extracted from (3), where * stands for visits to zero or more concepts.

a page of such a concept. On the other hand, a *bad* link represents a non-recommended concept, for which we expect that students are already missing some information or required knowledge to understand the concept. Hence, the average time spent by the students in such concepts may even be higher than for recommended concepts that were not visited. Finally, the *neutral* link represents an already learned concept. We then expect that the average time spent by the students on such concepts is lower than for the other concept statuses. It is interesting to check that the concept is already known, so that the student only needs to skim quickly through the page to find a specific piece of information of a link to be followed. Note that this idea is intrinsically related to the *informative pages* introduced in [49] and presented in details in Chapter 3.

3. One student following *bad* links

For navigation through links in an hypertext course that uses the link hiding technique in the adaptation engine, it is important to evaluate the UM state because the link hiding depends on the UM. Following *bad* links (for more information about what bad links means, see Section 3.1) is not prohibited

and does not block the student's navigation through the course. However, a large number of *bad* links being followed by the student could indicate that the UM is not capturing the knowledge of the user correctly. Therefore, the question we would like to answer in this evaluation is: is the user's knowledge being captured correctly by the system, also considering that it is changing all the time?

Our approach to evaluate the UM is simple: once we identify that a student has been following *bad* links constantly, we present her a survey to identify her needs and discover the cause of this access via bad links. Depending on the students' answers, we might detect that the UM may not have been updated by the system in an appropriate way. The student could have clicked on a link that is shown as a *neutral* or *good* while the UM still represents it as a bad link. This would be a bug (either in the software or in the link not being made adaptive). The link could also have been *bad* without the student noticing that the link is in fact a *bad* link (and expecting the link to be there and to be recommended). In the former case, the UM status is not being updated correctly. In the latter case, the adaptation design of the course may be flawed. Both cases require the attention of the designer of the course. This topic is part of our experiments presented in Chapter 4.

4. Different students following the same *bad* link

In the previous item we have discussed the case when a single student follows *bad* links constantly. The current topic follows the same spirit but has a subtle (albeit important) difference. Here we want to discuss the case where *many* students follow (the same) *bad* links. Considering that *bad* links are non-recommended concepts, a large number of students following the same non-recommended concept raises the doubt of why this link is not recommended if many students are clicking on it. Again, this could be a problem of the UM or the design of the course. On the one hand, giving another survey to the students would be sufficient to identify the needs of the students and the cause of this situation. On the other hand, it is important for the author of the course to get a feedback on which concepts students are visiting through bad links.

5. Student motivation

In a distance learning course, regardless of whether it is adaptive or not, it is always hard to identify if the students are motivated. The author of an adaptive course needs to know when students lose interest; more important, the author of the course wants to know the cause of this fade of motivation and if the adaptation engine or the course design is making the student feel confused and disoriented. We want to concentrate in two types of behavior: a student who enters a course and spends a long time navigating through the course pages without ever performing the tests, and a student who visits the test pages but does not complete them. The first case can be caused by many factors. For example, the student may simply not know where the test page is or may have not noticed that there exists a test to be performed. Also, the student may not feel prepared to complete the test or may have decided to give up. In order to investigate whether a student is likely to fall into one of these categories, we propose to analyze the total time spent by the student in the whole course and in each concept of the course. If the student has spent a large amount of time navigating through the course and accessing again and again a large number of concepts in a small period of time, then probably the student is motivated to pursue the course, but got lost and does not know what to do. This suggests that the test needs to be relocated to a better place in the course or needs to be highlighted. Such a behavior is also likely to be observed in the case when the student does not feel prepared to complete the test. However, in this latter case, we expect the student to keep visiting all the concepts referring to that topic over and over again.

There are several challenges to be tackled in this chapter. The most challenging problem is the *navigation paths analysis*, because it is a way of enabling future adaptation engines to change the adaptation autonomously. In Chapter 4 we present the findings from the evaluation of an adaptive course where students follow bad links. In Chapter 8 we present a tool (we call *Analysis Plug-in*) through which authors can analyze which concepts are accessed most via bad links. The *analysis plug-in* also presents the time spent on a concept for each student.

6 GALE: Open Generic Extensible Adaptive System

This chapter describes GALE system [22, 2], the Generic Adaptation Language and Engine that came out of the GRAPPLE EU FP7 project. The main focus of this chapter is not only describe GALE, but also present the extensible nature of GALE. Most of this chapter is taken from the PhD thesis of David Smits [22] and associated main publication [2] and is only included here because a good basic understanding of the GALE architecture is needed to understand our extensions (see chapter 8) and the evaluation of extensions made by students (see chapter 7). We also have a brief and comprehensive description of this chapter in [3].

The GALE system was recently introduced as an evolution of the AHA! system [18]. AHA! was originally designed to serve a hypertext course taught through the Web[9]. AHA! aimed at being generic and general purpose and was perhaps best described in the world's first adaptive paper [18] presented at the ACM Hypertext 2006 conference (through an adaptive talk). Knutov et al [15] describe many new adaptation techniques developed to date and provide a list of challenges for creating a new generic adaptive hypermedia system. That research has led to the GAF model (Generic Adaptation Framework) [20, 21], capable of dealing with several types of adaptive systems and applications, from "traditional" adaptive hypermedia to personalized search and recommender systems. GALE tackles this challenge through a very modular and extensible approach described in this chapter. We describe the core functionality of GALE that does not make any assumption of the application area and can best be described as an (empty) adaptation shell. We also describe some standard built-in elements that make GALE directly suitable for popular applications, where there is a hierarchical domain model, content based on some form of XML, and where there is some type of linking for which we want to have adaptation. We then demonstrate how certain types of extensions can and have already been made and used to make GALE more suitable for different types of applications, thus bringing GALE closer to what the GAF model envisions.

This chapter presents in Section 6.1 the global architecture of GALE. To show the generic aspects of GALE we concentrate on three aspects: the processor pipeline for selecting and adapting resources and links, the configuration through

which you can completely alter GALE's behavior, and the event bus, services and extensions through which you can make GALE compatible with different adaptation languages. Section 6.2 presents an application developed by a student to show the complementarity of the conceptual structure and the content of an adaptive application. We also list some other adaptive sites that were created by groups of students, with or without the graphical CAM authoring tool [38] developed in the GRAPPLE project and later renamed to the Course authoring tool. We also show some features of the adaptive PhD thesis written about, for and served by GALE [22].

6.1 The Architecture of GALE

Figure 10 shows the GALE architecture. We will concentrate on the purple part (bottom left) which deals with the adaptation to a resource and some of the blue part (top part of the right side) which deals with internal and external services to handle adaptation formats and user modeling services. If you are reading this chapter sequentially so far and you prefer an example-based description you should read Section 6.2 first. Our description focuses on the aspects of GALE that illustrate its flexibility and its extensibility.

The green part of Figure 10 handles the user logging in and GALE setting up a session and loading some information into a cache for the Domain Model (DM) and User Model (UM). After logging in HTTP requests for concepts can be processed.

Let us first consider the overall process of GALE handling a request from a user (which typically comes in as a request for a URL, sent from the browser through HTTP). We can visually represent this process as shown in Figure 11. The process involves interaction with the Domain Model (DM) of an application (describing the topic domain and adaptation) and the User Model (UM) which holds all the information GALE knows about each individual user. (Note that the use of colors is just for clarity. There is no correspondence between colors in figures 10 and 11.) Throughout this work we will often refer to the configuration of GALE which is done through a file we call `galeconfig`. It allows almost every aspect of the behavior of GALE to be changed to suit the needs of its users.

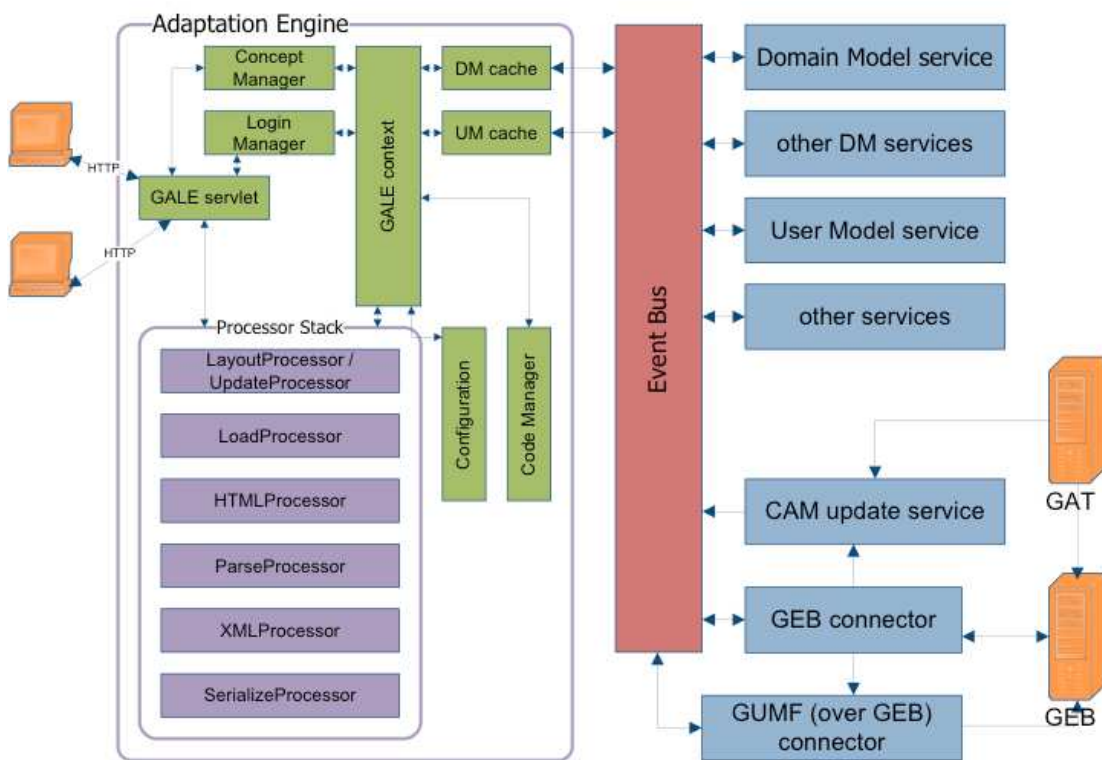


Figure 10: The architecture of GALE [2]

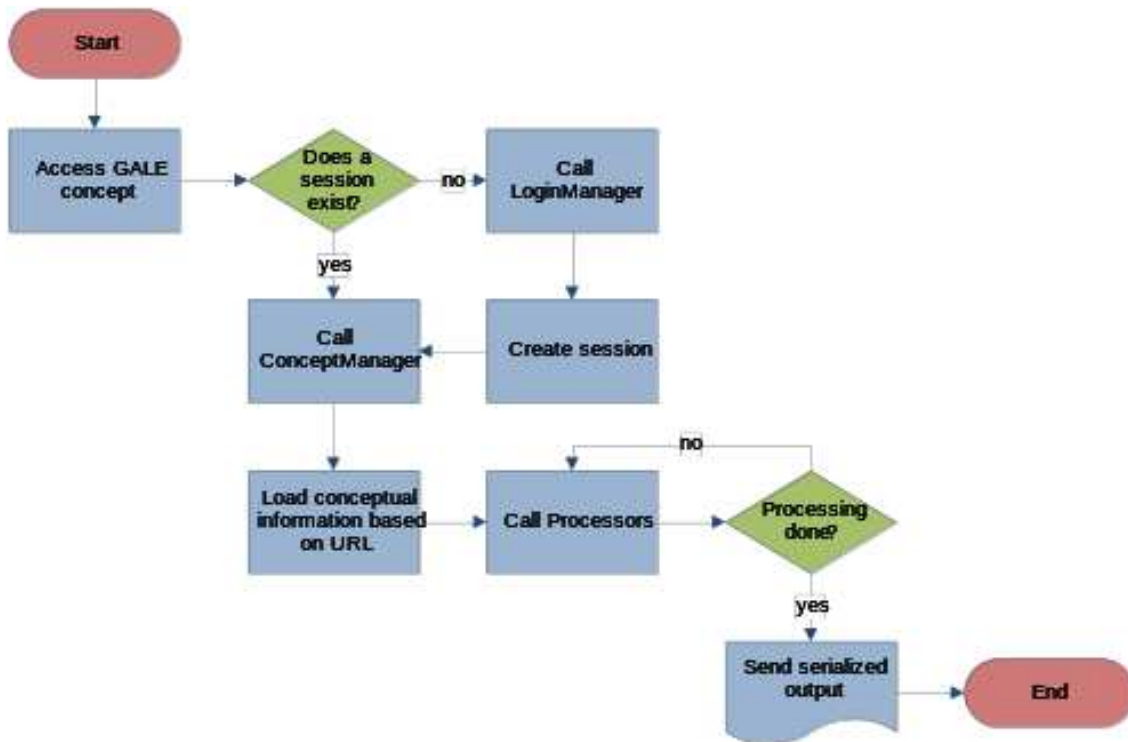


Figure 11: Handling a request for a concept [2]

1. If this is the first request the user sends since starting the browser no session will be associated with that request, so a session is initiated and a login procedure started. This is standard behavior for Web-based applications. GALE can work with several login managers, defined and enabled or disabled in *galeconfig* (see Section 6.1.2) but for stand-alone GALE use the login follows the following multi-step procedure:
 - (a) For a first request (without session information as there is no session yet) the user is still unknown. The login manager redirects to a servlet/page that prompts the user for a user id and password.
 - (b) The user id is passed on to the UM cache, to request the application-independent part of UM for this user. Internally GALE refers to this as the user entity.
 - (c) Since the UM cache will not have cached the user model yet, it will communicate with the user model service through the event bus.

The event bus makes it possible to distribute the processing load for adaptation and user modeling between different machines.

- (d) UM is needed by the login manager to verify that the user has provided the correct password. Next the login manager (servlet) returns a redirect to the original URL. As a result the user's browser will request the same concept again, this time with session information. (This repetition of the request is invisible to the end-user.)
2. GaleServlet now calls the concept manager in order to find out how to handle the request. If the request is for a concept, the concept manager will determine the identity of the requested concept and retrieve the domain and adaptation data for the concept from the DM cache which may need to load it from a Domain Model service. (Here again it is possible to offload the DM service to a different machine.) If the URL does not refer to a concept it is handled differently (e.g. a file can be simply retrieved and served, as in the case of an image).
3. Handling the concept is a multi-step sub-process that uses processors (more or less in the top to bottom order as shown in the purple part of Figure 10). GALE can be extended with new processors that can be used anywhere in the processing pipeline. (We describe these extensions in Section 6.1.3.) Section 6.1.1 describes the default processor pipeline (the default configuration in galeconfig). The processor pipeline will typically load a resource, process it in memory, and at the end serialize it again and send the result to the user's browser as an HTTP response.

6.1.1 The GALE Processor Pipeline

The processing of a concept (request) is done through a series of processors (each one handling the output of the previous one). Galeconfig defines the list of processors and thus also the order in which processors are invoked. One can easily insert a processor into the processor chain (as shown in Section 6.1.3) to add some functionality to GALE.

GALE uses its own language GAM to define an application's domain and adaptation model (DM). Section 6 shows some of that code. Here we use a subset

of GAM which we refer to as GALE EL code (EL stands for Expression Language). This is code to either evaluate an expression over information from DM and UM or to update values in UM. UM is the main information source for the personalization (or adaptation) offered by GALE. We sometimes use a property of a concept, which is a DM element, and sometimes an attribute, which is a UM element. The GALE EL code is in fact Java code in which some shorthand notation is used to refer to concepts, properties and attributes. The code manager is configured in *galeconfig*. In principle it is possible to write your own code manager to enable GALE to work with a different expression language, perhaps not based on Java but on another programming language. Section 6 gives several examples of GALE EL code. Here we just give a bare minimum introduction:

- $\{\#\text{suitability}\}$ refers to the value of the suitability attribute of the current concept ($\#$ always refers to a attribute).
- $\{\#\text{image?title}\}$ refers to the title property of the image attribute of the current concept ($?$ always refers to a property, whether a property of a concept or of an attribute). While using GALE the value of an attribute may change but properties of an attribute cannot change. (The properties are defined in the DM.)
- $\{->(\text{parent})?\text{type}\}$ refers to the type property of the parent of the current concept (following the parent relation). Again, this is defined in the DM: properties and relations cannot be changed on the fly.
- $\{\#\text{visited}\} = \{\#\text{visited}\}+1$; is a statement that assigns the result of $\{\#\text{visited}\}+1$ to the visited attribute of the current concept. It thus simply increments the visited attribute by 1.

Note that the properties, attributes and relations used to illustrate the syntax are not predefined in GALE. The “beauty” of GALE is that it only defines that there can be concepts with properties, attributes (possibly also with properties) and relations and does not say which concepts, properties, attributes or relations should exist. However, to make GALE more easily usable from the start there are some “additions” that are included in the GALE distribution that do make such

assumptions. For instance there is a navigation accordion menu available (static-tree-view) that presents part of a concept hierarchy, assuming that parent relations exist to define that hierarchy. Also, the *LoadProcessor* we describe below expects the existence of a resource attribute in order to decide which file to load.

We now explain the processors of the pipeline, mostly in the order they are called (which is the top to bottom order in Figure 10).

1. The first processor that is called is the *UpdateProcessor*. It signals an *EventManager* that the *access concept* event has occurred. The default *EventAccessHandler* executes the event code of the concept as defined in DM. Typically this event code will signal some UM update to the UM cache. For instance, the number of visits to the concept may be incremented. UM cache communicates over the event bus with the UM service. Event code associated with UM attributes may prompt the UM service to generate more UM updates. The visit to a concept may signal increasing interest in that topic and related topics or may indicate an increase in knowledge of the concept. The resulting changes to UM are posted on the event bus, and as a result are integrated in the UM cache. Important to note here is that 1) UM updates are calculated before generating adaptation (a design choice that corresponds to previous behavior of AHA! and motivated and explained in [18]), and 2) UM updates come from both the *UpdateProcessor* (in the adaptation engine) and from (event code) rules executed within the UM service. This distributed execution of UM updates is essential in GALE as GALE can also be used together with external UM services, including the GRAPPLE user Model Framework GUMF [70]. UM updates resulting from a concept access are handled synchronously (the adaptation engine waits for a response from the UM service). However, should an event access result in a message being sent to an external UM service such as GUMF it is handled asynchronously (so the adaptation engine will not wait for a response but will handle the UM update whenever it comes in). Figure 12 shows the process of UM updates in a more graphical way.
2. After the UM updates have been performed the *LoadProcessor* will retrieve the actual resource (file) associated with the concept. The name of the resource is found in the resource attribute of the concept. This “name” may refer to a

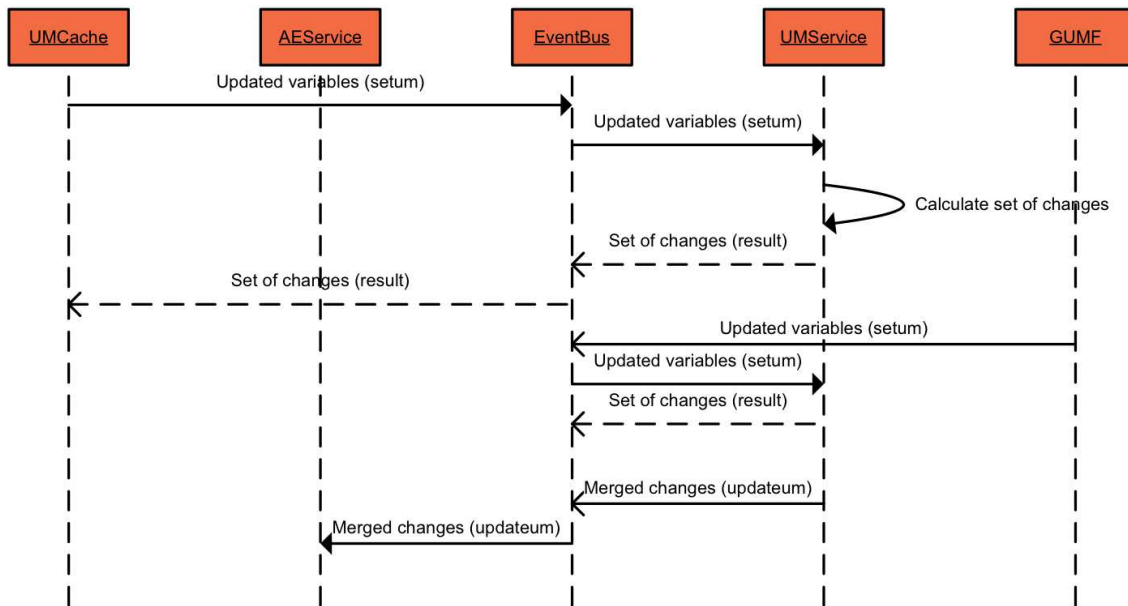


Figure 12: User model updates being handled by the adaptation engine and UM service(s) [3]

local file or a file to be retrieved from some server using http. It may also be a (GALE EL code) expression over DM and UM to “compute” the name of the actual resource. An *InputStream* is opened so that a subsequent processor can load and process the data. File name extensions are used to determine the mime type of the resource.

3. Optionally the *LogProcessor* then adds an entry to a global log file (*access.log* by default). The id of the user, date, request, referrer (that may be present in the HTTP request), the name of the requested concept and the resulting resource are logged, for possible later analysis. Note that in order to log the resource that is retrieved and adapted the *UpdateProcessor* must run before the *LogProcessor*. An implication of the chosen order is that the *LogProcessor* cannot log any user model state information that existed before the *UpdateProcessor* performed its updates. In Section 6.1.3 we show how to extend GALE to perform such additional logging if desired.

4. If the mime type of the resource is HTML (not XHTML) the *HTMLProcessor* uses the (open source) *Tagsoup*⁵ converter to convert the file to XHTML. The new *InputStream* now contains valid XHTML.
5. If the input (after step 4) is XML (also XHTML) the *ParseProcessor* converts it into an in-memory DOM tree, using the open source *dom4j*⁶ parser. We will not describe the processing of non-XML input further, but in *galeconfig* it is possible to insert processors for other data types if desired. (Such processors of course need to be developed as the GALE distribution only contains a processor for XML.)
6. The *XMLProcessor* walks through the DOM tree in order to perform adaptation where needed. The modules that may be used to perform adaptation to certain tags are loaded by the *XMLProcessor*. Which modules exist and to which XML tags they are applied is configured in *galeconfig*. Modules are provided to handle “if” tags, “object” tags, links, variables, and more. We explain a few modules in next section. Adding new modules to handle different tags, possibly in different XML formats, is relatively easy. Within the GRAPPLE project the addition of modules has been investigated for device adaptation and for adaptation to virtual reality [71].
7. Optionally, the *LayoutProcessor* generates a frame-like structure using tables, by creating an (in-memory) XML document that contains the views (any class that implements the *LayoutView* interface) embedded in a table that defines the layout. This XML document has a placeholder element where the actual content should be. A *CSSLayoutProcessor* is also available that uses *css* and *div* sections to lay out the browser screen. A *FrameLayoutProcessor* can also be used that uses an *iframe* element to display the actual content. The different processor choices make it possible to combine GALE with many other presentation structures (e.g. with portals or learning management systems). For simplicity we do not describe the differences between these processors in detail, nor their actual position in the pipeline (which also differs).

⁵See ccil.org/cowan/XML/tagsoup/ for more information on tagsoup.

⁶See dom4j.sourceforge.net for more information on dom4j.

8. When the DOM tree is adapted the *SerializeProcessor* generates the textual XML representation and presents that to *GaleServlet* as an *InputStream*. *GaleServlet* uses this to generate the HTTP response to be sent back to the user's browser.

For resource types that do not have specific processors associated with them, *GaleServlet* will create the final *InputStream* itself in order to send the content back to the browser as an HTTP response. This happens for instance with images embedded in HTML pages. It is also possible to “bypass” the processor pipeline to perform completely different processing and generate a result “page”. This is done by means of a *Plug-In*. Examples of predefined plug-ins are the *Form* plug-in that allows input from forms to be used to execute UM updates, the *Password* and the *Logout* plug-in with obvious functionality. A typical “extension” of basic GALE functionality is a plug-in to evaluate multiple-choice tests for use in educational applications. (This is not a real extension as it is already included in the GALE distribution.)

6.1.2 GALE Flexibility through Configuration

GALE was developed as part of the GRAPPLE EU project. GRAPPLE aimed at integrating open source or commercial Learning Management Systems (LMS) with Adaptive Learning Environments (ALE). The project aimed at making adaptivity available to a broad audience (mainly in technology-enhanced learning). To this end the adaptation engine needed to be made very generic and extensible, because developers of adaptive applications are expected to have special desires for adaptation functionality we might not foresee. Within the project itself adaptation in simulation and adaptation in virtual reality were already considered, as well as device adaptation. Different LMS may require a different way of embedding GALE output in their presentation. The different layout processors make this possible. To make GRAPPLE usable with many different LMS and to support life-long learning a repository of user-specific information was needed (the GRAPPLE User Modeling Framework, GUMF [70]). Different LMS store information in GUMF, and Section 6.1 already explained how input from such an external UM service is

captured by GALE. The integration process between LMS and GALE is described more in detail for instance in [72].

GALE uses the Spring⁷ “inversion of control” container to configure and instantiate all components. Without going too much in detail we describe the most important configuration elements here (omitting small things like where the access log file or some other files are stored or at which address GEB and GUMF are located). The GALE configuration is stored in the *galeconfig* file *galeconfig.xml* (in Tomcat’s *webapps/gale/WEB-INF* directory). What the configuration mostly does is associate names that have meaning as functional parts of GALE with a Java class name or Bean (implementing the functionality) and it also defines which properties configure the behavior of that functional part. We look at some parts in detail below.

1. The *processorList* defines processors that handle a request and adapt resources. The main processors have been mentioned in Section 6.1 already. Here we look specifically at two processors: *XMLProcessor* and *PluginProcessor*.
 - (a) The *XMLProcessor* performs transformations (for adaptation) to the DOM tree of an XML resource. *Galeconfig* contains a list of Modules that handle specific XML tags. Adding adaptation to a new tag can be done by creating a new module and associating it with the tag (and its namespace) in this list. We only present a small selection from the modules that are included in the GALE distribution:
 - The *IfModule* handles the `<if>` tag. It expects `<if>` to have an argument “*expr*” that is a Boolean expression in GALE code. It expects one or two child elements: a `<then>` and optionally an `<else>` element. The module replaces the `<if>` subtree by either the content of the `<then>` subtree or the `<else>` subtree. The *IfModule* thus realizes what is known as the *conditional inclusion of fragments* technique [15].
 - The *AdaptLinkModule* handles the `<a>` tag which is used just like the HTML `<a>` tag, but referring to a concept, not a page or

⁷See www.springsource.org for more details about Spring.

resource. GALE (actually the *LoadProcessor*) decides which page to retrieve and return based on the resource attribute of the concept. An optional `exec` argument can be used to associate a (UM) action with following the link. For instance: `exec="{tour#start}=true;"` could be used to say that following this link indicates the start of a tour, whereas accessing the same concept (tour) through other links does not. GALE supports the adaptive link annotation (and hiding) techniques through link anchor colors and through icons in front of and behind the link anchors. While this can be defined in *galeconfig* (for all adaptive applications running on the same server at once) one can define for each concept an attribute `#link.classexpr` to conditionally define a link class (corresponding to a link color in the application's stylesheet) and an attribute `#link.iconlist` that determines under which conditions (over DM and UM) which icons are placed. Any GALE EL expression can be used to conditionally generate the filename (URL) of an icon (or the name of the link class). Typical use of icons can be found in e.g. ELM-ART [13] and Interbook [8], using colored balls to indicate the suitability of a link and checkmarks to indicate the knowledge level of the link target concept. Typical use of link colors uses blue for recommended links, purple for suitable but already visited links and black for non-recommended links that appear hidden amongs black running text. But any combination of colors is possible and e.g. the adaptive PhD thesis [22] also uses orange and dark purple. The choice is completely up to the application designer.

- The *ObjectModule* inserts either a file specified through the `data` argument or a concept specified through the `name` argument. The `<object>` tag is preferred over the `<if>` tag when the same fragment needs to be conditionally included in many different pages. With `data="header.xhtml"` we can insert a header file in a page, whereas with `name="programming"` we insert whichever resource is associated with the concept programming (possibly involving evaluating expressions to select that resource).

- The *VariableModule* inserts either the value of a UM attribute or the result of a GALE expression in the page. `<variable name="#visited">` for instance can show the number of visits to the current concept; `<variable expr=${#visited}>` does the same, but now as expression.
- The *AttrVariableModule* is similar but inserts its result in the surrounding element's tag. `<attr-variable name="src" expr="${?image}">` uses the value of the image property of the current concept as the source (url) of an image to insert. Note that because of syntax restrictions of the XML language we could not use `|variable|` inside the `` tag itself. (Arguments of an XML element cannot contain XML elements.)
- The *ForModule* repeats a fragment of XML for elements in a list.

```

<for var='concept' expr='${-(parent)}'>
  <variable expr='${%concept?title}' /><br/>
</for>

```

This example inserts a list of the titles of the children of the current concept.

- The *PluginModule* generates a link to a plug-in. (It does not perform the plug-in code itself as this is done by the *PluginProcessor* described below. `<plugin name="logout">Logout<plugin>` results in a link through which the user logs out. Note that because GALE transforms `<plugin>` into a link anchor the description of the plugin (the word "Logout" in the example) must not contain an `<a>` tag but can contain other HTML tags if desired.
- The *MCMModule* is a nice example of how the "core" functionality of GALE can be extended to serve specific types of applications. This module (which we associated with a `<test>` tag in our applications) uses subelements that are called `<question>` and `<answer>` (with different parameters and content) to generate a multiple-choice test. We used this in a setup where we wrote more questions and answers

than are needed and the module selects a random subset of questions and answers (making sure to include correct answers of course). As the multiple-choice test is generated by a module it is embedded in an (XHTML) page that can contain any other content (and links) as desired. Also, the content of each question and answer can contain arbitrary tags as well.

- (b) The *PluginProcessor* handles plug-ins that perform some function and then generate “complete” output. Examples of plug-ins are the password and logout plug-ins (to change the user’s password and to log out) and *exec* to execute some GALE code and show the result (used mostly for debugging purposes). To facilitate educational applications we added the *mc* plugin to evaluate a multiple-choice test generated by the *MCMModule*.
2. The *hibernateDataSource* bean controls how GALE stores data (including DM and UM). As we use Hibernate⁸ GALE is independent of the database backend used. We have used two storage methods so far: *hsqldb* which stores data in a simple text format and *mysqldb* to store data in a MySQL database. MySQL (or any other real database like Postgres or Oracle) is recommended for a real server installation. In that case it is possible to use a separate machine for the database to offload the GALE server.
 3. Figure 10 shows that GALE uses an event bus through which the core GALE engine communicates with DM and UM and other services. Two implementations of this bus exist: one (*LocalFactory*) that uses method calls and thus requires all services to reside on the same server and one (*SOAPFactory*) that uses SOAP and can handle DM and UM services that run on different machines. The event bus is configured to communicate with a number of services. (This bus uses the Publish/Subscribe method.) The DM and UM services are most obviously needed, but several other services exist to implement compatibility of GALE with different authoring formats and tools for adaptive application, as described in the next section.

⁸See www.hibernate.org for more details on this Java persistence framework.

4. (13) The *loginManager* bean defines how GALE users can identify themselves. The *DefaultLoginManager* presents a simple form for username and password (explained in Section 6.1). The *LinkLoginManager* allows LMS users to automatically log in on GALE using the id they have on the LMS. The LMS uses a “secret” key to identify itself to GALE and essentially “promises” that the user is who (s)he says (s)he is. The *IdPLoginManager* makes use of Shibboleth⁹ to make GALE usable in a federation of institutes/companies that share a single sign-on facility. The *OpenIdLoginHandler* allows the reuse of an already existing id in any service using Open ID¹⁰.
5. The *codeManager* defines which language of adaptation code is used. All examples of GALE code used in this paper use GAM: Java with some shorthand to refer to DM and UM values, but it is possible to use very different code provided that a new code manager is developed. Note that in GALE there are actually two code managers (slightly different): one in the adaptation engine (dealing with concept event code and GALE expressions used in GALE XML tags) and one in the UM service (dealing with UM attribute event code).
6. (15) The *configManager* is a wrapper for handlers of different types of configuration: for processors, for link adaptation (with icons) and for presentation. The latter (*Presentation-Config*) defines which automatically generated parts can exist in the presentation of adapted pages. These parts are called views. In order to create a view it may be necessary to make some assumptions about the presence of specific properties or attributes. The GALE distribution comes with some predefined views: the *static-tree-view* which displays a menu over the hierarchy of concepts of the current application, based on the assumption that there are *parent* relations to define that hierarchy (see Figure 17), the *next-view* which generates a link to the next concept in a guided tour (see Figure 16), and the *file-view* which inserts (and adapts) a file with a fixed name. The file-view can be used to include a header and/or footer for instance (see Figures 16 and 17) instead of including an `<object>` tag.

⁹Consult the Shibboleth Consortium site at shibboleth.net for more details about Shibboleth.

¹⁰Consult the Open ID Foundation site at openid.net for more details about Open ID.

6.1.3 Extending GALE: from generic towards general-purpose adaptation engine

As we already showed in the previous (sub)section the generic core of GALE needs to be extended in order to obtain the desired functionality for a specific type of application. GALE can be extended in extreme ways, like by adding a new code manager. In this way, GALE was created to be a generic adaptation engine in the sense that its functionalities allow authors to use (or at least emulate) a great variety of adaptive techniques. Also, having a generic core, GALE can also be extended to be used for general purposes, i.e., in many different adaptive applications. The more common (and easier) ways to extend GALE include the following types of extensions:

1. Adding modules: The standard GALE distribution comes with modules for conditional inclusion of fragments (the *IfModule* and *ObjectModule*), for link adaptation (the *AdaptLinkModule*), for presenting information calculated from the DM and UM (the *VariableModule* and *AttrVariableModule*), for generating lists (the *ForModule*) and counting DM or UM elements (the *CountModule*) and for including more complex content parts generated from DM and UM by means of views (the *ViewModule*). Specific applications may require adapting different XML tags in specific ways. We already showed the existence of the *MCMModule* for multiple choice tests. GALE also has a *TextModule* for handling input in the Creole¹¹ syntax used in Wikis. In the GRAPPLE project modules were added to handle virtual reality input described in the X3D language (see [71] for details).
2. Adding views: Views generate content from DM and UM or any other source and are included in the presentation through the <view> tag (or any other tag associated with the *ViewModule*). Figure 17 shows the *static-tree-view* that produces an accordion menu over the hierarchy formed by parent relations. The link colors and colored balls are chosen based on the value of a suitability attribute in combination with a visited attribute. Students have experimented with different variations of the *static-tree-view* including one that shows a progress bar for the top n (configurable) levels in the concept hierarchy.

¹¹See wikicreole.org for more details on the WikiCreole language.

3. Adding plugins: Instead of just “adding” or “altering” something in a given resource you sometimes may wish to generate a complete presentation, e.g. as a response to completing a form (for which GALE offers the *FormPlugin*). A special case is the *MCPlugin* for multiple-choice quizzes which returns a score, and the *ExecPlugin* which executes arbitrary GALE EL code. When creating a new plug-in you should be aware that after the *PluginProcessor* has run all subsequent processors are skipped. (This guarantees that GALE cannot mess up what the plug-in has generated.)
4. Adding processors: A nice example of a reason to add a new processor was found in the research described in [49] and presented in Chapter 3. In this work the effect of link annotation is researched by checking whether users follow the advice given through link annotation or not. To study this in the hypermedia course taught at the TU/e the user sessions needed to be replayed (from the available log) because the annotation state of links (when followed) was not logged. To avoid this we wished to add this state to the information being logged by GALE. However, this state is computed from the user model, and the *UpdateProcessor* changes the basis for this computation before the *LogProcessor* runs. The solution was to add a processor that evaluates the annotation state and stores it in a user model attribute not touched later by the *UpdateProcessor*. The *LogProcessor* can then log that state for later analysis by researchers. Indeed, we used this data in the evaluation presented in Chapter 4.
5. Adding services: Before GALE numerous other adaptive systems have been developed. It would be a shame not to be able to reuse them in GALE, given that GALE is very generic and should be able to emulate (or be extended in order to emulate) the behavior of these other systems. Such emulation was already performed before in [41] where the AHA! system [18] was used to emulate Interbook [8]. To import applications in other formats one can develop new DM services to be placed on the event bus. The main approach here is that the new service should only be used to retrieve the domain model (conceptual structure) and the adaptation rules for an application, not the content (pages) of an application. The GALE distribution includes services

to accept applications defined in the CAM or Course Editor [38] developed in the GRAPPLE EU project and also applications in the format used by AHA! (version 3). Besides the GAM format GALE also supports an XML representation of its applications: the GDOM format. There is a service to read GDOM definitions and an export plugin that can generate a GDOM representation for any application (independent of which format was used to import the application). It is possible to develop compilers (to use in a service) for formats like Interbook, Scorm, and others, but this has not yet been done until now.

6.2 Examples of GALE Applications

This section presents an example-based introduction to (and motivation for) GALE. It reports on applications that were realized in GALE and on student projects to create applications and GALE extensions. In Chapter 7 we evaluate part of the applications developed by the students and presented in this chapter. An important remark is that this section can be read without first reading the previous section.

We will not describe how to define an overall adaptation strategy or how to design the adaptation based on pedagogical relationship types such as prerequisites. This is the subject of GRAPPLE's (CAM or Course) authoring tool described in [38].

GALE supports two approaches towards creating the content of an application, represented in Figures 13 and 14 (taken from [73]).

Figure 13 shows an author writing complete pages. This is a viable option for authors who wish to create adaptive applications without writing (or even seeing) any GALE code. This approach is well suited for an application in which pages do not have a common structure. The hypermedia course 2ID65 at the TU/e was originally created in this way and the GRAPPLE tutorial (at <http://gale.win.tue.nl/>) as well. This is also a good approach for making existing applications adaptive, for instance applications generated from databases (possibly from a content management system or a wiki).

Figure 14 shows an example where a cluster of pages share the same structure. (A and B are alike, X and Y are also alike.) When pages are created separately one

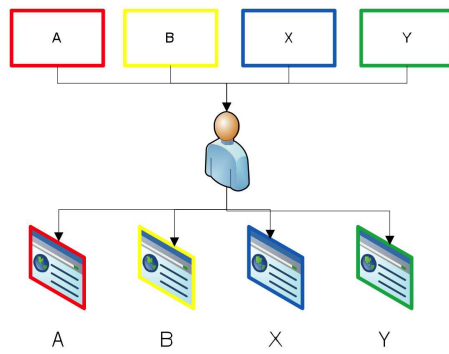


Figure 13: Creating pages separately [2]

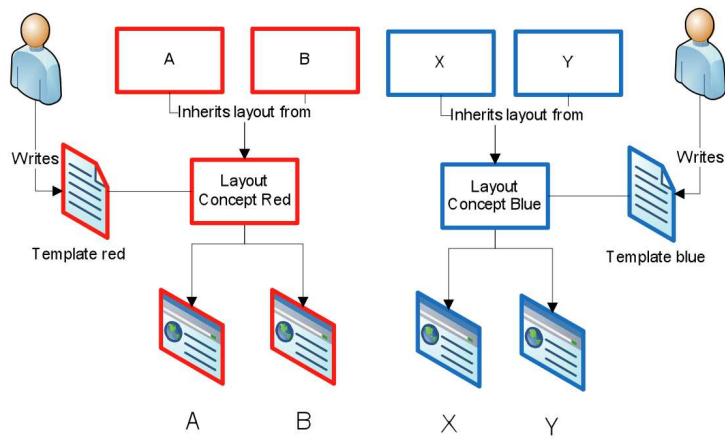


Figure 14: Authoring through template pages [2]

needs to be careful to replicate the style on all pages that should be similar, and changing the style involves changing all these pages individually. Therefore Figure 14 shows the use of one template page for each set of pages that should look alike. The templates are used to (virtually) create individual pages by indicating which bits to use where in the template. Changing the presentation of all pages that use the same template requires a single file to be edited. Figure 15 shows the presentation of a concept from the application Milkyway that was created by a master student at the TU/e [73] and that has frequently been used to illustrate GALE. All pages of Milkyway have the same look and feel, not only with the header and footer (and a navigation menu that we cut off) but also within the main content part of the page. Milkyway deals with stars, planets and moons, and has slightly different templates for each of these types of celestial bodies. Below we explain the template used for presenting planets.


- Milkyway
- Nebula
- Star
- Planet
 - Mercury
 - Venus
 - Earth
 - Mars
 - Jupiter
 - Saturn
 - Uranus
 - Neptune
- Moon

Paul De Bra (debra@win.tue.nl) has read 7 pages and still has 22 to read - [list of read pages](#) - [pages still to be read](#)
Options in stand-alone mode: [change password](#) [logout](#)

Jupiter

Is Planet of: [Sun](#)

Image of Jupiter



Information

Jupiter is the fifth planet from the Sun and the largest planet within the Solar System.[10] It is two and a half times as massive as all of the other planets in our Solar System combined. Jupiter is classified as a gas giant, along with Saturn, Uranus and Neptune. Together, these four planets are sometimes referred to as the Jovian planets.

The following [Moon\(s\)](#) rotate around Jupiter:

- [Callisto](#)
- [Ganymede](#)
- [Io](#)
- [Europa](#)

Next suggested concept to study: [Saturn](#)

Figure 15: Example of a page based on a template.

In a GALE application every *concept* can have an associated *layout*. In Figure 15 the header and footer are defined as part of that layout. The “main” part of the presentation is the page. The layout can be shared between concepts but can also be different for each concept. It can also be adaptively changed but we do not expect this to be common. There are different ways to realize this layout (within HTML): using frames, tables, iframes and simple “divisions” that are positioned based on a CSS style sheet. All possibilities are supported by GALE through different *layout processors*. You can see that the presentation in Figure 15 consists of multiple parts because the scroll bar does not cover the header and footer. In the figure you can see that the example page contains a *title* (Jupiter), a *typology* (Is Planet of: Sun), an *image with title*, an *information fragment* and a *list of links to children* (The following Moon(s) rotate around Jupiter.) Below we show how each part is defined, to give you a feel for creating pages containing GALE tags and GALE code. This shows that pages can be partly “constructed” out of DM and UM elements rather than being purely “authored”. All DM properties that are used (like “Jupiter”, “Sun”) and relations (like “*parent*” and “*isPlanetOf*”) have been created using the graphical GRAPPLE authoring tools (the Domain editor to be exact).

- In the header we see the name and email of the user. The code `<variable name="gale://gale.tue.nl/personal#name">` and `<variable name="gale://gale.tue.nl/personal#email"/>` extracts that information from UM and inserts it in the page. (The “gale:” pseudo-protocol written in what looks like a URL tells gale to retrieve information from the “GALE Context” and is easier to understand than the complete Java method call to retrieve the name or email. That method call could of course also have been written here as GALE code syntax is Java with shorthand to retrieve DM and UM information.)
- The title “Jupiter” is generated by means of the code `<variable expr="{?title}"/>` which inserts the title property of the current concept in the page. As shown in Section 6.1 a # sign refers to a user model attribute and the ? sign refers to a domain model property.
- The “Is Planet of: Sun” part is a bit more complex. “Planet” is actually the title of the parent concept and “Sun” is the title of the

concept to which Jupiter has an *isPlanetOf* relation. `<a><attr-variable name="href" expr="$->(parent)?title"/>` generates the link anchor tag. Calculating the value for the “href” argument cannot be done inside the `<a>` tag because XML tags are not allowed to appear inside another XML tag so it is done using `<attr-variable>`. This is similar to `<variable>` but inserts its output into the parent xml tag, in this case the `<a>` tag. The anchor text is `<variable expr="$->(parent)?title"/>` (which generates the word “Planet”). The link to “Sun” uses `<a><attr-variable name="href" expr="$->(isPlanetOf)?title"/> <variable expr="$->(isPlanetOf)?title"/>`. Like with “Planet” the same expression appears twice because “Sun” is the anchor text as well as the link destination.

- To insert the image (we ignore the title here) we use the code `<attr-variable name="src" expr="$ {?image}"/>` where the name (URL) of the image is part of the domain model (created with the graphical Domain editor). As you see the attr-variable tag can be used in combination with any other tag to specify an argument used in that other tag (the parent tag in the XML document).

- The information fragment is stored in a separate file and included in almost the same way as the image:

```
<object><attr-variable name="data" expr="$ {?info}"/></object>.
```

The name of the file is part of the domain model.

- The list of moons is generated using the `<for>` tag. (*Milkyway* additionally checks whether a planet has moons to avoid generating an empty list if it doesn't.) We only show the code for the list (not for “The following moons rotate?”)

```
<ul><for var="concept" expr="$<-(isMoonOf)"/>
  <li><a><variable expr="$ {%concept?title}"/>
    <attr-variable name="href" expr="&quot;%concept&quot;"/></a></li>
</for></ul>
```

All concepts with an *isMoonOf* relation to the current concept are associated, one by one, with the variable “concept”. A bullet list is generated with for

each such concept a list item with a link to the concept and with the title of the concept as link anchor text.

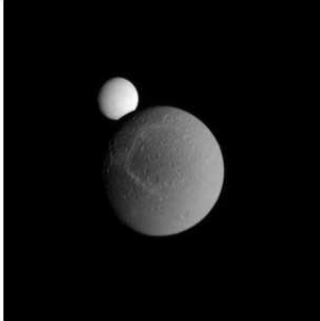
The Milkyway application has a “Course model” in which *prerequisites* are defined. For instance, a planet should be studied before studying its moons becomes recommended. Figure 16 shows the top part of the page about the concept “Moon”. This time we have not cut off the navigation menu. In this menu we see link annotation (with link colors and with colored balls) to indicate that only the four moons of Jupiter are recommended (as we only studied the planet Jupiter and no other planet that has moons).

- Milkyway
- Nebula
- Star
- Planet
- Moon
 - Moon_Earth
 - Phobos
 - Deimos
 - Europa
 - Io
 - Ganymede
 - Callisto
 - Titan
 - Miranda
 - Ariel
 - Umbriel
 - Titania
 - Oberon
 - Triton

Paul De Bra (debra@win.tue.nl) has read 7 pages and still has 22 to read - [list of read pages](#) - [pages still to be read](#)
Options in stand-alone mode: [change password](#) [logout](#)

Moon

Image of Moon



Information

A natural satellite or moon is a celestial body that orbits a planet or smaller body, which is called the primary. Technically, the term natural satellite could refer to a planet orbiting a star, or a dwarf galaxy orbiting a major galaxy, but it is normally synonymous with moon and used to identify non-artificial satellites of planets, dwarf planets, and minor planets. **Visited:** 9

Next suggested concept to study: [Europa](#)

Figure 16: Link adaptation based on prerequisites.

The link annotation shown in Figure 16 uses the same “good”, “neutral” and “bad” link classes (blue, purple, black) that were previously used in some AHA!

applications, as we can see in Chapter 3 and 4. In GALE (and in fact already in AHA! as well) you can define arbitrarily many link classes and arbitrary conditions for deciding which class a link should have. The default class and color scheme are suitable for the default adaptation (or pedagogical) models created using the GRAPPLE authoring tools, described in [38] for instance. The graphical authoring tools hide the complexity of the GALE adaptation rules from authors, but they do enable advanced authors to create their own arbitrarily complex rules, specified in the Generic Adaptation Model language (GAM).

In April 2012 David Smits defended the world’s first adaptive PhD thesis, about GALE, and served by GALE [22]. The domain model and adaptation rules for this thesis were written directly in GAM and the layout was defined using CSS (cascading style sheets). Figure 17 shows a screen shot from the thesis. It illustrates some interesting style differences between the thesis and the *Milkyway* application and shows some extensions to the core GALE functionality.

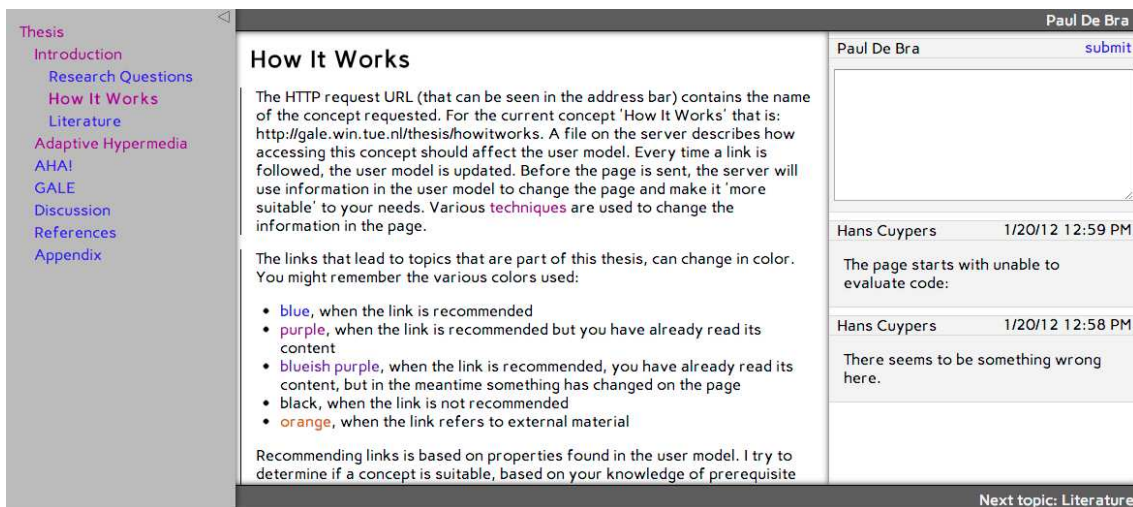


Figure 17: Screenshot of the adaptive PhD thesis about and served by GALE.

This “How it works” page from the thesis explains how the link adaptation makes use of five different link colors with a different meaning. It shows vertical lines to indicate the parts of the current page that are different from what was shown on this page during the previous visit. (So there not only is content adaptation but also some help to the end user to see which parts have been adapted.) The screenshot also shows an optional comments field where the reader can enter comments and

see all users' comments on this page. Only the author of the application can delete comments. There is some adaptable functionality as well: the menu on the left can be closed (by clicking on \triangleleft), the comments part can be hidden, some stereotype user profile can be selected ranging from layperson (e.g. family reading the thesis) to expert (e.g. review committee member), and the profile can also be reset to start over and try different adaptation. The thesis also has layout that automatically adapts to the device used by the reader, for instance moving the menu to the top when a narrow-screen smartphone is used.

The above description shows that it is possible to create very different types of application and presentation in GALE, making it truly “general purpose” as an adaptive Web-based hypermedia platform. At the TU/e some groups of students developed their own GALE applications, using any content they desired, to see whether the flexibility and extensibility of GALE can really be used by sufficiently technically skilled authors. All students were given the *Milkyway* application [73] and the *adaptive PhD thesis* [22] as examples. In 2011 most groups used the graphical GRAPPLE authoring tools [38] for defining the adaptation. In 2012 most groups used the GAM language directly, after receiving feedback from the 2011 students about the different authoring approaches. Chapter 7 presents an evaluation of these GALE applications developed by the group of students from the year of 2012.

Below we list the subjects of the applications the students developed, and also whether the graphical GAT authoring tool was used or the GAM language, whether page templates were used (instead of writing each page individually and completely) and whether a GALE layout manager was used or layout was done using CSS. In the table below the first 10 projects were from 2011, the latter from 2012.

All of these examples used adaptive navigation support through *prerequisite* relationships, some added *conditional inclusion of fragments*, some *stretchtext*, some used template pages, some used content from a wiki, some required specific handling of forms input (e.g. the restaurant menu and radio player). Contrary to what one might expect students preferred to use the GAM language over the graphical authoring tools. This may be explained by the type of students (computer science) who valued full control over ease of use of the authoring interface.

In 2012 master students created their own GALE extensions in order to see whether the claimed extensibility of GALE is actually achievable and usable in practice. We present an evaluation of these GALE extensions in Chapter 7.

App. Subject	Authoring Tool (GAT vs GAM)	Template pages used? (yes/no)	Layout (GALE/CSS)
Photography	GAT	yes	GALE Layout
Logic	GAT	yes	GALE Layout
Biomes of the world	GAM	no	CSS
Disney Virtual Tour	GAM	yes	CSS
Restaurant Menu	GAT	yes	CSS
Web Technology intro	GAT	yes	GALE Layout
Movies	GAM	yes	CSS
The “happy” Butcher	GAT	yes	CSS
TU/e Information System	GAT	yes	GALE Layout
HTML/CSS intro	GAT	yes	GALE Layout
Research Methods book	GAM	no	CSS
Relativity Theory	GAM	no	CSS
Elementary Electro Physics	GAM	no	CSS
C-programming	GAM	no	CSS
Radio Player	GAM	no	CSS
Food Recipes	GAM	no	CSS
Latex	GAM	no	frames
Currency Trading	GAM	yes	CSS
Linux Commands	GAM	no	CSS
Do-It-Yourself Tools	GAM	no	CSS
Sports Center	GAM	no	CSS

Table 7: Adaptive applications developed by (groups of students)

7 GALE Extensibility Evaluation: A Qualitative Approach

The flexibility of the AHA! system [18] was already shown through a few applications. For example, a master student at Virginia Tech used AHA! to create an adaptive alcohol abuse tutorial [74], with a look and feel that was completely different from the hypermedia course for which AHA! was originally designed. Another good example of this flexibility was used in [41] where AHA! was used to emulate the look and feel and adaptation of Interbook [8].

7.1 Introduction

GALE's flexibility (of configuration in *galeconfig* and the ability to add new Java classes) was a desire of the developers of adaptive applications, because they wished not only to be able to express their own adaptation strategies and adaptation effect but also to add new functionality like adding adaptation to Virtual Reality [71], adding visualization of student progress or extending logging to perform more in-depth analysis of student behavior.

Developers can thus extend GALE functionality by extending the generic core of GALE [3]. It can be done for any specific type of application, as we presented in Chapter 6. The main goal of GALE is to become a generic and extensible adaptive system that can be used by many researchers and educators without all being forced to use the same type of presentation and adaptation based on the same type of rules. For this reason, it is important to be able to reuse (or, at least, emulate the behavior of) other adaptive systems developed so far. The service extension is a good way of emulating or extending other adaptive systems. Services have been added to import adaptive applications developed for other adaptive systems. Such a new service is used to retrieve the domain model (conceptual structure) and the adaptation rules for an application, but not the content (pages) of an application. The GALE distribution includes services to accept applications defined in the CAM or Course Editor [38] developed in the GRAPPLE EU project and also applications in the format used by AHA! (version 3). Besides the GAM format, GALE also supports an XML representation of its applications: the GDOM format. There is a

service to read GDOM definitions and an export plugin that can generate a GDOM representation for any application (independent of which format was used to import the application). It is possible to develop compilers (to use in a service) for formats like Interbook, Scorm, and others, but this has not yet been done until now.

In Section 7.3 we analyze the quality of the applications and extensions developed by master students in computer science during a six-week project. Our qualitative approach consists on the analysis of the answers and the comments of 16 master students of the same group, where the students commented about their experience in authoring and developing in GALE. Section 7.4 discusses the results obtained by the qualitative evaluation of the extensibility and modularity of GALE. A brief overview of this chapter can also be found in [75].

7.2 Dataset

The data considered for this chapter consists of the questionnaire answered by 16 master students enrolled in the course called “Adaptive Systems” offered in 2012. The students were in the master programs of Computer Science and Engineering, Business Information Systems, Embedded Systems and Human-Technology Interaction. The students had 4 weeks of classes, 20 hours in total. The classes intended to prepare the students to a later (group) assignment (as well as for an exam, given in week 5). In the course students learned about adaptive Web-based systems and adaptive technology-enhanced learning in general as well as about the specifics of GALE as a generic platform to create adaptive applications.

The questionnaire was made available online to all students, shortly before their project presentation (and left open for a few days). The students were not required to complete the questionnaire and the system did not log their identity. The anonymous nature of the questionnaire was intended to give students more “freedom” to answer the questions honestly.

The assignment was composed of the creation of an adaptive application, using any content the students desired, and the development of a new GALE extension in order to see whether the claimed extensibility of GALE is actually achievable and usable in practice. Students had six weeks to develop their application and extension and to then present their work. Beside the classes, all students were

given the Milkyway application [73], the GRAPPLE (and GALE) tutorial and the adaptive PhD thesis by David Smits [22] as examples.

In total 11 groups of students worked on different adaptive applications. The subjects the students chose for developing an the application are the following:

1. A tutorial/course on C-programming;
2. An introduction into Currency Trading;
3. A tutorial on Do-It-Yourself Tools;
4. A course on Elementary Electro Physics;
5. An adaptive selector for Food Recipes;
6. A Latex tutorial;
7. A tutorial about Linux Commands;
8. An adaptive Radio Player;
9. An introduction into Relativity Theory;
10. An adaptive version of a (commercial) Research Methods Book
11. An introduction into the facilities of a Sports Center.

All of these applications (except for the radio player) used adaptive navigation support through *prerequisite* relationships. Other features used by the students in some projects were *conditional inclusion of fragments*, *stretchtext*, template pages (to ensure consistent layout), content from wiki sites, and specific handling of form input (e.g. radio player). Contrary to what one might expect, students preferred to use the GAM language over the graphical authoring tools that were developed in the GRAPPLE project [38]. This may be explained by the type of students (each group contained at least one computer science student) who value full control over ease of use of the authoring interface. The applications presented by the students were of a high quality, which means that in a few weeks they were able to have an adaptive application based on existing content, with fully functional adaptation. Based on

this assertion, we start our questionnaire analysis by asking about the authoring experience in GALE.

The questionnaire was answered by 16 students. The youngest one was 22 years old and the oldest one was 28 years old. The average age was roughly 24 years. 14 out of 16 already had experience in Java language programming (required to make extensions to GALE).

7.3 A Qualitative Evaluation Approach: Authoring in and creating extensions for GALE

This section presents our analysis over the students' answers for the questionnaire about authoring and creation of extensions. We start our analysis with the students answers for the question "how difficult was the authoring as a whole, on a scale of 1 to 5?". (So 1 means very easy and 5 means very difficult.) 14 students answered, with an average score of 3.5. Only 20% of the students considered authoring to be easy (and none very easy), which corresponds to earlier findings in the GRAPPLE project (that were worse even though in the project a supposedly more user-friendly graphical authoring tool was used). Before analyzing the open comments about the whole authoring process, we analyze some aspects of authoring in GALE more in depth to understand what is easy and what makes authoring difficult.

Students were also asked about how difficult the authoring process was in terms of specific aspects of the system. Table 8 shows the answers of the students, where level 1 is the lowest level (very easy) and 5 is the highest level (very difficult). The numbers presented in the table indicate the number of students who chose that level, with percentage in brackets.

It is interesting to note that 50% of the students considered the *Installation and Deployment* aspects the most difficult part of the authoring process (marked as level 4 or 5) and none of them considered the aspect of creation (of *concepts* and *relationships*) to be very difficult. The aspect of authoring (and development) is one topic that needs to be taken care of in future work. One important remark is that the teachers of the course received, during the six weeks of the project, only one email about the *installation and deployment* topic, and they were visited by 2 students, on different days, for personal help on installation and deployment of

Aspects	Levels					
	1	2	3	4	5	No Ans.
Concepts	3 (18%)	5 (31%)	3 (18%)	3 (18%)	0 (0%)	2 (12%)
Concepts and Relationships	6 (37%)	1 (6%)	5 (31%)	2 (12%)	0 (0%)	2 (12%)
Adaptation Rules	0 (0%)	4 (25%)	7 (43%)	2 (12%)	1 (6%)	2 (12%)
Installation and Deployment	0 (0%)	5 (31%)	1 (6%)	4 (25%)	4 (25%)	2 (12%)
Test (quiz)	2 (12%)	2 (12%)	7 (43%)	2 (12%)	0 (0%)	2 (12%)

Table 8: Number of answers about the aspects of authoring regarding the ease of creation in a level from 1 (very easy) to 5 (very difficult).

GALE system. So most students found the installation difficult but did manage it without help from the teachers. Regarding ease of use, we highlight two interesting comments by the students. One student said that “We spent more time trying to understand how to implement things and debugging than actually doing relevant work.” Another one said “The setup and usage documentation should be more clearly written... Another thing could be the compiling and setup process of the entire project.” These comments combined with the other three where assistance was given by the teachers show that on the one hand we do need to make installation and documentation better, but on the other hand also show that authors of adaptive applications can be greatly helped by having pre-installed GALE servers. As GALE allows many adaptive applications (with different presentation style and different adaptation) to co-exist on a single server authors in general need not be concerned with installation. The students only needed to set up their own server because they were asked to program extensions to GALE (a more technical task than authoring an adaptive course).

The *adaptation rules* aspect shows that 10 out of 16 students consider this aspect not very easy. 10 students marked level 3 or higher for their answers. It is another aspect where we have to make some effort to make it easier and well documented. Considering that all students used a text editor to write their own configuration file

(the GAM file) instead of the graphical tool, it is very understandable that they had to do some effort in order to start writing and authoring with the GAM language, since they were not used to that language. The language can be partially written using Java code, but it is not mandatory as explained in the previous section. In any case, not all students were familiar with Java: only 87% of them said in the questionnaire that they have some experience with Java. An additional problem is that the syntax sometimes differs from what they were used to in Java. The students commented on some of their difficulties regarding the GAM language in the questionnaire. One said “I think it is not a good idea to let such a person (who does not have experience in the GAM language) edit GAM files in order to perform authoring.” Another student complemented “Sometimes it felt a little bit inefficient, but the possibilities were available...”. They were worried about the structure of the GAM file and about the language they had to write in, but they agreed in one thing: the level of experience in programming makes a difference. The first student went further on and said that “...such user should need a high level of programming experience before he/she will feel comfortable authoring in such way (using GAM language instead of graphical tool).” While the second said “The syntax of the GAM file was a little bit weird at the start of the development, but I got used to it quite fast.”

Table 8 shows the aspect *concepts* and, also, *concepts and relationships* with low level of easiness (very easy) on authoring them. As we can see in Table 8, more than 80% of the answers are marked in level 3 or lower, with 50% of the answers marked in level 1 or 2 for the *concept* aspect, while the *concepts and relationships* aspect got roughly 44% of the answers in level 1 or 2. These aspects are the basis of the adaptation, and it is important to have almost half of the students creating the concepts and its relationships in an easy way in a short period of time.

The difficulties in using GAM emphasize the importance of providing templates of adaptation rules for common situations (like e.g. dealing with prerequisites). The “inheritance” feature of GALE allows authors to concentrate on defining concepts and relationships (which was found to be easy by the students) and to simply inherit adaptation rules (as writing them was found difficult). The GALE distribution thus needs to be expanded with more examples of adaptation rules that can be reused.

As part of their applications some groups included a multiple-choice test (especially the ones developing a course or tutorial). The results for the test (quiz) aspect, presented in Table 8, suggest that the students did not need much effort to use the MC plugin and module. A multiple-choice test is represented as an HTML form (with some additional tags to identify questions and answers), is associated to a concept and is then presented as a regular course page. The “form” elements for questions and answers are generated by a GALE module and the MC test is evaluated through a GALE plug-in. Since the students all saw examples of MC tests (in source code) they could very easily learn to create their own tests.

The students also developed extensions to GALE. Some noteworthy extensions that were made are:

- One group developed an extension to the multiple-choice test module and plugin contained in the GALE distribution. In GALE an MC test is associated with a concept and the result of the test is a “knowledge” score for that concept. The students changed the code in order to link individual questions to concepts.
- Another group also tackled the multiple-choice functionality to extract explanations of answers from pages of the application (instead of copy/pasting them into the quiz).
- Some groups added new views, as used for an automatically generated navigation menu. One group added a progress bar to the top-level items in the navigation menu, and one group even created a completely new graphical (graph-based) navigation menu.
- The adaptive radio used a new plugin to randomly select a radio station with a bias towards previously liked stations.

In the questionnaire the first question about the development of extensions was “How difficult was the development of GALE extensions?”. 13 students answered this question, with an average close to 3.5. (4 students marked easy or very easy and 7 difficult or very difficult.) To better understand the students point of view about the difficulty in implementing GALE extensions we asked about the ease or

difficulty of specific aspects of designing and implementing extensions: installation, Java package explorer, code comments, developing clean code, deployment and test. Table 9 presents the students’ answers for these six aspects on a scale of 1 to 5, where 1 is the lowest level (very easy) and 5 is the highest level (very difficult).

Aspects	Levels					
	1	2	3	4	5	No Ans.
Clean Code	1 (6%)	4 (25%)	4 (25%)	3 (18%)	1 (6%)	3 (18%)
Java Package Explore	2 (12%)	0 (0%)	4 (25%)	1 (6%)	5 (31%)	4 (25%)
Code Comments	1 (6%)	2 (12%)	3 (18%)	3 (18%)	4 (25%)	3 (18%)
Installation	0 (0%)	1 (6%)	2 (12%)	5 (31%)	4 (25%)	4 (25%)
Deployment	1 (6%)	1 (6%)	2 (12%)	4 (25%)	5 (31%)	3 (18%)
Test	0 (0%)	3 (18%)	3 (18%)	4 (25%)	4 (25%)	2 (12%)

Table 9: Number of answers about the aspects of development regarding the ease of implementation in a level from 1 (very easy) to 5 (very difficult).

The first aspect we would like to point out is whether the GALE source code could be considered to be clean code, which means whether the code was written in an understandable way to allow for easy maintenance afterwards. Understanding the existing GALE code is important for creating extensions as the GALE code is a good “example” of the type of code the students had to develop. On average the students claimed the GALE code was written in a way that is not too difficult nor too easy to understand. (Note, they only commented on the existing GALE code, not the code of their extension.)

The *java package explore* aspect was not approved by the students, since 6 out of 12 of the valid answers (students who gave an answer) considered this aspect to be difficult or very difficult, and only 2 students considered that an easy aspect. This aspect was a big surprise for the GALE developer experts, since they had considered the java packages “self-explanatory.” For example, the GALE processor classes are implemented in the *nl.tue.gale.ae.processors* package, or the GALE configuration classes are implemented in the *nl.tue.gale.ae.config* package. Each functional part of GALE is located in a different directory and has a corresponding name. One

point that the students might have considered as part of this aspect is the place where the GAM file should be, as commented by a student: “The setup should be more clearly written. For example, the use and path placement of the default gam files (concept.gam).” This aspect was not discussed in detail with the students, but it is something that need to be clarified before the next experiments.

The *code comments* aspect is a topic in which the GALE developer experts really agree with the students, since the experts believe that the code relies a bit much on being self-explanatory and could indeed benefit from more comments. There are some classes that do not have explanatory comments about what the class does or what some lines of code do. That is also what the students marked in their answers. Most of them considered the code comments not that good, but a few of them mentioned that the comments are good or really good. One student pointed out that “There was hardly any code documentation, save for in-line comments.”

It is clear from the students’ comments and answers that the three aspects of the ease of implementation installation, deployment and test are the worst part for the development in GALE. Table 9 presents the results about these three aspects of implementation. We intend GALE to be usable as a basis for many adaptive system research projects, with other developers implementing and extending GALE for their specific research area. In this case, the GALE system’s source code should really be easier to understand. Besides the group assignment content, which was considered by the professors and tutors of the Adaptive Systems course to be of a good level, the students did not consider GALE an easy system for further (research) development. The main complaint of the students was the installation and setup process. One said “The installation of the various tools is quite elaborate. A series of instruction movies for different platforms (Windows, Mac,...) would be helpful to stimulate a jump start.” Another student complained “Getting GALE system to work in Eclipse is an absolute nightmare.” The setup of GALE was quite often a source of complaint as we can see in another student’s comment “I had a very hard time setting up GALE on a Macbook.” These comments came somewhat as a surprise to us, especially coming from computer science students, because we have installed GALE and compiled and run it (with or without Eclipse) on different Windows, Linux and Mac systems before giving it to the students.

The deployment and test aspects of implementation were also a reason for complaining. When a student was asked about suggestions, comments or complaints about the development of GALE extensions he said “Provide clear instructions of how the whole complex editing trajectory should be followed if nothing changes.” Another student complained about the tests he/she was doing “What is also weird is that sometimes the very same code doesn’t work while it worked a day before...”. Both students did not give more details about what happened to let us improve the test aspect of the implementation, but we understand that with the current level of debugging skills of our students problems encountered during the development of GALE extensions can be hard to find.

The group assignments did receive good grades but the students’ evaluation of the authoring and development process indicate that this was mainly due to a lot of effort by the students, not by the GALE environment making it easy. The students’ feedback through the questionnaire provides good directions to improve the quality of GALE, both as a platform for creating adaptive applications or courses and as a platform that can be extended to suit different needs and desires by adaptive systems researchers.

7.4 Discussion and Conclusions

It is important to remark that, despite the problems faced in the use of GALE, the students were able to develop applications and extensions of a high level of quality. The students used different adaptive techniques and methods; for example, some of them added *conditional inclusion of fragments*, some added *stretchtext*, some used template pages, some used content from a wiki, and some required specific handling of forms input. These show that GALE can serve a lot of different types of adaptive applications.

The analysis of the questionnaire answers provided us a very deep insight about the authoring and development process in GALE as experienced by master students in the main field of computer science (and closely related fields). The main criticism made by the students is related to the *installation* and *setup* of both GALE and the development environment. More than 50% of the students considered the process of installing and deploying an application or an extension in GALE a very

difficult process. Summarizing the comments of students, one said “We spent more time trying to understand how to implement things and debugging than actually doing relevant work.” Fortunately this criticism can be remedied by improving the installation procedure and documentation. Also, authors of adaptive applications do not need to perform any of these step, only developers and system managers need to do that.

Although students complained that the use of GALE (as author and as developer) was difficult, more difficult than we anticipated, our experiments with student-generated applications have shown that defining and implementing adaptation in a variety of adaptive applications is not really so hard: the students successfully completed their projects within the 6 week period that was allocated. We can conclude that within a few weeks designers with computer science skills can indeed make use of the extensibility and modularity of GALE to extend GALE in order to make it suitable for different types of applications. The extensibility of GALE was an explicit research and design goal because we intend GALE to be usable as a basis for many adaptive (hypermedia) system research projects. Each project can create additions to GALE for their specific research purpose but does not have to design an adaptation platform from scratch, wasting valuable development time. Since students were able to start and complete extensions to GALE in just a few weeks GALE is fulfilling its promise of being usable as a basis for adaptive systems researchers, potentially saving months or even years of effort typically spent on developing an adaptive system from scratch, only to demonstrate the effect of some particular new adaptive feature or type of application.

8 New Tools for Analyzing Navigation, Test and Quiz Logs

This chapter presents in detail the newly developed tools for analyzing navigation logs and test and quiz results. These tools were developed as a GALE plug-in, from the scratch, instead of using or integrating on-line analytical processing (OLAP) tools or any known Web analytic tools. We chose this option in order to have a new GALE extension that can be used as an example for future developments, and to keep GALE as a self-contained tool, with no external tool with external support needed. The importance of such a tool for the authors is best explained using Chapters 3 and 4 as examples. These chapters present the evaluation of two versions of an adaptive course, and the main goal is to analyze the structure of the course from the point of view of the student's navigation. In general, the authors would like to check whether the students follow the structure of the course, meaning for instance whether the students follow *good* or *bad* links, and whether the students performed the tests (at the right moment, and how well they fared).

Chapters 3 and 4 present statistical measurements generated by an external tool. Indeed, the *empirical hub coefficient* (EHC), the *informative pages* and the correlation coefficient (between the out-degree of a concept and its EHC) were generated by an external tool. It was possible for us to use an external tool because we had access to the AHA! and GALE logs. (If you have GALE installed locally, you will also have access to the GALE logs, since its configuration is stored in the GALE config file.) However, in general, it is difficult for authors to make use of the logs due to the following problems and restrictions:

1. You need access to the application logs;
2. You need an external tool to analyze the logs;
3. The analysis is made “offline”, which is not ideal since GALE serves on-line courses and applications.
4. Security and privacy. The AHA! and GALE log files contain all log data for that server since the log is not split into different files for different

applications, so authors analyzing the log have access to the log of other authors' applications.

An analysis tool built into the server could tackle these problems.

The main goal of the analysis tools (log, test and quiz tools) is to summarize the log data and present it in an “understandable” way to the authors of an adaptive application. In GALE these tools are implemented as plug-ins. The log analysis tool gives to the author an important overview of the students' navigation over the structure of the course. For instance, it can be used to answer our first two research questions: 1. does the annotation/hiding adaptive technique influence the choices of the students? 2. how does the interplay between the link structure of an adaptive course and the rules employed by the adaptation influence the choices of the students? One can say that the first question is not entirely answered since we do not know why the student is, for example, following a *bad* link. In this case, the author should ask the student why they are following these links using the *quiz* plug-in (we did that in the evaluation presented in Chapter 4). The analysis of such answers can be done by the quiz analysis tool. The evaluation presented in Chapter 4 was made partially with the support of the tool described in this chapter.

We have written Sections 8.1 and 8.2 in such a way that you can read them in any order. Section 8.1 is a more technical explanation and Section 8.2 describes the statistical measurements presented by these tools and present some examples. If this thesis were adaptive this navigation freedom would have been obvious, and redundancy between the sections could have been avoided to a larger extent.

8.1 Technical Support

The GALE system allows programmers to extend its functionality by implementing their own *Java Classes* for modules, processors, plug-ins, services, etc., as described in Chapter 6. For the analysis tools we implemented plug-ins. A plug-in in GALE is an implementation that has the *pluginProcessor* associated with it. A plug-in is capable of “bypassing” the (remainder of the) processor pipeline to perform completely different processing and generate a result “page” that is not processed or adapted further by any processor that comes “later” in the pipeline. Chapter 6 presents an insight on how to implement a plug-in in GALE. For further references

about creating a GALE plug-in, read David Smits' thesis [22]. A plug-in in GALE generates a “complete presentation”. In more technical terms, a plug-in allows the programmers to have direct control over the client HTTP request and response. Good examples of how it works are the *logout plug-in*, where the user is logged out of the application, or the *MCPlugin* that evaluates a multiple test and presents a score to the user. In both cases they generate a response to the client and skip all the subsequent processors (more details in Chapter 6) that would for instance generate a layout with header and navigation menu. When a request is handled by a plug-in the presentation will not show such header or menu (unless the plug-in itself generates them).

The GALE log is stored in a file on the GALE server. It is important to mention that the navigation, test and quiz logs are written in different files. The log files are logged “on-demand” (we mean: for every request to a concept, test or quiz the *LogProcessor* is called which generates a log entry and flushes its output buffer to guarantee completeness of the log even in the case of a web server crash). GALE uses the Spring¹² “inversion of control” (IOC) to configure and instantiate all the components. With Spring IOC the programmers can add the new *LogManager* to retrieve and analyze the log file by means of a newly implemented *Java class* in two steps: 1. Configure the Java class (in Spring it is called *bean*) in the config file (*galeconfig.xml*) to use the *LogManager bean property*; and 2. Inject the *bean property* in the implemented Java class.

The navigation log file contains the following data: the *date and time*, the *user*, the *concept*, the *resource* (URL of page), the *view*, the *link class* and the *referrer*. These data are better explained in Section 4.1. The *view*, the *link class* and the *referrer* are data that we missed in the evaluation presented in Chapter 3. The *link class* presents the annotation state of links followed by the user. We would like to mention that to log this state of links it was necessary to implement a new *processor*. The reason to implement this processor is that the annotation state is computed from UM, and the *UpdateProcessor* changes the basis for this computation before the *LogProcessor* runs. In that case, the new processor have to come before the *UpdateProcessor* in the pipeline. The new processor stores the link class in an UM attribute that is not touched later by the *UpdateProcessor*. The *LogProcessor* can

¹²See www.springsource.org for more details about Spring.

then log that state for our analysis. Note that simply reversing the order of execution of the *UpdateProcessor* and *LogProcessor* (to log the annotation state before the UM is changed) is not possible because the *resource* or page to be presented to the user is adaptively determined based on the UM state, and it is determined based on the *new* UM state after running the *UpdateProcessor*. Running the *LogProcessor* first would cause the wrong or perhaps even an undefined *resource* value being logged.

The implementation and use of the analysis tools as plug-ins in GALE can be considered an easy task for a Java programmer, since the plug-in is written as a Java class, and it extends the abstract class called *AbstractPlugin* which implements the *Plugin* interface. GALE contains a number of plug-ins already and these can also serve as example code for programmers creating new plug-ins. It is important to remark that there are three different plug-ins for the log, test and quiz analysis.

To make a plug-in available in GALE it is necessary to configure that in the *galeconfig.xml* file. You must insert the new plug-in Java class in the list of plug-ins known by the *PluginProcessor*. When creating a new plugin it is important to be aware that after the *PluginProcessor* has run all subsequent processors in the pipeline are skipped. The example below shows the lines inserted in *galeconfig.xml* to make the log analysis tool available in GALE.

```
<entry>
  <key>
    <value>analysis</value>
  </key>
  <bean class="nl.tue.gale.ae.processor.plugin.AnalysisPlugin">
    <property name="logManager" ref="logManager" />
  </bean>
</entry>
```

When a GALE request (URI) ends with `?plugin=analysis` the *AnalysisPlugin* will be called. The *AnalysisPlugin* is the implementation of the log analysis tool.

The mechanism that reads the navigation, test and quiz logs works in the same way. The GALE *LogProcessor* stores log information in files. This is fine for adding entries to log files but the file I/O mechanism is not really suitable for the analysis plug-ins to read the log files each time the plug-in is called. For better performance

the analysis tool reads the log file the very first time it is called and then saves the data in an in-memory HyperSQL database¹³. The next time the analysis tool is called, it verifies whether there is an update (new lines are included) in the log file. If this is true, the tool only reads the new lines and inserts them in the database. An in-memory database could reduce the performance of the server, but in our experience the biggest log file we read was not longer than 1GB and did not pose a problem for the running server.

For security reasons and to keep the confidentiality of the users, only authors of an application have access to the analysis tool and the tool also restricts that access to the data of the authors' own applications. Figure 18 shows a page with a menu in the header (top right position) used to access the analysis tools. To analyze the log of an application, such as the 2I080 course presented in the Chapter 4, the author has to access the 2I080 course and then (s)he can call the tools.

The *AnalysisPlugin* is called when an author selects the *Log Analysis tool* from the menu presented in Figure 18. Figure 19 shows a list (*hardcoded*) of possible statistical overviews the *AnalysisPlugin* can generate. The four topics shown in Figure 19 are HTML links which call the *AnalysisPlugin*. These links pass the URL as a parameter to the plug-in. The parameter is called *action* and it identifies the topic in the list. For example, when the author clicks on the first link topic *#Visits per CONCEPT*, the GALE request (URI) ends with `?plugin=analysis&action=concept`. In this example the *action* parameter has the value *concept*.

The analysis plug-in gets the HTTP request (using the GET method) and parameters. A query in the database is generated depending on the parameters. The query returns a *result set*. As a plug-in, the analysis tool generates a result "page", which means that the response is a HTML page. The HTML page is created based on the *result set*. Figure 20 shows the example above, where the *action* parameter has the value *concept*. In the example, the *result set* returns the number of times a concept was visited by the users. The test and quiz plug-ins also use such parameters to retrieve data from the corresponding logs.

The last two columns of the table presented in Figure 20 contains links to drill-down the measure per incoming link or per outgoing link, respectively. We present

¹³See www.hsqldb.org for more details about HyperSQL Database.

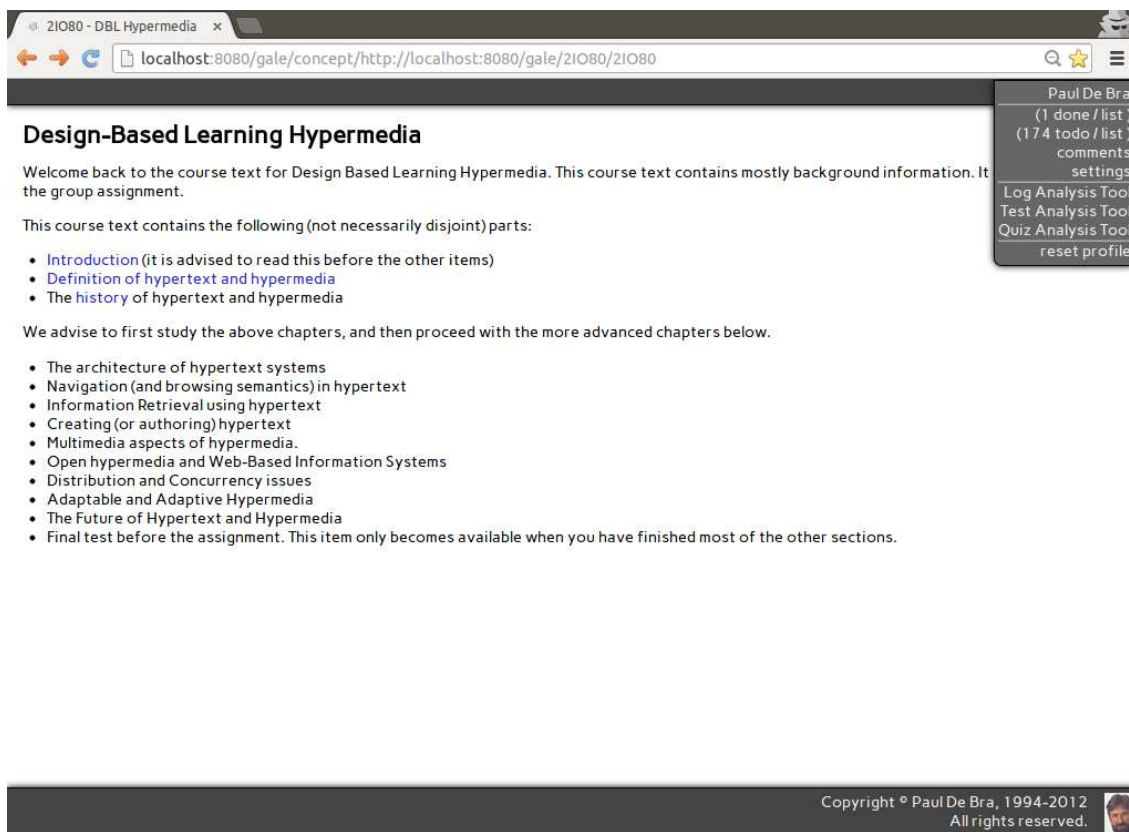


Figure 18: Screenshot of a concept with links to the analysis tools within the menu in the top right position.

more detail about these measures in the next section. At this point we would like to remark that these links also pass parameters through the GALE request, which is `?plugin=analysis&action=concepts_incoming&conceptParam=XXX` and `?plugin=analysis&action=concepts_outgoing&conceptParam=XXX`, for the incoming link and outgoing link, respectively. The XXX value depends on the line of the table, for example, the value of XXX in the first line is *21080*, and in the second line is *test-all*.

The next section presents the possible statistical overviews that can be generated for GALE applications.

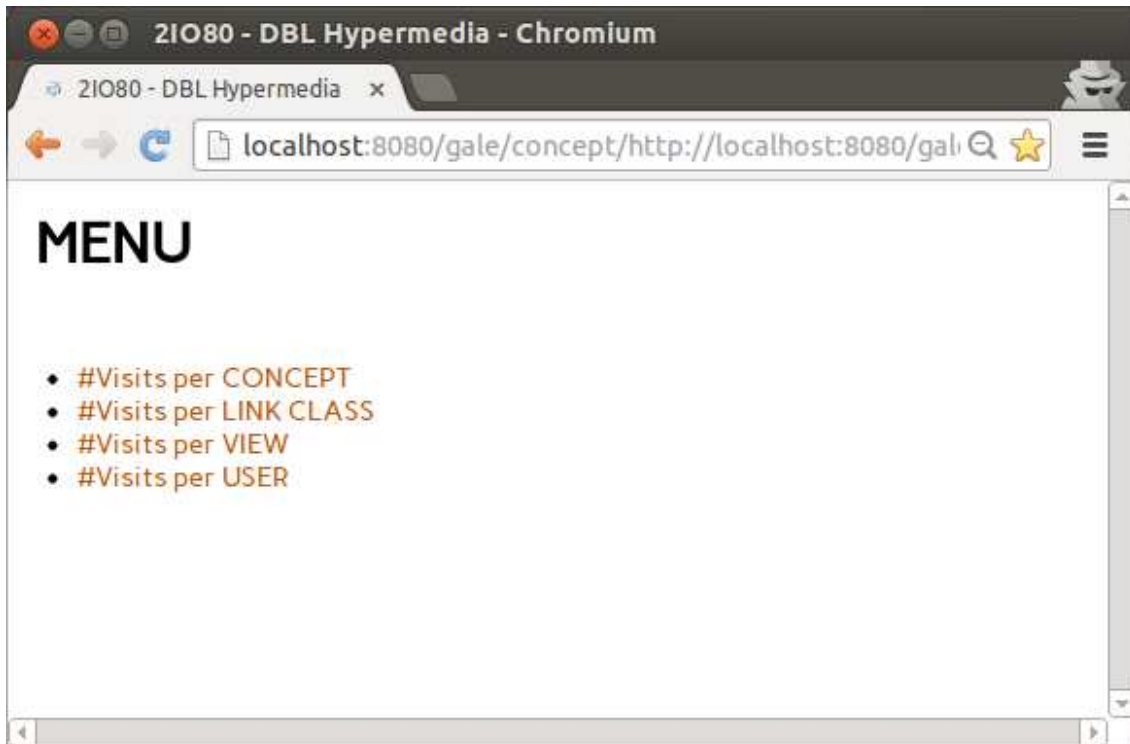


Figure 19: Screenshot of the menu options for the Log Analysis Tool.

8.2 Statistical Measures

In this Section we present the statistical measurement used in the evaluation of an adaptive course and how it would help authors improve their own courses. Earlier in this chapter we said that the new tools were developed to keep GALE as a self-contained tool and to give users an example on how to implement new extensions to GALE (more precisely a GALE plug-in). We could have developed new tools presenting *basic* statistical measurement, e.g., the number of users that accessed an application or the number of tests performed in a course), but the power of these tools is similar to the OLAP tools: to roll-up and/or drill-down the data. To differ from OLAP tools we implemented the quiz and test analysis tools, and what we call *the EHC and informative pages table* (for *EHC* and *informative pages* refer to Section 3.4, Section 4.4 or summarized in [49]).

The analysis we have done in Chapter 4 were made partially using the tools presented in this section. To better explain how such tools could, for instance, assist the author of the hypermedia course of Chapter 4, we recall the research questions

#Concepts Accesses

1	2IO80	1583	Incoming	Outgoing
2	test-all	720	Incoming	Outgoing
3	history	645	Incoming	Outgoing
4	intro	557	Incoming	Outgoing
5	architecture	522	Incoming	Outgoing
6	definition	471	Incoming	Outgoing
7	navigation	462	Incoming	Outgoing
8	authoring	354	Incoming	Outgoing
9	retrieval	316	Incoming	Outgoing
10	adaptivehypermedia	261	Incoming	Outgoing
11	openhypermedia	229	Incoming	Outgoing
12	multimedia	227	Incoming	Outgoing
13	distribution	225	Incoming	Outgoing
14	tower	198	Incoming	Outgoing
15	test-history	189	Incoming	Outgoing
16	hyperdocument	187	Incoming	Outgoing
17	xanadu	172	Incoming	Outgoing
18	structural-analysis	170	Incoming	Outgoing
19	ham	166	Incoming	Outgoing
20	test-introduction	160	Incoming	Outgoing
21	dexter	147	Incoming	Outgoing
22	future	147	Incoming	Outgoing
23	test-definition	144	Incoming	Outgoing

Figure 20: Visits per Concepts: the number of visits for a concept and links to drill-down the measure per *Incoming* link or per *Outgoing* link.

we have made in this thesis: 1. does the annotation/hiding adaptive technique influence the choices of the students? 2. how does the interplay between the link structure of an adaptive course and the rules employed by the adaptation influence the choices of the students? These questions will lead us to a good example on how these tools will help the authors. Before we answer these questions, we present the statistics implemented in the log analysis tool. In general, the statistics presents the number of visits of a specific measure grouped by other(s) measure(s). The tool presents the following statistics:

1. Visits per Users: it presents the number of visits (or clicks) per user; Clicking on the statistical number presented, the author is given a list with all entries for that specific user;

2. Visits per Concepts: it presents the number of visits for a concept. Here, the author can check the most visited concept. Authors can access the visits per concepts per incoming link or the visits per concepts per outgoing link. Figure 20 shows an example of this measure for an *adaptive course* served by GALE.
3. Visits per Concepts per Incoming Links: it presents the number of visits for a concept grouped by incoming concept (the concept from which the user came). This measure allows the author to check from which concept the user came from to visit the current concept.
4. Visits per Concepts per Outgoing Links: it presents the number of visits for a concept grouped by outgoing concept (the concept accessed after visiting the current concept). In Chapter 3 we discuss the idea of *informative pages*. Recall that an *informative page* is a page with many links and information about the links, where the information intends to help the user choose the correct link to follow.
5. Visits per Link Class: it presents the number of visits (or clicks) to a concept grouped by link class (for more detail on link class see Chapter 3). Chapters 3 and 4 show the importance to the author to know which concepts are accessed through bad links. For this reason, the statistic *visits per link class* shows the number of visits (or clicks) grouped by link class.
6. Visits per Link Class per User: it presents a list of visits grouped by user and link class, i. e., the number of clicks in a specific link class grouped by user. This measure is discussed in Chapter 5, where the author can find the students who follow bad links.
7. Visits per Link Class per Concept: it presents a list of visits grouped by concept followed by the user using a specific link class, i. e., the number of clicks on a (link to a) concept using a specific link class. This measure is discussed in Chapter 5, where the author can find the different students who followed the same bad link.

8. Visits per View: it presents the number of clicks in a link grouped by view (for more detail on system views see Chapter 6). An application can have different views. For example, the *Adaptive Course* presented in Chapter 4 has three views: content, where the text of the concept is presented, menu, where the structure of the course is shown as a tree of concepts and sub-concepts, and, header, where configuration menu and user information are presented. This measure shows in details how the views were used by the users in the application. The measure counts the number of times a user followed a link from within each view and also which (links to) concepts were followed from within each view. This is an important measure for the authors since they can understand how the users are navigating through the application. For this measure the author can see the visits per views per user and visits per view per link class.
9. Visits per View per User: it presents the number of clicks in a specific view grouped by user. For example, David Smits presents in his *thesis application* served by GALE a view called *next topic view* [22]. This view presents a link to the next concept in the concept tree order (this order was created by Smits), which means that if the thesis were a book the next topic would be the “next page” of the book. Considering that it is an adaptive application, and experts navigated through that, it would be interesting to check whether users used the *next topic view* or not.
10. Visits per View per Link Class: it presents the number of clicks in a specific view grouped by link class. Chapter 4 shows how students follow the advice given through link annotation and link hiding in the different views of the *adaptive course*. In this manner, authors can distinguish between bad links followed through annotated links versus bad links followed through hidden links. It is also possible to check whether students look for good links when they are navigating through the course pages or if they consider all the links that appear in menu-like structures. From this measure, authors access the visits per view per link class per user.
11. Visits per View per Link Class per User: it presents the number of clicks on links in a specific view using a specific link class grouped by user.

The first level is composed by: Visits per User; Visits per Link Class; Visits per View; and Visits per Concept. The tool present a link and the author can click and drill-down in the statistics. For example, when the author visits the statistic *visits per link class*, he can drill-down to check which users visited a specific link class or he can also drill-down to check which concepts were visited using a specific link class.

Another important remark is that from the lowest level (or the deepest level) the author can drill-down to the log entry list. For example, if the tool is showing the visits for *conceptA* where users accessed that using a *bad* link, the author can drill-down and check the complete information (student ID, date, activity, session, score, view, link class and referrer, presented in Chapter 4) for all log entries where the concept is *conceptA* and the link class is *bad*.

We show in this section how an author could use the analysis tools to answer our first two research questions (or, at least, try to answer these questions). Start with the question “how many *bad* links the students followed and from which view?” To answer this specific question the author need to access the log analysis tool and click on the statistic called *Visits per View*. At this moment, the tool will show the possible views of the course (the *hypermedia* course in Chapter 4 has the *content*, *header* and *static-tree-view (menu)* views), the number of visits for each of them and two links to drill-down the data: *Visits per View per User* and *Visits per View per Link Class*. Thus, to find out how many bad links were followed from each view it is necessary to follow the link: *Visits per View per Link Class* for each of the views. The *hypermedia* course would show 471 visits via bad links in total, where 262 from the menu view and 209 from the content view. Going further and clicking on the content view via bad links we find that 105 out of 209 clicks came from the *welcome* page. Roughly 50% of the bad link accesses from the content view came from the *welcome* page. Considering the link hiding/annotation and going deeper in this analysis, we find out that 87% of the bad link visits are via link annotation, as shown in Chapter 4. This is a simple example to show the importance of such statistics to the authors of a course. To get a complete answer to our research question we need more than only a simple question, we could have done a few more questions to try to infer whether the link hiding/annotation influence the student’s choice. For example, we asked: a. which concept is more visited via bad links? b.

how many users accessed a concept via bad links? c. how many users accessed a concept via bad links more than 5 times? These questions would start the analysis of the influence of the adaptive technique on the student's choice, but we are aware that we need more than only a few statistics to understand the student's behavior.

Our second research question is related to the link structure of the course and its influence over the student's choice. In our analysis in Chapter 3 we define the *empirical hub coefficient* (EHC) and *informative pages* (for more details refer to Section 3.4 or [49]). To confirm our definition of EHC and informative pages and our hypothesis that the link structure influence the student's choice, we implemented the EHC and informative pages tables, which shows, exactly, Tables 5 and 6 and the correlation coefficient between the EHC and the out-degree of the concepts in Chapter 4. Section 4.4 presents the analysis we have made for these tables.

We also developed a tool for the analysis of the tests, we call that *test analysis tool*. Before presenting the tool, we would like to remark that a test in GALE is of multiple choice form. The author inserts as many questions and possible answers (for each question) as he wants. There is no limit to the possible number of answers for each question, and different questions may have a different number of (correct and wrong) answers. The questions and answers are presented to the user randomly. The author can indicate the number of questions and (correct and wrong) answers that will be presented for the user.

First the author chooses a test, then the tool presents a list of questions and answers for that test. Figure 21 shows an example of the tool where the questions of a test called *test-history* are presented. The tool shows the number of times a question was answered by the users. In Figure 21 the words "*Appearance times: 4*" shows that the first answer (and also the second, in this case) was answered 4 times by the users. For each question we show the number of times the answers were shown for that question and, also, the number of times the users marked that answer. For example, in Figure 21 the first question and second answer shows *Intermedia 3/2*, where 3 means the number of time this answer appeared for that question and 2 means the number of times a student chose that answer. The correct answers are shown in bold.

The test analysis tool is an opportunity for the authors to check the users' answers and mistakes and it can be used by the author, depending on the questions

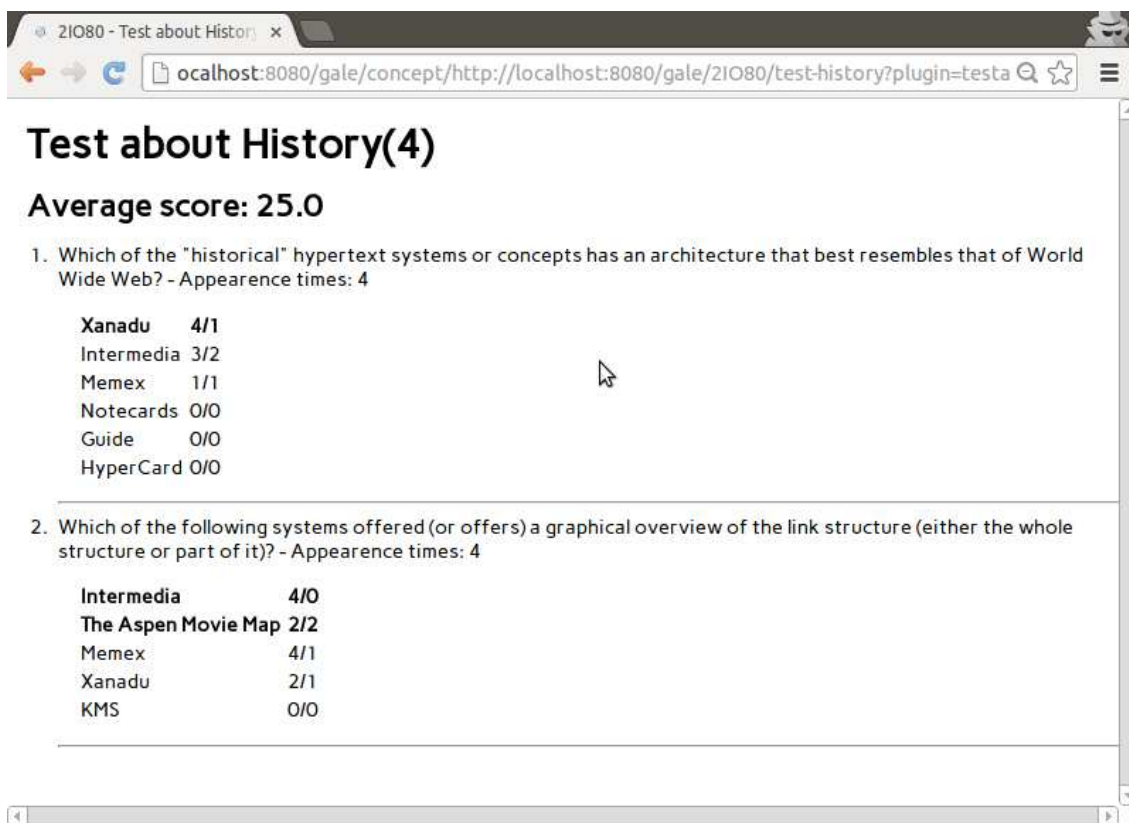


Figure 21: A screenshot of the test analysis tool. The test called *test-history* is shown as an example of the tool.

he created, to review the concepts related to that question. For example, an author can see that in *question1* 85% of the users marked the wrong answer *answerA*. If *question1* is related to *conceptA*, it indicates that, probably, the approach used to explain this concept needs a review, since 85% of the users marked the wrong answer.

The quiz log analysis is another tool that we developed. A “quiz” in GALE is like a test, with no randomly selected question and answers, no right answer and, consequently, no score. An example on how and why to use a quiz is presented in Chapter 4. There we used a quiz to ask the students why they were following bad links. The quiz log analysis tool presents a simple table with the quiz questions and answers with the number of times each answer was marked by the users.

In the next chapter we discuss, in more detail, the work we intend to design and implement to make these tools more powerful.

9 Conclusion and Future Work

In this thesis we discussed the evaluation of an AH system and AH applications. The field of AH is already an important field in the academic research area with applications in e-learning, e-culture, entertainment, recommender systems, etc. We have concentrated on e-learning (or technology-enhanced learning to use a more modern term). The users, authors and developers are constantly bringing new ideas and needs to the AH field. In most of the cases these needs and ideas are subject to an evaluation.

9.1 Contributions on Adaptive Course Evaluation

Briefly, the main goal of all the evaluations presented in this thesis is to analyze the students navigation and students' feedback to assist authors of an adaptive course to improve the quality of the course and, also, to the developers of the GALE system to improve the quality and usability of the system.

We have performed an analysis of the presentation classes of the links in the “Hypermedia Structures and Systems” course (to simplify we call it as the *Hypermedia* course) offered by the Eindhoven University of Technology since 1994, focusing on detecting whether students follow the so-called *bad* links and if their navigation behavior is affected by the presence of *bad* links. We found that pages that can be classified as authorities (i.e. that have many incoming links) make the students in some sense curious and cause them to investigate the pages even when they are not yet recommended by the system. The overall conclusion of our study is that when an author wishes for students to follow the adaptive guidance, *authorities* that are not recommended from the start should be avoided. We also found out that students who start the course late in the agenda tends to ignore the link hiding/annotation technique and follows more *bad* links than other students.

In a later analysis of an updated version of the *Hypermedia* course (which was carried out six years later with a different set of students) we also gave a questionnaire which the students who performed the course answered. As in the previous analysis, we found that students who start studying late in the course period tend to follow more *bad* links than students who use the entire period. We also noticed in the log and questionnaire analysis that students often followed

bad links because they were curious about the linked non-recommended concept or because they are exploring to know more about what will be offered in the course before they start studying the material in more depth. We have observed that students navigating via *bad* links have a tendency of continuing to use *bad* links in the subsequent pages.

In the analysis of the updated *Hypermedia* course, we also noticed that there are more log entries during the exam (4,495 entries) than during the whole period leading up to the exam (3,976 entries). An important work to be done in later courses is to stimulate the students right from the beginning of the course term. The absence of classes or other contact with a tutor during the course period leads students to postpone their study activity until right before the exam.

Our hypothesis says that the adaptive technique is an effective model of adaptation and students' guidance. Our findings in the first evaluation lead us to an educated guess about the influence of the hiding adaptive technique to guide the students through the course: most of the students did not follow the system's suggestion because of their curiosity. Regarding the second evaluation, we observed that, when used with care, the *link hiding* technique is an effective way to guide the students through the course. However, our analysis also revealed that the *link annotation* technique is not as effective as the hiding technique in guiding the students.

9.1.1 Hubs and Informative Pages

We have also carried out a structural analysis of the course, comparing the definition of *hubs* to new empirical measures, called *empirical hub coefficient* (EHC) and *informative pages*. Intuitively, *hubs* are pages that contain a large number of links to others pages. We cannot apply this concept directly in our setting, since in an adaptive application the number of links of a web page may change over time. The EHC of a page (or, more precisely, a concept) X is the ratio between the number of times students clicked on a link of X to go to a different page (or concept) and the number of times that students accessed page X. We also introduce the notion of *informative pages*, which are pages that not only have a large number of links, thereby being good candidates for *hubs*, but also contain enough information to *guide* the student through the links of the page.

We have observed that *empirical hubs* appear in the course we studied, and most of these *hubs* turned out to also be *informative pages* that guided the navigation of the students through the large number of links that *empirical hubs* contain. We also identified the presence of non-informative *hubs* that give the student many choices but no guidance. An author should ensure that all *hubs* have information to guide the students through the links in the page. In other words, an author should turn the hubs into *informative pages*.

9.2 Contributions on GALE Extensibility Evaluation

Another important goal of this thesis is the evaluation of the extensibility and modularity of GALE, since it was created to be a generic open source extensible adaptive engine. Our hypothesis is that GALE reach its goal as a generic and extensible adaptive system serving different kinds of adaptive applications. In this case, we got feedback from master students through a questionnaire and by analyzing the applications and extensions developed by them. The students had a few problems developing their own applications and extensions. Despite the problems faced in the use of GALE, the students were able to develop applications and extensions of a high level of quality. The students used different adaptive techniques and methods; for example, some of them added *conditional inclusion of fragments*, some added *stretchtext*, some used template pages, some used content from a wiki, and some required specific handling of forms input. These show that GALE can serve a lot of different types of adaptive applications.

The questionnaire analysis about the development and creation of adaptive application by master students gave us a very deep insight on that task. The main criticism made by the students is related to the installation and setup of both GALE and the development environment. More than 50% of the students considered the process of installing and deploying an application or an extension in GALE a very difficult process. If we were able to summarize the comments of all students in one sentence we would choose this one “*we spent more time trying to understand how to implement things and debugging than actually doing relevant work.*” We are aware that this comment does not represent the opinion of all students, but we observed that most students shared this criticism. Fortunately this criticism

can be remedied by improving the installation procedure and documentation. Also, authors of adaptive applications do not need to perform these steps, only developers and system managers need to do that.

Although students complained that the use of GALE (as an author or developer) was difficult, more difficult than we anticipated, our experiments with student-generated applications and extensions have shown that defining and implementing adaptation in a variety of adaptive applications is not really so hard: the students successfully completed their projects within the 6 week period that was allocated. In fact, within the 6 weeks of the project, students not only created their own adaptive applications, but also implemented different kinds of GALE extensions, as shown in Chapter 7. The conclusion from this work is that within a few weeks designers with computer science skills can indeed make use of the extensibility and modularity of GALE to extend GALE in order to make it suitable for different types of applications. The extensibility of GALE was an explicit research and design goal because we intend GALE to be usable as a basis for many adaptive (hypermedia) system research projects. Each project can create additions to GALE for their specific research purpose but does not have to design an adaptation platform from scratch, wasting valuable development time.

9.3 Limitations

In this thesis we investigated a few aspects of adaptation that influence the choice of the students for an adaptive course created and served by AHA! and GALE. We are aware that our investigation here has some limitations; e.g., we could only analyze two adaptive courses. Another important point is that the evaluations of adaptive courses here were made in two different versions of the same course. Consequently, being limited to the findings of this analysis and to the results obtained in these courses, we cannot generalize our results to other adaptive courses, even ones developed with AHA! or GALE.

Another limitation we would like to point out regards the group of students used in this evaluation. The experiments were made with three different groups: two groups of bachelor students and one group of master students. The experiments with the group of bachelor students were made in the evaluations in Chapters 3

and 4, with 76 and 46 students, respectively. Despite not being the ideal scenario, it gave good insights on the influence of the adaptation on the students choice. The group of master students appeared in the evaluation presented in Chapter 7, and they were only 16. Nonetheless, this study gave us a good starting point to prepare quality materials to improve the user experience in developing new applications and extensions.

The small number of courses we analyzed, and the small number of students present in each course, prevent us from drawing general conclusions regarding the methods of evaluation suggested in this thesis. In one hand, if the evaluation turns out to produce poor results, this does not imply that AHA! and GALE are bad platforms: the bad results could have been caused because the author of the course did not do a good job in designing the course. On the other hand, if the outcome of the evaluation is good, it could still have been the case that AHA! and GALE are bad authoring platforms: the course could have been evaluated well simply because the authors of those specific courses were particularly effective in designing and implementing the course. Consequently, our results in this thesis consists of a proof of concept, with some lessons to be learned from our case studies.

9.4 Future Work

In Chapter 8 we discussed the analysis tools that help authors of adaptive applications in GALE investigate the log navigation, test answers and quiz answers of their users. The analysis tools were partially used in this thesis to evaluate the adaptive course in Chapter 4. For future work we intend to improve the number of statistics measures presented in these tools. Here we list some of them:

1. The students who visited the related concepts after failing the test.
2. The students who visited the related concepts before answering the test.
3. The percentage of students that failed the test and visited at least one of the related concepts.
4. The percentage of students that failed the test and did not visit any of the related concepts.

5. The percentage of students that passed the test and visited at least one of the related concepts.
6. The date and time the quiz or test was answered.

In GALE the logs are stored in text files (including navigation logs, test and quiz results). It consumes much time and effort to read and write in these files. An improvement for GALE is to implement new Java classes to log using a database. This implementation will reduce the time needed to develop tools for analyzing the logs, and also speed up the time to process the logs.

At this moment, only the authors are authorized to access the analysis tools for their own applications. We are aware that this is not the ideal. For future work we intend to create a mechanism where the author can allow other users to access the analysis.

A structure that will make the navigation tool more powerful is to implement a filter and a sort mechanism. With these two mechanisms the user can create his own statistical measures. For example, one can wish to filter the data to see only the students that visited a given concept, or that visited a certain concept via a bad link, or that visited that concept on a given date.

We stated in Section 9.3 our limitations, one of which being that the results observed in this thesis are restricted to the evaluation of only two adaptive courses. For future work we are interested in carrying out the evaluation of other courses, using the techniques presented in Chapter 5 and 4. Link structure analysis of others courses would confirm the definitions of *empirical hubs* and *informative pages* we propose in this thesis.

For the AH system field there are many challenges to be tackled. The privacy of the users is one thing that needs to be improved, such as having a client-side user model, created and maintained through the browser, without the need to install any software on the client. Another challenge is to have part of the computation for the adaptation in the client side. Even a “lowly” tablet or phone has enough processing power to perform the adaptation for a single user. The last thing, and most challenging, the adaptation should in the future become able to correct itself. Automated analysis of the behavior of the entire user population should reveal that certain “optimization” of the adaptation is possible. Realizing such *adaptation of*

the adaptation or *meta-adaptation* requires interesting new research, especially to ensure that the adaptation can only change “for the better”.

References

- [1] BRUSILOVSKY, P., KARAGIANNIDIS, C., AND SAMPSON, D. Layered evaluation of adaptive learning systems. *International Journal of Continuing Engineering Education and Lifelong Learning* 14 (2004), 2004.
- [2] SMITS, D., AND DE BRA, P. GALE: a highly extensible adaptive hypermedia engine. In *Proceedings of the 22nd ACM conference on Hypertext and Hypermedia* (New York, NY, USA, 2011), HT '11, ACM, pp. 63–72.
- [3] DE BRA, P., KNUTOV, E., SMITS, D., STASH, N., AND RAMOS, V. F. C. Gale: a generic open source extensible adaptation engine. *New Review of Hypermedia and Multimedia* 19, 2 (2013), 182–212.
- [4] BATES, T. *Technology, e-learning and distance education*. Studies in Distance Education. Routledge, 2005.
- [5] FRASSON, C., GAUTHIER, G., AND MCCALLA, G. Intelligent tutoring systems. In *Proceedings of the Second International Conference on Intelligent Tutoring Systems* (Berlin, 1992), Springer-Verlag.
- [6] BRUSILOVSKY, P. Intelligent tutoring systems for World Wide Web. In *Proceedings of the Third International WWW Conference* (Darmstadt, 1995), pp. 42–45.
- [7] BRUSILOVSKY, P., SCHWARZ, E. W., AND WEBER, G. ELM-ART: An intelligent tutoring system on World Wide Web. In *Proceedings of the Third International Conference on Intelligent Tutoring Systems* (London, UK, 1996), Springer-Verlag, pp. 261–269.
- [8] BRUSILOVSKY, P., EKLUND, J., AND SCHWARZ, E. Web-based education for all: a tool for development adaptive courseware. *Computer Networks and ISDN Systems* 30, 1–7 (1998), 291–300. Proceedings of the Seventh International World Wide Web Conference.
- [9] DE BRA, P. Teaching hypertext and hypermedia through the Web. *Journal of Universal Computer Science* 2, 12 (1996), 797–804.

- [10] KAY, J., AND KUMMERFELD, R. J. An individualised course for the C programming language. In *Proceedings of the Second International WWW Conference* (Chicago, 1994).
- [11] GONSCHOREK, M., AND HERZOG, C. Using hypertext for an adaptive help system in an intelligent tutoring system. In *AI-ED, 7th World Conference on Artificial Intelligence in Education* (1995), pp. 274–281.
- [12] BUSH, V. As we may think. *The Atlantic Monthly* (1945).
- [13] BRUSILOVSKY, P. Methods and techniques of adaptive hypermedia. *User Modeling and User Adapted Interaction* 6, 2–3 (1996), 87–129.
- [14] BRUSILOVSKY, P. Adaptive hypermedia. *User Modeling and User Adapted Interaction* 11, 1–2 (2001), 87–110.
- [15] KNUTOV, E., DE BRA, P., AND PECHENIZKIY, M. AH 12 years later: a comprehensive survey of adaptive hypermedia methods and techniques. *New Review of Hypermedia and Multimedia* 15, 1 (2009), 5–38.
- [16] DE BRA, P., HOUBEN, G.-J., AND WU, H. AHAM: a Dexter-based reference model for adaptive hypermedia. In *Proceedings of the Tenth ACM Conference on Hypertext and Hypermedia* (New York, NY, USA, 1999), HYPERTEXT '99, ACM, pp. 147–156.
- [17] WU, H. *A Reference Architecture for Adaptive Hypermedia Applications*. PhD thesis, Eindhoven University of Technology, The Netherlands, 2002.
- [18] DE BRA, P., SMITS, D., AND STASH, N. The design of AHA! In *Proceedings of the seventeenth conference on Hypertext and hypermedia* (New York, NY, USA, 2006), HYPERTEXT '06, ACM, pp. 171–195.
- [19] DE BRA, P., AERTS, A., SMITS, D., AND STASH, N. AHA! meets AHAM. In *Proceedings of the Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (London, UK, 2002), AH '02, Springer-Verlag, pp. 388–391.

- [20] KNUTOV, E., DE BRA, P., AND PECHENIZKIY, M. Generic adaptation framework: a process-oriented perspective. *Digital Information* 12, 1 (2011).
- [21] KNUTOV, E. *Generic adaptation frameform for unifying adaptive web-based systems*. PhD thesis, Eindhoven University of Technology, The Netherlands, 2012.
- [22] SMITS, D. *Towards a Generic Distributed Adaptive Hypermedia Environment*. PhD thesis, Eindhoven University of Technology, The Netherlands, 2012.
- [23] DE BRA, P., PECHENIZKIY, M., VAN DER SLUIJS, K., AND SMITS, D. GRAPPLE: integrating adaptive learning into learning management systems. In *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications 2008* (Chesapeake, VA, USA, 2008), J. Luca and E. Weippl, Eds., pp. 5183–5188.
- [24] KAPLAN, C., FENWICK, J., AND CHEN, J. Adaptive hypertext navigation based on user goals and context. *User Modeling and User-Adapted Interaction* 3, 3 (1993), 193–220.
- [25] CALVI, L. Formative evaluation of adaptive CALLware: a case study. In *Proceedings of the International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (London, UK, 2000), Springer-Verlag, pp. 276–279.
- [26] TOTTERDELL, P., AND BOYLE, E. The evaluation of adaptive systems. In *Adaptive User Interfaces*, D. Browne, P. Totterdell, and M. Norman, Eds. Academic Press, London, 1990, pp. 161–194.
- [27] BRUSILOVSKY, P., KARAGIANNIDIS, C., AND SAMPSON, D. The benefits of layered evaluation of adaptive applications and services. In *Empirical Evaluation of Adaptive Systems: Proceedings of Workshop at the Eighth International Conference on User Modeling, UM2001* (2001), S. Weibelzahl, D. Chin, and G. Weber, Eds., pp. 1–8.

- [28] BRUSILOVSKY, P., AND EKLUND, J. A study of user model based link annotation in educational hypermedia. *Journal of Universal Computer Science* 4, 4 (1998), 429–448.
- [29] KLEINBERG, J. M. Authoritative sources in a hyperlinked environment. *J. ACM* 46, 5 (September 1999), 604–632.
- [30] DE BRA, P., AERTS, A., BERDEN, B., DE LANGE, B., ROUSSEAU, B., SANTIC, T., SMITS, D., AND STASH, N. AHA! the adaptive hypermedia architecture. In *Proceedings of the Fourteenth ACM Conference on Hypertext and Hypermedia* (New York, NY, USA, 2003), HYPERTEXT '03, ACM, pp. 81–84.
- [31] OBERLANDER, J., O'DONNELL, M., MELLISH, C., AND KNOTT, A. Conversation in the museum: experiments in dynamic hypermedia with the intelligent labelling explorer. *The New Review of Hypermedia and Multimedia* 4 (1998), 11–32.
- [32] STOCK, O., ZANCANARO, M., Busetta, P., Callaway, C., KRÜGER, A., KRUPPA, M., KUFLIK, T., NOT, E., AND ROCCHI, C. Adaptive, intelligent presentation of information for the museum visitor in PEACH. *User Modeling and User-Adapted Interaction* 17 (2007), 257–304. 10.1007/s11257-007-9029-6.
- [33] WANG, Y., STASH, N., AROYO, L., GORGELS, P., RUTLEDGE, L., AND SCHREIBER, G. Recommendations based on semantically enriched museum collections. *Web Semant.* 6, 4 (November 2008), 283–290.
- [34] DE BRA, P., SMITS, D., AND STASH, N. Creating and delivering adaptive courses with AHA! In *Proceedings of the First European Conference on Technology Enhanced Learning - EC-TEL'06* (Crete, 2006), vol. 4227, Springer, pp. 21–33.
- [35] HSIAO, I.-H., SOSNOVSKY, S., AND BRUSILOVSKY, P. Guiding students to the right questions: adaptive navigation support in an e-learning system for java programming. *Journal of Computer Assisted Learning* 26, 4 (2010), 270–283.

- [36] CONLAN, O., HOCKEMEYER, C., WADE, V., AND ALBERT, D. Metadata driven approaches to facilitate adaptivity in personalized elearning systems. *The Journal of Information and Systems in Education 1* (2002), 38–44.
- [37] HENZE, N., AND NEJDL, W. Adaptivity in the KBS Hyperbook system. In *2nd Workshop on Adaptive Systems and User Modeling on the WWW, workshop held in conjunction with the WWW Conference (WWW8) and the International Conference on User Modeling* (Toronto, CA, 1999).
- [38] HENDRIX, M., AND CRISTEA, A. I. Design of the CAM model and authoring tool. In *4th European Conference on Technology-Enhanced Learning* (2009), EC-Tel 2009.
- [39] BALÍK, M., AND JELÍNEK, I. Towards semantic web-based adaptive hypermedia model. In *PhD Symposium at the European Semantic Web Conference (ESWC)* (Tenerife, ES, 2008).
- [40] DE VRIEZE, P. *Fundamentals of Adaptive Personalization*. PhD thesis, Radboud University Nijmegen, The Netherlands, 2006.
- [41] DE BRA, P., SANTIC, T., AND BRUSILOVSKY, P. AHA! meets Interbook and more... In *Proceedings of the AACE ELearn 2003 Conference* (2003), pp. 57–64.
- [42] SEEFELDER DE ASSIS, P., AND SCHWABE, D. A general meta-model for adaptive hypermedia systems. In *Proceedings of the 3rd International Conference on Adaptive Hypermedia and Adaptive Web-Bases Systems* (2004), P. De Bra and W. Nejdl, Eds., AH '04, pp. 433–436.
- [43] CHIN, D. N. Empirical evaluation of user models and user-adapted systems. *User Modeling and User-Adapted Interaction 11*, 1-2 (2001), 181–194.
- [44] PARAMYTHIS, A., WEIBELZAHN, S., AND MASTHOFF, J. Layered evaluation of interactive adaptive systems: framework and formative methods. *User Modeling and User-Adapted Interaction 20* (December 2010), 383–453.
- [45] BOYLE, C., AND ENCARNACION, A. O. An adaptive hypertext reading system. *User Modeling and User-Adapted Interaction 4* (1994), 1–19. 10.1007/BF01142355.

- [46] MEYER, B. Retail user assistant: Evaluation of a user-adapted performance support system. In *Human-Computer Interaction*, B. Blumenthal, J. Gornostaev, and C. Unger, Eds., vol. 876 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 1994, pp. 45–55. 10.1007/3-540-58648-2_25.
- [47] WEBER, G., AND SPECHT, M. User modeling and adaptive navigation support in WWW-based tutoring systems. In *International Conference on User Modeling (1997)*, Springer, pp. 289–300.
- [48] HÖÖK, K. Evaluating the utility and usability of an adaptive hypermedia system. *Knowledge-Based Systems 10*, 5 (1998), 311–319. Intelligent User Interfaces.
- [49] RAMOS, V. F. C., AND DE BRA, P. The influence of adaptation on hypertext structures and navigation. In *Proceedings of the 21st ACM conference on Hypertext and hypermedia* (New York, NY, USA, 2010), HT '10, ACM, pp. 77–82.
- [50] PARAMYTHIS, A., TOTTER, A., AND STEPHANIDIS, C. A modular approach to the evaluation of adaptive user interfaces. In *Proceedings of the Workshop on Empirical Evaluations of Adaptive Systems, held in the context of the 8th International Conference on User Modeling (UM'2001)* (Sonthofen, Germany, 2001), S. Weibelzahl, D. Chin, and G. Weber, Eds., pp. 9–24.
- [51] WEIBELZAHN, S., AND WEBER, G. Evaluating the inference mechanism of adaptive learning systems. In *Proceedings of the 9th international conference on User modeling* (Berlin, Heidelberg, 2003), UM'03, Springer-Verlag, pp. 154–162.
- [52] PARAMYTHIS, R., AND WEIBELZAHN, S. A decomposition model for the layered evaluation of interactive adaptive systems. In *Proceedings of the 10th International Conference on User Modelling, Lecture Notes in Artificial Intelligence 3538* (2005), Springer, pp. 438–442.
- [53] KITCHENHAM, B. A., PFLEGER, S. L., PICKARD, L. M., JONES, P. W., HOAGLIN, D. C., EMAM, K. E., AND ROSENBERG, J. Preliminary guidelines

- for empirical research in software engineering. *IEEE Transaction Software Engineering* 28 (August 2002), 721–734.
- [54] YANCEY, J. M. Ten rules for reading clinical research reports. *American Journal of Orthodontics & Dentofacial Orthopedics* 109, 5 (1996), 558–564.
- [55] WELCHE, G. E., AND GABBE, S. G. Review of statistics usage in the American journal of obstetrics and gynecology. *American Journal of Obstetrics and Gynecology* 175, 5 (1996), 1138–1141.
- [56] MCGUIGAN, S. The use of statistics in the British journal of psychiatry. *The British Journal of Psychiatry* 167, 5 (1995), 683–688.
- [57] GENA, C. Methods and techniques for the evaluation of user-adaptive systems. *Knowledge Engineering Review* 20, 1 (2005), 1–37.
- [58] GENA, C., AND WEIBELZAHN, S. The adaptive Web. In *Lecture Notes in Computer Science*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds. Springer-Verlag, Berlin, Heidelberg, 2007, ch. Usability engineering for the adaptive Web, pp. 720–762.
- [59] VAN VELSEN, L., VAN DER GEEST, T., KLAASSEN, R., AND STEEHOUDER, M. User-centered evaluation of adaptive and adaptable systems: a literature review. *The Knowledge Engineering Review* 23, 03 (2008), 261–281.
- [60] DE JONG, M., AND SCHELLENS, P. J. Reader-focused text evaluation. *Journal of Business and Technical Communication* 11, 4 (1997), 402–432.
- [61] WEIBELZAHN, S. Evaluation of adaptive systems. In *8th International Conference on User Modeling* (Berlin, 2001), vol. 2109, Springer, pp. 292–294.
- [62] ROMERO, C., VENTURA, S., JOSE ANTONIO, D., AND DE BRA, P. Personalized links recommendation based on data mining in adaptive educational hypermedia systems. In *Creating New Learning Experiences on a Global Scale. Second European Conference on Technology Enhanced Learning, EC-TEL 2007* (2007), Springer, pp. 293–305.

- [63] BOTAFOGO, R. A., RIVLIN, E., AND SHNEIDERMAN, B. Structural analysis of hypertexts: identifying hierarchies and useful metrics. *ACM Trans. Inf. Syst.* 10, 2 (April 1992), 142–180.
- [64] GRIMMETT, G., AND STIRZAKER, D. *Probability and Random Processes*. Oxford University Press, August 1992.
- [65] RAMOS, V. F. C., DE BRA, P., AND XEXÉO, G. Qualitative and quantitative evaluation of an adaptive course in GALE. In *Scaling up Learning for Sustained Impact*, D. Hernandez-Leo, T. Ley, R. Klamma, and A. Harrer, Eds., vol. 8095 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2013, pp. 301–313.
- [66] RAMOS, V. F. C., DE BRA, P., AND XEXÉO, G. B. A continuous layered evaluation of adaptive hypermedia systems (chapter 8). In *Evaluation in e-Learning*, Y. Psaromiligkos, A. Spyridakos, and S. Retalis, Eds. Nova Publishers, 2012, pp. 145–158.
- [67] KARAGIANNIDIS, C., AND SAMPSON, D. G. Layered evaluation of adaptive applications and services. In *AH '00: Proceedings of the International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (London, UK, 2000), Springer-Verlag, pp. 343–346.
- [68] WEIBELZAHN, S., AND WEBER, G. Advantages, opportunities, and limits of empirical evaluations: Evaluating adaptive systems. *Künstliche Intelligenz* 3 (2002), 17–20.
- [69] ORTIGOSA, A., AND CARRO, R. M. The continuous empirical evaluation approach: evaluating adaptive web-based courses. In *UM'03: Proceedings of the 9th international conference on User modeling* (Berlin, Heidelberg, 2003), Springer-Verlag, pp. 163–167.
- [70] ABEL, F., HENZE, N., HERDER, E., AND KRAUSE, D. Interweaving public user profiles on the web. In *Proceedings of the 18th international conference on User Modeling, Adaptation, and Personalization* (Berlin, Heidelberg, 2010), UMAP'10, Springer-Verlag, pp. 16–27.

- [71] TROYER, O., KLEINERMANN, F., PELLENS, B., AND EWAIS, A. Supporting virtual reality in an adaptive web-based learning environment. In *Learning in the Synergy of Multiple Disciplines*, U. Cress, V. Dimitrova, and M. Specht, Eds., vol. 5794 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2009, pp. 627–632.
- [72] VAN DER SLUIJS, K., HIDDERS, J., LEONARDI, E., AND HOUBEN, G.-J. GAL: A generic adaptation language for describing adaptive hypermedia. In *Proceeding of the International Workshop on Dynamic and Adaptive Hypertext: Generic Frameworks, Approaches and Techniques (DAH'09) in conjunction with ACM Hypertext 2009* (2009), DAH'09, pp. 13–24.
- [73] PLOUM, E. Authoring of adaptation in the GRAPPLE project. Master Thesis, Eindhoven University of Technology, The Netherlands, 2009.
- [74] DEVDUTTA, B. AlcoZone: An adaptive hypermedia based personalized alcohol education. Master Thesis, Electrical and Computer Engineering Department, Virginia Tech, USA, 2006.
- [75] RAMOS, V. F. C., BRA, P., AND SMITS, D. GALE extensibility evaluation: a qualitative approach. In *Proceedings of the World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education (ELEARN) 2013* (Las Vegas, NV, USA, 2013), T. Bastiaens and G. Marks, Eds., AACE, pp. 296–305.