RC23662 (W0507-077) July 8, 2005
Computer Science

IBM Research Report

Dynamic Learning Experience: A System for Just-in-Time Learning

Robert Farrell, Ernest Fuller, Raymond O'Donnell, Yael Ravin, Steve Rolando, John C. Thomas
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Dynamic Learning Experience: a system for just-in-time learning

Robert Farrell, Ernest Fuller, Raymond O'Donnell, Yael Ravin,
Steve Rolando, John C. Thomas
IBM
({robfarr, eafuller, rodonnel, ravin, stever, jcthomas} @us.ibm.com)

Abstract

The report describes the work of a joint team of researchers and training practitioners who developed and validated Dynamic Learning Experience (DLE), a system for just-in-time learning. DLE defines a novel approach to content re-use and a new delivery model. The team created a software tool and a methodology for turning technical manuals and informational presentations into modular learning objects. A Web-based learning environment was implemented and linked to the repository of learning objects. We tested the system in two pilot implementations within IBM, with very positive user feedback. We also conducted a controlled study and found that users of DLE performed significantly better on a work-related task than users of a standard search engine.

Keywords
e-learning, learning objects, just-in-time learning, dynamic learning experience

Introduction

In today's volatile and highly competitive economic climate, pressure is increasing on people to acquire new skills and rapidly absorb information. As a consequence, even if workers are highly motivated to educate themselves on new topics, they do not have the time to take a full course of instruction. Instead, they often turn to the Web or the intranet and search in an ad-hoc way for what they need. Information is difficult to organize for effective learning. The knowledge acquired is often disconnected, and therefore easily forgotten and "dis-integrated" with practice.

We believe that a new approach is required – one which leverages knowledge workers' high motivation to actively participate in their own learning while addressing their time constraints and the limitations of current search mechanisms. We base this view in part on known observations that adult professionals under time pressure often like being in control of their learning (Brookfield, 1986). In this paper we describe Dynamic Learning Experience (DLE), a system and methodology we created to test this approach. DLE enables content creators to extract information from primary document sources and quickly create short, self-contained web-based educational materials as learning objects and associate them with descriptive metadata. Using DLE, learners can then search for relevant learning objects, assemble them into individualized instructional sequences, and experience the sequences as "custom courses."

We hypothesize that if learners are able to construct their own learning paths from modular learning objects, there would be a number of benefits:

- Increased motivation to learn, based on self-directed learning and increased emphasis on meta-cognitive skills, such as goal setting and planning;
- Improved job satisfaction, when learner feels trusted and empowered;
- Improved effectiveness on job tasks, as courses created dynamically can be tailored to address specific learning gaps. Learners faced with significant task demands know more about their own skill deficiencies and knowledge gaps than others who are not familiar with their work context.
- Better completion than the typical rate of under 50 percent for online courses, because the courses focus the learner on a small set of resources relevant to a given topic;
- Lower cost to develop material, based on modularization and reuse of learning objects.

In the following sections we describe the system we built, the population we addressed, the pilot tests and the controlled study we conducted to confirm our hypotheses and measure the value of the system.

The Dynamic Learning Experience System

Using *Dynamic Learning Experience*, users identify learning objects of interest using an on-line search feature (Figure 1). They can query by keywords (e.g., "e-business", "languages"), intended uses for the desired learning material within an instructional context (e.g., introduction, procedures), level of difficulty (e.g., easy, difficult) and other features. Relevant learning objects that are returned by the search[1] can then be individually selected and assembled to form a custom course, as shown in Figure 2. Alternatively, a "Dynamic Assembly" option can automatically select a coherent subset of learning objects to meet the desired course duration.

The system arranges the selected learning modules in a logical, pedagogical sequence of lessons, determined by a uniform sequencing policy: topics are arranged in a logical learning progression (see Farrell, Liburd & Thomas 2004), and within each topic, rhetorical and instructional principles are used to order modules from introduction through conclusion.

Custom courses can be easily shared with any user on the Intranet by sending an e-mail with an embedded hyperlink leading directly to the online course. Users can begin learning right away. First-time users are requested to confirm profile information. Shared courses are copied to each user's individual personal course catalog stored on the server, so they can be further personalized by the recipient. The lifecycle metadata indicating the original author and assembly date are stored to maintain some degree of attribution to the originator.

[1] For more detail on the XML search engine adapted for use in this system, see 02002.

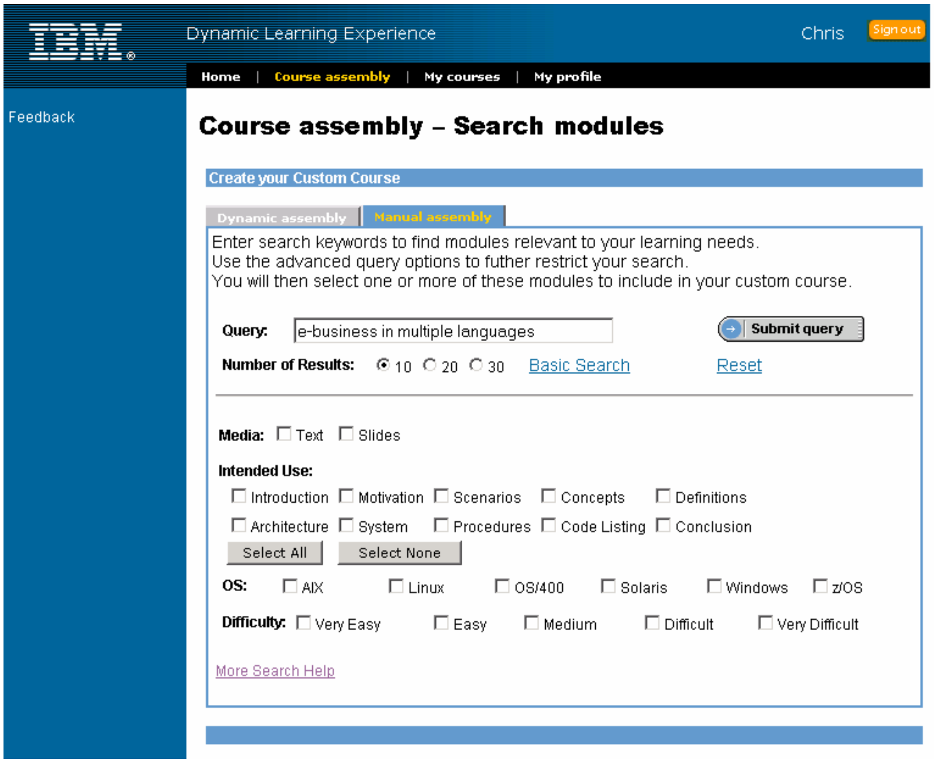

Figure 1: The Search Page

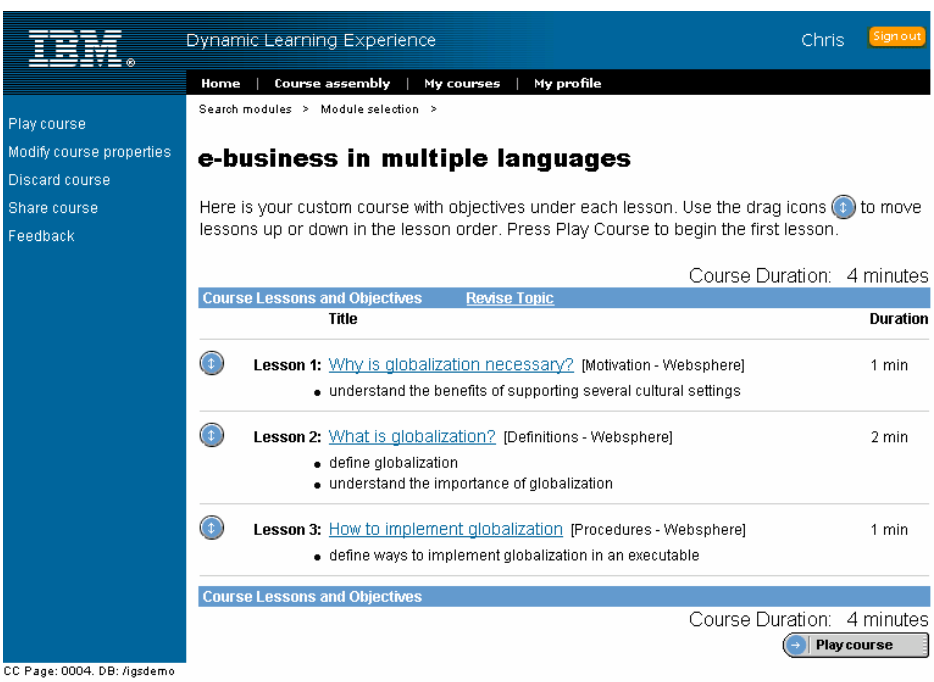

Figure 2: A Custom Course

The IT Practitioners Targeted for Study

We focused on a subset of IBM's Global Services practitioners, a population of about 150,000 professionals providing information technology (IT) services to customers in 160 countries. They represent adult, motivated, learners with some IT background who have specific educational needs that are job-related. They are learners who are motivated to find, organize, and comprehend.

We formed a joint team with members from the IBM Global Services Learning (IGSL) organization and from IBM's Research Division. IGSL is an organization that provides recommendations and programs for continuous learning and sustained professional development to IBM Global Services practitioners. IBM Research is an IT research organization, with more than 3000 scientists and engineers.

In 2002, IGSL defined a "speed to knowledge" initiative, to broaden and deepen employees' knowledge in key areas of the business. The intent was to quickly provide access to small amounts of specific information based on a search of all available content. After testing a few products and prototypes available on the market, IGSL concluded that their ideal solution should include not only learning modules, but also a mechanism for modularizing existing non-modular content, as well as a more sophisticated search mechanism. This provided the context for the creation of DLE.

Pilot Tests of Dynamic Learning Experience

Content Creation

We chose content about Web Services and the IBM WebSphere product, both of great interest to the IBM technical community. For the first pilot, we used seven "how-to" manuals known as Redbooks (http://www.redbooks.ibm.com), comprising of over a million words of text and over two thousand images. Based on the table of contents, we automatically segment the books into modules, pre-populated with titles and an estimated study time. Experts review the suggested modules and add metadata using a spreadsheet-like tool. After several trial runs with experts, we decided on a linear model whereby any entry in the table of contents could start a module. This model was easy for our subject matter experts to understand and apply. (But see Douglas 2001 for other models of learning object development.) The additional metadata entered by subject matter experts include:

- learning objectives
- a difficulty level on a five-point scale
- intended use fields to capture the rhetorical role of this section (e.g., introduction versus conclusion), the cognitive level of the material (e.g., concepts versus procedures), and/or the specific resource type (e.g., architecture diagram versus a listing of definitions)
- intended audience (a primary and secondary job role in the organization)
- relevance to a development environment, operating environment, or product.

The information from automatic derivation and manual entry is then combined into an IEEE standard Learning Object Metadata file.

For the second pilot test, we expanded the coverage to a total of five hundred learning objects, created out of sixteen Redbooks. We extended our tools to process information from slide presentations, and included material from seven slide presentations, with a total of 2900 slides. We added eleven high-level topics to help organize information across the learning material. These included issues (e.g., security) and deployment steps (e.g., debugging, administration) that are common across multiple products. These topics became a primary point of semantic integration across the diverse set of source of materials. Finally, we added comprehension and inference questions to each module, to improve retention.

Pilot tests

The goal of the first pilot was to assess the utility and impact of converting existing knowledge content into small learning modules assembled into custom courseware by practitioners. The learning environment was instrumented with feedback and logging features. Data were correlated across these features to provide a deeper insight into the learning experiences of individual users. A thorough analysis of the log data was performed to look at patterns and trends. The goal of the second pilot was to determine if employees would spontaneously use the system for focused learning and to determine how the system was actually being used by employees.

Results

A total of 114 users from across the company signed up for the first pilot, which ran for one month. Of these, seventy-four users used the system for at least an hour. All users filled out evaluation forms. Our evaluation shows that users were highly satisfied with the system. Eighty-one percent answered with a positive rating (4 or 5 out of 5) to the question "What is your overall satisfaction with this method of learning?" Users found the system easy to use, fast, and selective. Eighty-one percent reported that this system would enhance their knowledge/skills. Fifty-two percent said they would prefer this method of learning to other methods, although many said they would prefer classroom instruction if it were available. Only one user preferred working with the technical manuals directly. Over ninety percent found creating and navigating custom courses easy to do. During the pilot, users provided twelve positive comments and only two negative ones. These were very encouraging results.

In the second trial, 250 employees from twenty-five countries used the system. One-hour interviews were conducted with a focus group of eighteen. Users reported highly favorable experiences with the system. All wanted to know when the system would be available on a production basis. At least one person used the system to prepare presentations for a customer engagement. Several people used the system as a reference tool as they were performing work. One person used the system to generate advanced organizers for further study. Users reported that the system helped them to quickly come up to speed on their selected topic of interest. Without exception, all of the users thought the concept behind the system was sound.

Discussion

A detailed analysis of the logs revealed that many users were spending a lot of time trying to find modules on particular subjects. We believe that some of the failures were due to lack of material on the topics they were interested in. We addressed this in the second pilot by substantially extending the coverage of the system.

The data also suggest that users often made their learning choices based upon surface features rather than deeper process or relational features, a potential problem if we want to drive users toward deeper understanding. We felt these problems might be alleviated if the dynamic organization of modules better reflected a logical learning path. Thus, we introduced more high-level topics for a better logical organization of modules across data sources.

Controlled Study

It was unclear at the conclusion of the pilots whether the system was superior to a regular search engine used with modular content. This comparative value is particularly important to establish given that human labor is involved in the addition of metadata required for sequencing. We conducted a controlled study to compare the performance on a work-related task of two groups – a control group using a standard search engine and an experiment group using DLE against the same set of learning objects.

Experimental Design

Twenty-six subjects participated. Most were classified as "IT Specialist": consultants, sales specialists, managers, etc. There was a great deal of variability in the breadth and depth of their general programming background, years of experience, as well as in their pre-existing knowledge. Approximately half the subjects were randomly assigned to each group, roughly comparable on these measures. Almost all subjects were remote and "observed" by an experimenter who followed their actions on the Web site and listened to their thinking on a conference call. In the "Query Only Condition," subjects used the search engine to enter query terms and then browse through the materials returned, take notes, and cut and paste. In the "Custom Course Condition" the system also assembled learning objects into courses. The metadata in the two conditions differed, with the Custom Course group able to see metadata for each module (learning objectives, description, difficulty, and duration). The major interaction modes in the two conditions were quite different. The Query Only condition was similar to a typical search engine, while the Custom Course condition afforded additional steps of deciding which search results were relevant based upon the available metadata and asking the system to rearrange the search results into a course.

Subjects were given a scenario-based task to motivate them to learn. The scenario was extremely challenging (to help avoid ceiling effects) but also open to some interpretation so that subjects would have a great deal of latitude in the materials they chose. Subjects were asked to construct a high level design for turning an application into a Web Service. They had about an hour to study before working on the task. Their designs were quantified and evaluated in several ways: the total output (number of pages, boxes, arrows, etc); the similarity of their solution to a solution designed by an expert; and a "grade" assigned by three experts in a manner blind to condition.

Results

We found positive significant correlations between the Custom Course Condition and all the quantitative measures. Full statistics can be found in Thomas & Farrell (in publication). The Custom Course group produced more design behavior that the Query Only group and their solutions had more common terms with the expert solution. There was a significant correlation between the Custom Course Condition and qualitative measures as well. Custom Course designs received significantly higher grades from the experts, who tended to agree overall on quality.

In all the measures there was significant correlation between self-reported experience and measures of quality. We had to "correct" for this effect by looking at residuals.

We used several measures of behavior as well. In one, we counted the number of words used by each group to reflect cognition ("learn", "know", "think") versus words reflecting search ("get", "find"). The vocabulary distribution may be indicative of the cognitive experience of the subjects: the Custom Course group concentrated on learning, while the Query Only group concentrated on searching.

Applications of DLE

Our data suggest that DLE could be useful for improving performance on a wider range of tasks requiring time-constrained comprehension, such as those typically faced by technical professionals when studying for a comprehensive certification or preparing a detailed technical presentation. Users noted that many job groups could benefit from this approach, such as consultants, who often need to find out about a wide variety of topics quickly, or sales people, who need to find out about competitive offerings.

The results also indicate that the system would be useful for project teams. Many users specifically liked the custom course sharing feature, as a way of having a project team communicate and develop shared understanding, especially since teams are often global and team members are remote from each other. Finally, other users suggested that custom courses would be quite useful when someone changes jobs and has to come up to speed quickly on a new product or product family.

Users also reported that they were often using the system in a "multi-tasking" or "interrupt-rich" environment and that the system was responsive to this situation.

Limitations of DLE

The range of experience in each of the two groups was considerable, suggesting that we need to run a much larger number of subjects in order to minimize the effect of these determinents of performance. In addition, any of the following factors might reduce the value of the system to users:

- learner who are not motivated to learn or want to "game" the system
- learners who know very little about the field and cannot formulate good queries
- learners who have no familiarity with computers or how search engines work
- learning goals that involve a fundamental rethinking of belief structure
- a field of study with a heavy emphasis on acquiring new sensory-motor skills
- a field of study that requires a heavy emphasis on interpersonal, social, and/or imaginative skills; e.g., negotiation, debate, or creative writing.

None of these limitations is probably fundamentally insoluble, but each would require additional functionality added to the system and/or the context of use of the system.

References

Brookfield, S. D. (1986) Understanding and Facilitating Adult Learning. San Francisco: Jossey-Bass.

Carmel, D., et al. (2002) An Extension of the Vector Space Model for Querying XML Documents via XML Fragment, The 25th Annual ACM Special Interest Group in Information Retrieval (SIGIR) Conference, Tempere, Finland.

Douglas, I. (2001) *Instructional Design based on Reusable Learning Objects: Applying Lessons of Object-oriented Software Engineering to Learning Systems Design*. The 31st ASEE/IEEE Frontiers in Education Conference. Reno, NV.

Farrell, R. et al. (2003) *Learner-driven Assembly of Web-based Courseware,* Proceedings of E-Learn 2003, Phoenix AZ.

Farrell, R., Liburd, S.D., and Thomas, J. C. (2004) Dynamic *Assembly of Learning Objects*. Proceedings of the World-Wide Web Conference, New York, NY.

Thomas J. C. and R. Farrell in publication) An Experimental Investigation of the Effectiveness of Dynamic Assembly of Learning Objects

Yael Ravin, PhD, leads IBM's research in learning. Before assuming this position, Ravin managed research in knowledge management. She has published two books, holds several patents, and has presented at several international conferences.

Contact for Robert Farrell, Yael Ravin and John Thomas:
T. J. Watson Research Center
19 Skyline Drive
Hawthorne, NY 10532, USA