

Understanding Digital Library Adoption: A Use Diffusion Approach

Keith E. Maull
Institute of Cognitive Science
Dept. of Computer Science
University of Colorado
Boulder, Colorado 80309
maull@colorado.edu

Manuel Gerardo Saldivar
Institute of Cognitive Science
School of Education
University of Colorado
Boulder, Colorado 80309
saldivar@colorado.edu

Tamara Sumner
Institute of Cognitive Science
Dept. of Computer Science
University of Colorado
Boulder, Colorado 80309
sumner@colorado.edu

ABSTRACT

With the growth in operational digital libraries, the need for automatic methods capable of characterizing adoption and use has grown. We describe a computational methodology for producing two, inter-related, user typologies based on use diffusion. Use diffusion theory views technology adoption as a process that can lead to widely different patterns of use across a given population of potential users; these models use measures of frequency and variety to characterize and describe these usage patterns. The methodology uses computational techniques such as clickstream entropy and clustering to produce both coarse-grained and fine-grained user typologies. A case study demonstrates the utility and applicability of the method: it is used to understand how middle and high school science teachers participating in an academic year-long field trial adopted and integrated digital library resources into their instructional planning and teaching. The resulting fine-grained user typology identified five different types of teacher-users, including “interactive resource specialists” and “community seeker specialists.” This typology was validated through comparison with qualitative and quantitative data collected using traditional educational field research methods.

Categories and Subject Descriptors

H.1.2 [Information Systems]: Models and Principles—*User/Machine Systems*; H.3.7 [Information Systems]: Information Storage and Retrieval—*User Issues*; J.1 [Computer Applications]: Administrative Data Processing—*Education*

General Terms

Experimentation, Human Factors

Keywords

Technology adoption, diffusion of innovation, use diffusion models, educational digital libraries

1. INTRODUCTION

Educational digital libraries have evolved from being primarily research-oriented enterprises to encompass a large number of operational library sites, such as NSDL, Merlot, and DLESE in the US, SchoolNet in Europe, and the National Digital Learning Resources Network in Australia, to name just a few. As library efforts continue to mature, there is a growing need for efficient and scalable methods to characterize their uptake and adoption, their impact on teacher and student practices, and ultimately their impact on student learning. In this article, we describe a computational methodology for automatically identifying and characterizing different patterns of digital library adoption and use.

This methodology instantiates a particular theoretical model of technology adoption, use diffusion [22], which in turn builds on prior work on the diffusion of innovation [20]. Diffusion of innovation is one of the most researched and widely employed social science models; it has been used to study the adoption of agriculture innovations such as new corn varieties [20], health innovations such as water purification and disease treatments [ibid], and the very rapid adoption of digital consumer products [17]. Diffusion of innovation theory provides a lens for understanding the different factors that influence a person’s decision to use, or not use, an innovation, and when in the product lifecycle they might adopt. For instance, an “early adopter” farmer might be motivated by the thought of potential harvest gains to be the first farmer in the region to try out a new corn variety, whereas a “late majority” farmer would wait until most of the surrounding farms had already adopted the new corn. Characterizing the adoption of contemporary information services, such as educational digital libraries, however, is more complex than identifying when the farmer planted the corn or when the consumer bought the digital device. Did a teacher that used NSDL or MERLOT one time ‘adopt’ the library? In acknowledgement of this complexity, instead of focusing on when an innovation is first used, use diffusion examines both *how* and *how much* an innovation is used to identify different adopter categories. It recognizes that both the depth and breadth of usage will vary widely across different users, and that successful adoption will take many different forms.

Building on techniques from web analytics and data mining, the proposed computational methodology employs two different algorithms, in a two-step process, to develop both coarse-grained and fine-grained views of user behavior. These

algorithms rely on detailed web site usage logs, where each individual action in the interface is recorded and associated with a unique user identifier, for their input. In the first step, one algorithm uses frequency of use and variety of use to sort users into different use diffusion adopter quadrants, such as “intense use”, “limited use”, “specialized use”, etc. In the second step, a clustering algorithm is employed to develop a finer-grained understanding of the different patterns of generalized and specialized use within and across these quadrants. One challenge with operationalizing use diffusion in a computational method is modeling variety in a way that is application independent; we chose to use Shannon entropy [21], a mathematical construct from information theory, to model variety. The rationale for this choice, as well as the underlying mathematics, is described in detail later in this article.

To illustrate both the utility of the computational methodology and how to apply it in practice, we present a case study where we used this methodology to better understand how middle and high school science teachers integrated digital library resources from NSDL and DLESE into their instructional planning and teaching practice. Teachers were provided with a web-based planning tool that enabled them to customize their district’s adopted curriculum with digital library resources to better meet diverse learner needs and to share their customizations with other teachers in their district [25]. This planning tool, the Curriculum Customization Service, was deployed to all middle and high school Earth science teachers within a large urban school district in the Midwest US for a full academic year; the Service was carefully instrumented to record every user action and detailed usage logs were collected. This case offers an excellent testbed for this methodology for two reasons. First, the total potential user population is known and quantifiable, enabling us to more easily assess overall rates of uptake and adoption. Second, this deployment was also studied using traditional educational field research methods such as surveys, interviews, and classroom observations. Thus, we can compare the coarse-grained and fine-grained views of user behavior identified by the computational methodology with findings from the field study to better assess the accuracy and validity of the methodology’s output. While this case is clearly focused on understanding educational use of digital libraries, we believe the proposed methodology to be generalizable and applicable to a wide variety of digital libraries, information services, and learning applications.

2. BACKGROUND AND RELATED WORK

This research draws on theories and computational techniques from several disciplines in order to better understand digital library adoption use. First, as previously described, we draw on technology adoption and diffusion theories, which are historically rooted in social science, to inform the purpose and overall functioning of our computational methodology. Second, we discuss related research to understand user behavior and user typologies, describing how our approach compares to other efforts that are similarly focused on developing automatic methods.

2.1 Adoption and Diffusion Theories

Technology adoption occurs when an individual decides that a given technological innovation has utility and can add

value to his or her activities – such as teaching – if that innovation is somehow incorporated into those activities [24]. Thus, much theoretical work to date has focused on understanding the cognitive, affective, and contextual factors that influence a potential user’s decision-making process. One prominent family of theories offers extensions or refinements of Roger’s innovation diffusion theory [20, 6, 18]. To Rogers, technology adoption is fundamentally a function of the communication channels and social systems of which one is a part. This theory suggests that (1) within a social system there are typically five different “adopter categories” describing the different characteristics that users bring to bear when considering whether to adopt an innovation, and (2) that these characteristics influence how innovations move through the social system in predictable ways. Another model is the concerns-based adoption model developed by Fuller [10] and Hall [11]. As its name implies, this model is focused on individuals’ concerns, which are defined as the specific reasons – situated in one’s socio-cultural context – that one might have to adopt or not adopt technology. A third major approach is the technology acceptance model, or TAM [5, 27]. This was one of the first models to take into account individuals’ self-efficacy and expertise. Proponents of this model argue that individuals’ self-perception of their ability to use technology and their ability to judge whether a technology has utility for them are important factors for understanding technology adoption behaviors.

These prior models have contributed greatly to our understanding of technology adoption; however, they all share a couple of weaknesses. First, none of them take into account discontinuance; i.e., people often stop using a new technology after they have tried it out a few times. Second, they provide very little insight into actual use of the new technology. Since most of these methods rely on self-reported survey data, they more often predict variance in self-reported use rather than actual use [26]. Thus, more recently, theoretical attention has shifted from focusing on the decision component of adoption towards understanding adoption as a process that can lead to different patterns of use. Models that account for use fall into the category of “use diffusion” models. These models attempt to characterize the way, and the degree to which, people make use of the new technology. For example, once a consumer purchases a cell phone, to what degree does he or she actually use the phone? What features are used and how are these features used? Can different types of usage be recognized and compared?

Formally proposed by Ram and Jung [19] and subsequently updated and expanded by Shih and Venkatesh [22] to accommodate more robust and predictable descriptions of usage, use diffusion extends the traditional notion of adoption diffusion by focusing on system usage patterns. The work in this paper builds on the Shih and Venkatesh use diffusion model, which suggests two dimensions to patterns of technology use. The first dimension, frequency, provides a measurement of how much a technology is used. In the web context, for example, frequency of web site use might be defined as the number of sessions a user has generated over some period of time. This frequency measure can be a very useful indicator of their interest in site content. The second dimension of the use diffusion model is what the authors call variety. This dimension measures the range of use of

a technology; did the consumer make use of most of the features of their new cell phone, or only two or three? In the web context, unlike frequency and number of sessions, there are no standard measures of variety. For this study we model variety as clickstream entropy. The use diffusion model thus produces four categories of use of a technology as shown in Figure 1. When plotted along these two dimensions, the population of users is thus segmented into these four “adopter categories”: intense use, limited use, specialized use, and non-specialized use. A user may move to and from adopter categories depending, for example, on the time interval considered and the granularity of the data used for analysis (e.g. an individual session versus lifetime sessions).

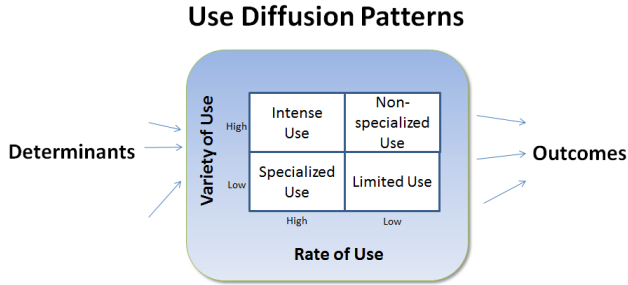


Figure 1: Use diffusion model proposed by Shih and Venkatesh [22]

2.2 User Typology Modeling

The outputs of our computational methodology are two, inter-related, user typologies: (1) a course-grained view of the user population segmented into use diffusion adopter categories and (2) a fine-grained view of the same population segmented along the same two dimensions but using more detailed measures for variety and frequency. Classifying and categorizing users into groups is a common task in user behavior analysis. The output of adoption models is often a set of adopter categories, which are a particular type of user typology. The categories produced by the use diffusion model are analogous to those produced by Roger’s diffusion of innovation model (i.e., “early adopter” or “late majority”).

Recently, a great deal of user typology research has been directed towards understanding and classifying Internet users, the common tasks they perform, and the details of their online behaviors. The Pew Research Group, for example, is developing user typologies describing the technology and Internet use patterns of Americans [14]. In their typology distinctions are made, for example, between “Light But Satisfied” users, individuals who use some technology but for whom technology does not play a central role in day-to-day life, and “Omnivores,” who embrace technology fully and participate heavily in online activities. Typologies of media users have been extensively explored by Brandtzæg [3], and with the rise of social media systems, user typologies within this context are a growing area of research [2]. In education research, recent work by Eynon and Malmberg illustrates how typologies can inform design and implementation rec-

ommendations: they are using Internet usage typologies of young students to more effectively integrate new technologies into classroom practices [9].

Other research is focused on using computational methods for identifying user typologies. Within educational data mining, clustering algorithms are being used to group students into different categories based on skills sets [1] or performance on a test or assessment [8]. Xu, et al. use clustering to identify and classify usage types of teachers [30]. They examine features of teacher-generated projects within the Instructional Architect tool to create a typology of users based on the kinds of projects that they produced. As these examples illustrate, clustering algorithms are generally used to assign group membership among items with common attributes or features in large data sets. The computational methodology presented in this article also uses clustering to identify and classify usage types of teachers. This research differs from prior efforts in that the features selected for the clustering algorithm are theoretically motivated by the use diffusion model.

3. COMPUTATIONAL METHODOLOGY

The two-step methodology for this research is constructed to discover usage patterns and user typologies. The first step captures coarse-grained user categories, while the second step determines fine-grained typologies of system use. The next two sections will examine the details of each step, their inputs, processes and outputs.

3.1 Step 1: Use Diffusion Patterns

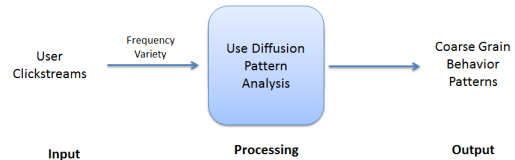


Figure 2: Step One Overview

To understand how use diffusion patterns are modeled, it is important to more fully examine the frequency and variety dimensions of the model. In this work, frequency is modeled as the number of user-initiated web sessions. While other frequency measures can be considered, this measure provides a good initial approximation of overall system use: fewer sessions imply lower system use, while more sessions imply higher system use. Variety, on the other hand, is more challenging to model because it is difficult to develop application-independent approaches to the concept. For the first step, we chose a variety metric that is based on aggregate user clickstreams. Intuitively, the clickstream of a particularly user approximates their broad usage of the system. Furthermore, *over time* clickstreams become regular – that is they become more predictable as users develop normal patterns of use within the system. By applying entropy – and specifically Shannon entropy [21] – over the lifetime clickstream of each user, a basic notion of variety is developed that gives an approximate measurement of user behavior. Entropy has been used extensively in many systems to calculate measures of randomness and to approximate the amount of information being communicated in a system.

The clickstream entropy calculation borrows from the idea that low-entropy models are considered more predictable than high-entropy models. Thus, clickstreams with low entropy imply low variety, which when applied in the context of system usage provide us a variety measure that is robust, while at the same time application-independent, generalizable, and extensible. Finally, entropy calculations are computationally trivial, making them both easy to compute and compare. Now that we have the domain independent model in place for describing variety, let us further motivate the idea for clickstream entropy in the web usage context. For a given website, a user generates a path of click interactions through that site. We use entropy to describe such a clickstream, with the primary interest in developing a measure of predictability from within a site. Our goal is not to *predict* the next click, rather we are interested in a coarse-grained metric on how *hard* it would be to predict the next click. The intuition behind this is simple: if entropy can model the lower bound on the predictability of a clickstream, then such predictability could be used analogously to assess the variety of a clickstream. We will now complete the mathematical preliminaries for entropy calculations in the clickstream case.

Let $S_{\alpha_k}^{u_i}$ be a clickstream of urls α_k for a given user u_i 's session on a website. Consider the set of urls in $S_{\alpha_k}^{u_i} = \{\alpha_{k_1}, \alpha_{k_2}, \dots, \alpha_{k_n}\}$, where $\alpha_{k_i} \neq \alpha_{k_j}$, and their relative probabilities in $S_{\alpha_k}^{u_i}$. Let $p_{\alpha_{k_i}}$ be the probability of α_{k_i} and $\sum_{i=1}^n p_{\alpha_{k_i}} = 1$.

The Shannon entropy of a clickstream $S_{\alpha_k}^{u_i}$ is thus

$$H(S_{\alpha_k}^{u_i}) = - \sum_{i=1}^n p_{\alpha_{k_i}} \log_2 p_{\alpha_{k_i}}.$$

For a given user u_i , the clickstream entropy C over *all* click streams $S_{\alpha^*}^{u_i} = \{S_{\alpha_1}^{u_i} \cup \dots \cup S_{\alpha_k}^{u_i} \cup \dots \cup S_{\alpha_n}^{u_i}\}$ and $p_{\alpha_{k_i}^*}$ is the probability of the url α_{k_i} over *all* click streams. The lifetime entropy for all clickstreams for a user is defined given by

$$C_{u_i} = H(S_{\alpha^*}^{u_i}).$$

The result of this first step can now be obtained. With the lifetime frequency and variety – modeled as session entropy C_{u_i} – for a each user, use diffusion quadrants of the first step of the methodology can be completed. Each user is binned into a quadrant based on the frequency and variety and thus placed into one of the coarse-grained diffusion patterns previously described.

3.2 Step 2: User Typology Modeling

While use diffusion patterns provide domain-independent quadrants of generalized usage behavior, to understand fine-grained user behavior we apply data mining algorithms, specifically clustering. Having discussed the challenges of the variety variable in step 1 above, in the second step features were selected from the clickstream data that expand variety and frequency in order to discover more detailed views of user behavior. Since clickstream data provides a good metric for computing variety through entropy, by selecting

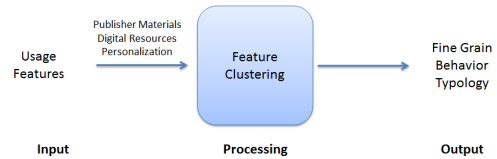


Figure 3: Step Two Overview

features that model variety in more detail, such as the usage of specific system components, we will develop a higher fidelity view of user behavior. These new refinements and the application of clustering expands the large grained use diffusion patterns into a fine-grained user typology that continues to model frequency and variety.

The first aspect of refinement requires that we expand frequency and variety by choosing more detailed features to model those variables. In the case of this study, frequency is expanded to include both sessions and total time spent within the system. Variety is expanded to include eight features that capture varying usage across system components that include publisher materials, digital library components and user contributed/social features of the system. While there are a number of clustering algorithms to choose for this second step, these experiments are based on the widely used model-based expectation maximization (EM) algorithm [7]. EM works by iteratively examining the parameters of each object instance to be clustered and builds a probability distribution that best explains where each object instance should belong. After many iterations, the model settles into a set of clusters that are represented by the model parameters derived for each cluster. The resulting output are clusters comprising a fine-grained user typology. The EM algorithm is used in these experiments because it is fast, robust and typically converges quickly. Furthermore, cluster shapes (e.g. circles, ellipsoids) may vary to include more flexible cluster membership.

4. CASE STUDY

For the remainder of this paper, we will focus on the case study that examines the use of the Curriculum Customization Service (CCS). The Curriculum Customization Service is a National Science Foundation funded program overseen by Digital Learning Sciences (DLS) – a joint institute of the University of Colorado at Boulder and the University Corporation for Atmospheric Research. DLS began development of the CCS in early 2008 and in July 2009, the CCS was made available to all Earth science teachers in a large urban school district in the midwest. Over 100 teachers were trained on the CCS for use in the 2009-2010 school year.

The CCS provides four major features to the end user (see Figure 4). First, it provides users with Web-based access to digital versions of the paper-based student textbooks, teacher manuals, and curriculum guides that comprise the Earth science curricula for both Grade 6 and Grade 9. The manuals and guides outline the state standards that must be met, explain how the various units in the Earth science curricula are connected to state standards, and provide additional supplementary materials for teacher use, such as ac-

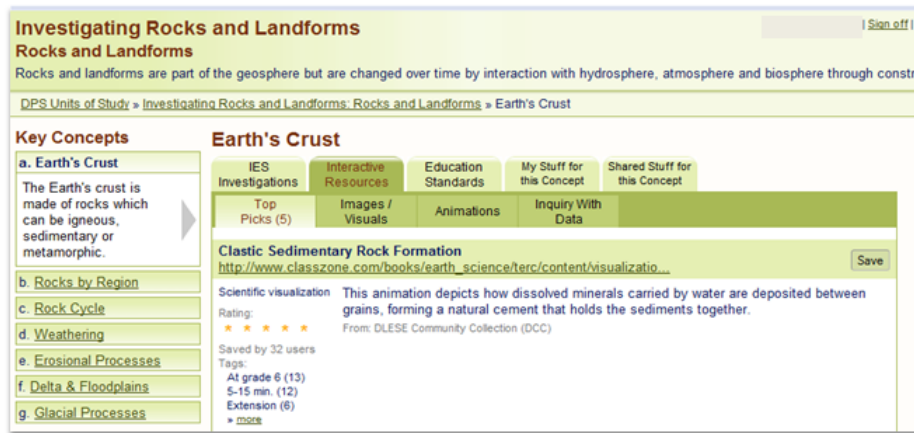


Figure 4: The CCS offers four major capabilities: access to publisher material (IES Investigations tab), access to digital library resources (Interactive Resources tab), personalization capabilities (My Stuff tab) and community features (Shared Stuff tab). The Interactive Resources component is opened above showing the top recommended digital resources.

tivities, teaching tips, and student assessments. These materials are all grouped under a single user interface component and are organized by key concept, which allows teachers to organize their lessons in a manner that flexibly meets the learning needs of their students. The digital versions of these curricular materials are identical to what was already available to teachers in paper form but can now be accessed from any computer with a Web connection.

Second, the CCS integrates digitized publisher content with interactive resources available from the Digital Library for Earth System Education (DLESE), a collection of Earth science related digital resources that are part of the National Science Digital Library. By clicking on the Interactive Resources tab, a user can see recommendations for animations, video clips, classroom activities, and other digital resources that pertain to the given key concept. The interactive resources available via the CCS have been vetted by the experts who manage the DLESE collection. Moreover, these resources are filtered by the system to ensure that they align with the Earth science curricula. Thus, when a teacher accesses a DLESE resource on, for example, volcanoes, the resource not only has been determined to have educational value by a subject matter expert, it has also been tied to the specific science concepts that must be taught as well as the science standards that must be met.

The third major feature of the CCS is an interactive Web 2.0 capability, whereby teachers can save digital resources recommended to them via the Interactive Resources component or they can upload their own resources to an area called My Stuff, thus storing teacher-developed materials in the same space as interactive resources from DLESE or from the curriculum for easy access. Once a resource is saved to My Stuff, teachers have the option to share a copy of the resource to an area of the CCS called Shared Stuff, which is accessible to any CCS user who clicks on the Shared Stuff component associated with a given key concept.

The final major feature enables teachers to share materials

with their peers. When a digital resource is added to Shared Stuff, the teacher who originally uploaded the resource, as well as other CCS users, can add searchable tags or keywords so that any search of the CCS system for those tags will list all resources tagged with the same key word or phrase. Finally, CCS users can add 'star ratings' to resources so that other users can determine how their colleagues rate a given resource. A resource that many users rate highly – four or five stars – might be more likely to capture the attention of system users than a resource with a low rating. Hew and Hara argue that this kind of sharing may be a catalyst for enabling improvement of practice [13], because such knowledge sharing tends to be tied to shared, situated instructional goals and challenges and is thus more likely to be relevant to a teacher's immediate, short-term needs. A more detailed description of the CCS, including results from the field trial, are described further by Sumner, et al [25].

4.1 Data Source and Step 1

The data source for this study was clickstream data of 98 users from 9 months of interaction within the CCS. The data for these experiments were particularly interesting because the user interface environment contains many dynamic client-side components which would not ordinarily be captured in a traditional web server log. The system was therefore instrumented to extract additional information from clicks on these dynamic client-side interface components. The CCS clickstream data log therefore includes detailed tracking of user activity that provides rich user interaction data. Since we were able to capture this data over a relatively long period of time, we were able to analyze actual user behavior as the users worked with the system in a natural, unhindered manner. Furthermore, we were able to examine how publisher materials, which are already core to standard teacher practice, are complemented by digital library resources, giving us a broader view of how teachers integrate these digital library resources into their traditional practices.

As a result of step 1, we obtained initial frequency and variety computations. Table 1 shows the descriptive statistics resulting from the first step of our experimental methodology. The data show that the mean entropy of our population is 5.5 and the mean frequency, here measured as the lifetime sessions that a user logged, is 37.85.

Entropy	Lifetime Sessions
Min.: 1.000	Min.: 1.00
1st Qu.: 4.787	1st Qu.: 7.00
Median: 5.749	Median: 23.00
Mean: 5.506	Mean: 37.85
3rd Qu.: 6.409	3rd Qu.: 48.00
Max.: 7.691	Max.: 171.00
Std. Dev.: 1.256	Std. Dev.: 42.201

Table 1: Summary statistics for the data plotted in Figure 5, $N = 98$.

When we apply the use diffusion model to the data by marking the quartiles at the means of each axis, we obtain a use diffusion pattern that shows that there is a large ($n = 38$) limited use group. Similarly, there are nearly as many ($n = 30$) intense users (entropy > 5.5 , frequency > 37.85) – those users that exhibit, both a larger amount of variety and frequency. The plot also shows that there is a potential outlier in the specialized category. This data point was included in our initial analysis because we want to focus on the complete data set; a larger data set would help to reveal whether this point is an outlier or a single instance of a larger pattern. As

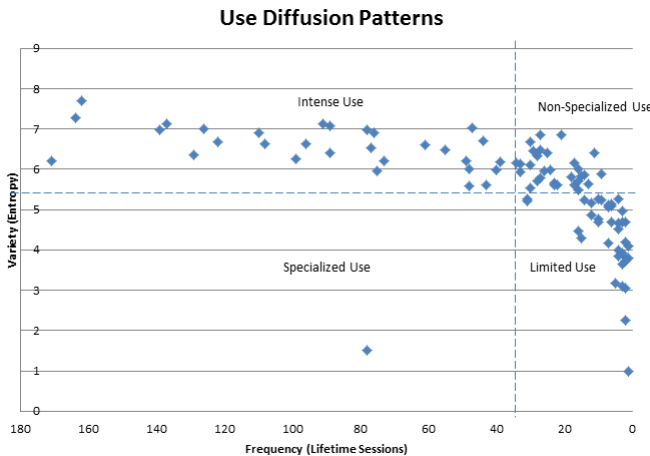


Figure 5: Use Diffusion pattern from step 1 showing frequency and variety of the data source.

can be seen in Figure 5, the value of the use diffusion pattern modeling is that it provides a comprehensive overview of coarse-grained behavior within the system. The dotted line divides the graph at the means of each axis, thus creating the four use diffusion quadrants. We can quickly see the distribution of basic usage patterns within the system. At the same time, however, the pattern does not produce enough information to determine the specific details of use. To do this, we must turn to the second step of the methodology.

4.2 Feature Selection and Step 2

We began the second step by choosing 10 features of our data for further analysis. The features for this step of the methodology represent four major functions of the system: (1) use of digital library-related system functions, (2) use of traditional publisher materials and related system functions, (3) system functions that involve personalization, and (4) user-contributed functions. Table 2 summarizes the system features used in the second step of our proposed methodology.

#	Feature Label	Description
1	Sessions	Total lifetime sessions
2	Hours	Total system hours
3	IR Activity	Total activity within interactive resources
4	IR Saving	Total interactive resource saving behavior
5	Shared Stuff Activity	Total user-contributed content activity
6	Shared Stuf Saving	Total user-contributed content saves
7	My Stuff Activity	Total ‘My Stuff’ activity
8	My Stuff Saving	Total ‘My Stuff’ saving behavior
9	Publisher-Teacher Materials	Total activity within publisher-teacher materials
10	Publisher-Student Materials	Total activity within publisher-student materials

Table 2: CCS cluster experiment features and descriptions.

4.2.1 Digital Library Features

The CCS is specifically designed with the goal of providing access to high-quality, digital library resources. Analyzing usage of the embedded digital library within the CCS should therefore provide useful data about typical teacher practices around digital resources. To capture this behavior, we selected features that detail the clickstream patterns of the embedded digital library resources within the system. These resources were presented in the user interface under four sub-categories: Top Picks, Animations, Images/Visuals and Inquiry with Data. Each category contained resources from DLESE that were either selected as highly relevant to the subject materials and unit focus (Top Picks) or contained metadata the were of the appropriate type, scope and topic (Animations, etc.) The items presented under each of these views were derived directly from DLESE web services [28] and appropriately presented to the user. Clickstreams into this component are tracked with the **IR Activity** and **IR Saving** features.

4.2.2 Publisher Materials Features

Publisher materials are included as a core component of the system functionality. The majority of the publisher’s items represent digitized versions of paper-based materials, whether they be book chapters, supplemental materials such as hand-outs, assessments, etc. The features here are therefore convenient digital proxies for real-world paper-based analogues. Clickstreams into this component of the system and corresponding sub-components were tracked and organized with the **Publisher-Student Materials** and **Publisher-Teacher Materials** features. The Publisher-Student Ma-

terials include publisher materials like digital versions of the student textbook, while Publisher-Teacher Materials include supplemental publisher materials such as instructional support materials.

4.2.3 Personalization Features

The CCS provides functionality to allow teachers to personalize the contents of their accounts, in particular, users are provided with the ability to save digital materials that they find of interest. Once saved, items may be retrieved for further review and may even be shared with others if desired. Saving is considered a personalization feature because it provides direct control over items that may fit a specific need (either at the time of save or in the future). Furthermore, saving implies an interest in the saved item, and while that interest may only last for a short time, it nonetheless acts as a marker of personalization behavior. Personalization behavior was captured with the saving behavior of the system through the **My Stuff Saving** feature. For example, teachers are able to save embedded digital library resources, such as animations, images, visuals and top picks, for the units of study they may be interested in. The **My Stuff Activity** feature represents the total activity performed within the My Stuff features of the system.

4.2.4 Community Behavior Features

There are features of the system that promote community-centric behaviors. For example, resources and other materials that users find interesting can be shared with the community at large, in a kind of community pool of resources called “Shared Stuff”. The feature has many implications when considering the nature of communities of practice of K-12 educators, who are often encouraged to share materials, pedagogical strategies and best practices amongst their peers. The community behaviors are captured with the **Shared Stuff Activity** and **Shared Stuff Saving** features.

4.3 A User Typology

Our second step relies on feature clustering to develop a user typology. There are many clustering algorithms to choose from, but for this set of typology experiments we chose the model-based EM algorithm [7]. Elsewhere we describe other experiments using other clustering algorithms and parameters [16]. Using the Bayesian Information Criteria to determine the optimal number of clusters for a given data set, it was determined there were 5 clusters to be discovered in our experiments. The details of the clusters are shown in Table 3. It should be noted that the clusters presented here show representative members of each cluster and do not necessarily represent actual users. That is, for each cluster in the table, the values present the ideal parameters that fit the object instances that belong within the cluster.

As can be seen in Table 3, cluster 1 characterizes the low-use pattern (low variety, low frequency) of the step 1 use diffusion pattern. The users in this cluster have produced very few hours of total use within the system. Furthermore they do not seem to be using the full range of system features. On the other hand, the experiments revealed an intense user cluster (cluster 3) that shows robust use of the system. Indeed, this “intense user” category seems to have used the system in full – exercising nearly every aspect of the system

Feature	Cluster				
	1	2	3	4	5
Sessions	4.934	47.096	119.551	26.131	83.897
Hours	1.0125	26.702	52.586	8.427	24.984
IR Activity	2.483	74.253	95.537	26.801	17.109
IR Saving	0.321	4.703	26.039	1.218	0.649
Shared Stuff Activity	3.382	25.753	81.159	17.558	82.174
Shared Stuff Saving	0.356	4.314	18.942	1.877	3.701
My Stuff Activity	0.123	4.864	19.116	1.106	0.860
My Stuff Saving	0.741	18.623	11.492	0.388	4.913
Publisher-Student Material	1.421	21.715	98.143	14.212	74.404
Publisher-Teacher Material	4.781	28.790	86.070	14.649	92.306
<i>N</i>	31	10	8	35	14

Table 3: Experimental cluster results showing the parameters obtained for each feature and cluster.

and logging both large numbers of sessions and significant hours of use. Two specialized user groups emerged from the typology. While they both show about the same number of total hours within the system, cluster 5, identifies users who spend a great deal of time with the community features and publisher materials of the system, while cluster 2 shows users who spend far more time within the interactive resources and embedded digital library component of the system. These clusters are valuable to understand, because it shows that some users find the embedded digital resources as important as other users find the community features. Finally, cluster four shows a group of users that had a session count that was above the median, but less than that of the intense and specialized users. This cluster also exhibits broad use of most of the features, with slightly more use of interactive resources. Table 4 summarizes each cluster, giving the key characteristics of the cluster and the diffusion pattern that the cluster belongs to. Furthermore, we have created fictitious typology labels to provide an easy way to remember the cluster characteristics.

4.4 Validation of Results with Field Research Findings

It is important to note that the methodology that produced the user typology clusters is not simply a new kind of analysis; rather, it presents an opportunity to bridge different analytical approaches. Although the clusters that emerged from our computational approach are based on clickstream data, these clusters are not arbitrary aggregations of user behaviors. They map onto real world usage patterns that emerged from the traditional educational research techniques used in our study of the CCS. Thus, the findings from our computational method are validated by the findings from our field research of CCS use ‘in the wild.’

Data were collected via traditional educational research methods, including surveys administered to all 124 Earth science teachers in the district, interviews of approximately 20 per cent of teachers, and classroom observations of approximately 7 per cent of teachers. An analysis of these

Cluster	Diffusion Pattern “Typology Label”	Characteristics
1	<i>Limited Use</i> “Uninterested Non-Adopter”	Very low over all system use.
2	<i>Specialized Use</i> “Interactive Resource Specialist”	Heavy use of interactive resources relative to other system features. Tends to access system weekly.
3	<i>Intense Use</i> “Ardent Power User”	Heavy and robust overall use of all features. Tends to access system daily.
4	<i>Non-Specialized Use</i> “Moderate Generalist”	Moderate overall system use. Shows slightly more use of interactive resources than other system features. Tends to access system several times monthly.
5	<i>Specialized Use</i> “Community Seeker Specialist”	Makes heavy use of Shared Stuff features and Publisher materials relative to Interactive Resource Activity. Tends to use the system weekly.

Table 4: An initial user typology derived from step 2, showing the diffusion patterns and characteristics of each usage type.

self-reported data, conducted independently of the computational analysis of clickstream data we have described above, suggested that teachers in our study fell onto a spectrum of system use behavior, from low-frequency users (corresponding to Cluster 1) to moderate (Cluster 4) and heavy (Cluster 3) users. The final two clusters fell on the moderate-to-heavy side of the spectrum.

Rogers’ theory of diffusion of innovation [20], discussed in our introductory section, predicts that the earliest users of a new technology (‘innovators’ and ‘early adopters’) comprise approximately 16 per cent of all users. Moore [17], who revised and extended Rogers’ work, further argues that a technology cannot become ‘mainstream’ within a given population until it is adopted by at least half of all potential users. Our Step 2 analysis indicates that about two-thirds of teachers in the district adopted the system to a significant degree with one-third of users (represented by clusters 2, 3 and 5) making heavy use of the system. Even if we confine ourselves to a theoretical framework, such as Rogers’, that focuses on quantifying when different segments of a target population adopt an innovation, our findings indicate that the CCS has been strongly adopted and is heavily used by the target population.

We discuss the motivations behind teachers’ use of the CCS in more detail elsewhere [16], but we will note here that teachers’ CCS use was, according to teachers’ self-reported data, heavily influenced by the perceived learning needs of their students. ‘Carlie,’ a teacher in the Ardent Power User cluster, stated in an interview: *Since [many of my] students are visual learners, or hands-on learners [who are] growing up in this technological age... [Engaging with interactive resources] seems to really drive home the idea or topic [being taught]. I ask them if [interactive resources] help them, if whatever they just viewed made the material make more sense, and oftentimes [they respond] “Wow, yeah, that to-*

tally made sense. Can we see it again?” Asked to speak more generally about her impressions of the CCS, Carlie stated: *[The CCS is] a centralized location where I can go to find that extra material that I know nine times out of ten [is] useful for me. [The CCS] actually has cut down on [my] random searching on the Internet.* (All teacher names are pseudonyms.)

It can be argued that identifying low, moderate, and intense user clusters is hardly a profound finding – after all, any set of system users can be divided into ‘low’ and ‘heavy’ user categories if the major variable of interest is frequency of system use. Recall that use diffusion theory calls for an analysis that incorporates frequency of use with an evaluation of how a given technology is used. Looking through the lens of use diffusion, the most salient finding from our computational methodology is that two of the clusters – Cluster 2, the Interactive Resource Specialist, and Cluster 5, the Community Seeker Specialist – represent ‘variety of use’ behaviors similar to user behaviors identified by our field research. This enables us to validate the findings of our computational methodology with real world data while at the same time giving us deep insights into how teachers integrated digital library resources into their instructional practices.

For example, Interactive Resource Specialists were moderately heavy users who spent most of their time in the Interactive Resources/Embedded Digital Library component of the CCS. In response to an interview question that asked him to discuss why he used the CCS, teacher ‘Corey’ told us: *With the CCS, I had a more centralized place to search for resources that I knew were aligned to the specific activities I was preparing and using in my classroom... My teaching practices have always focused on student engagement. I found that the CCS made it easier for me to find [interactive] resources with which I could capture student attention.*

In contrast, Community Seeker Specialists had CCS usage frequencies comparable to Interactive Resource Specialists but spent most of their time in the Shared Stuff and Publisher components of the system. In response to an open-ended survey question that asked respondents to explain the value of accessing Shared Stuff, teacher ‘Sheila’ stated: *Looking at the Shared Stuff uploaded by others gives me ideas about how I can present particular concepts in my [own] classroom.* In response to the same question, ‘Norma’ commented: *[Resources in Shared Stuff] have given me different perspectives on the different topics [in the curriculum] and thus enabled me to teach more effectively.*

5. DISCUSSION

As the case study shows, two useful views have emerged for understanding adoption. The first step of the methodology shows promise for quickly understanding application-independent large-grain usage patterns. The step second outputs a meaningful typology based on clusters that emerge from refined application-dependent feature data. The identification of clusters that correspond well with interview and self-reported data suggests there is promise and utility to the methodology we have introduced.

While we have obtained promising results from this study, there are two major weaknesses to consider. First, our

methodology has only been applied to a single population. It is plausible that the results that we obtained are specific to the population under examination. A different typology may emerge as a result of larger data populations, for example, new specialized user types might emerge. Second, in order to validate that our results are robust and generalizable, it will be important to study the methodology with different applications and user populations.

As can be seen in the case study, both of the steps of the methodology yielded interesting results. It appears, however, that the entropy modeling of the first step is lacking sufficient discriminatory power when describing specialized users. When compared with the clustering results of the second step, two clearly specialized user groups emerge, yet step one did not separate the specialized group enough to identify the emerging groups of the clustering experiments. More work needs to be done to develop an entropy model that separates the specialized users with better precision. It could furthermore indicate that specialization does not always imply low variety – and this is clear when considering the details of the system usage when variety is expanded to include more variables. It will be valuable to revisit our entropy model to develop a slightly more sensitive metric for variety, in particular for specialized usage.

6. BROADER IMPACTS

While this research has focused on demonstrating the value of our methodology for the study of digital library adoption, we believe it can help address other educational challenges, chief among which are teacher professional development, the correlation of teacher practice to student learning, and the evaluation of teaching practices.

Extant research indicates that one of the major barriers to effective integration of technology into teaching practices is the dearth of quality professional development and training vis-à-vis technology [12]. Teachers often complain that even when technology-related training is made available, they are often confronted with training that does not meet their needs as practicing educators because it assumes that they do or do not already possess a given set of skills. By modeling teachers' system use behavior, one will be able to understand inter-user differences and target system training and professional development to users' true needs. For example, CCS users in the Community Seeker Specialist cluster might be presented with training that would help them integrate more interactive resources into their teaching because such users are known to not spend much of their CCS usage time exploring interactive resources. For teachers in the Interactive Resource Specialist cluster, such training would be redundant because they are already making extensive use of the interactive resources available via the CCS.

While a consensus has emerged among policy makers and education researchers that effective adoption of technology can indeed have a positive impact on teacher practice [23, 15], the relationship between teachers' adoption of technological systems and student achievement is still not well understood. As part of our larger study of the CCS, we are analyzing the standardized test scores of students taught by teachers who used the CCS to determine what impact the teachers' use of the CCS had on student outcomes. In the

long term, the ability to correlate system use behaviors with students' academic achievement will make it possible to focus on teacher use behaviors that most benefit students; in turn, these behaviors could be taught to other teachers.

Evaluating teachers' instructional practices is a very difficult task that requires a massive commitment of human and financial resources [4]. For example, most evaluation systems rely on administrators to observe teachers in the classroom, a very labor-intensive approach that can result in teachers being awarded or denied tenure or pay raises based on the short period of time during which they were observed. Further, variance between evaluators' adherence to evaluation rubrics and bias in the evaluation instruments themselves can make it difficult to assess the validity of the evaluation process [29]. Alternatively, asking teachers to self-report their teaching practices might provide evaluators with additional information but such data would be subject to all the limitations that come with asking individuals to report on their own behaviors. Neutral 'third party' data – such as the clickstream data to which we applied our computational methodology – can help bridge the gap between what evaluators observe during the relatively brief periods they visit a teacher's classroom and the teacher's activities when he or she was not being formally evaluated. A teacher evaluation process that incorporates traditional observational and self-reported data with usage data produced by systems like the CCS would give administrators, policymakers, and teachers themselves deeper insights into instructional practices; further, such a hybrid system would produce a rich new understanding of teaching-related best practices that then could be shared with other educators.

7. CONCLUSION

By applying models of adoption and use diffusion alongside data mining techniques, we have developed a two-step methodology for discovering patterns within clickstream data that reveal both general and specific typologies of digital library user behavior. The application of this methodology to data from an academic year-long field trial in a large urban school district provided a rich opportunity to study the adoption and usage behavior of embedded digital library resources by middle and high school science teachers. The methodology shows considerable promise for extracting useful behavioral patterns “in the wild;” the resulting fine-grained user typology maps well with results emerging from traditional educational field research methods. The proposed methodology, while requiring more data to further validate, provides a valuable contribution towards our effort to develop automatic methods for studying digital library adoption and use. When fully realized, this methodology also has the potential to be extended to other applications and areas, such as informing teacher professional development, understanding the impact of digital library applications on student learning, and developing new approaches to the evaluation of teacher practice.

8. ACKNOWLEDGEMENTS

This material is based upon work supported in part by the NSDL program in the National Science Foundation under Awards #0734875 and #0840744 and by the University of Colorado at Boulder.

9. REFERENCES

- [1] E. Ayers, R. Nugent, and N. Dean. Skill Set Profile Clustering Based on Student Capability Vectors Computed From Online Tutoring Data. In *Proceedings of the 1st International Conference on Educational Data Mining*, pages 210–217, Montreal, Canada, June 2008.
- [2] F. Benevenuto, T. Rodrigues, M. Cha, and V. Almeida. Characterizing user behavior in online social networks. In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference*, pages 49–62, Chicago, Illinois, USA, 2009. ACM.
- [3] P. B. Brandtzæg. Towards a unified Media-User typology (MUT): a meta-analysis and review of the research literature on media-user typologies. *Computers in Human Behavior*, 26(5):940–956, September 2010.
- [4] C. Danielson and T. L. McGreal. *Teacher Evaluation To Enhance Professional Practice*. Association for Supervision and Curriculum Development, Alexandria, VA, USA, 2000.
- [5] F. D. Davis. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS quarterly*, 13(3):319–340, 1989.
- [6] G. Deffuant, S. Huet, and F. Amblard. An Individual-Based model of innovation diffusion mixing social value and individual benefit 1. *American Journal of Sociology*, 110(4):1041–1069, 2005.
- [7] A. P. Dempster, N. M. Laird, D. B. Rubin, et al. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.
- [8] A. K. Dominguez, K. Yacef, and J. R. Curran. Data mining for generating hints in a python tutor. In *Proceedings of the 3rd International Conference on Educational Data Mining*, pages 91–100, Pittsburgh, PA, June 2010.
- [9] R. Eynon and L. Malmberg. A typology of young people’s internet use: Implications for education. *Computers & Education*, 56(3):585–595, April 2011.
- [10] F. F. Fuller. Concerns of teachers: A developmental conceptualization. *American Educational Research Journal*, 6(2):207–226, 1969.
- [11] G. E. Hall. The Concerns-Based approach to facilitating change. *Educational Horizons*, 57(4):202–208, 1979.
- [12] K. Hanson and B. Carlson. *Effective Access: Teachers’ use of digital resources in STEM teaching*. Gender, Diversity, and Technology Institute, Education Development Center, Inc., 2005.
- [13] K. F. Hew and N. Hara. Empirical study of motivators and barriers of teacher online knowledge sharing. *Educational Technology Research and Development*, 55(6):573–595, 2007.
- [14] J. Horrigan and Pew Internet & American Life Project. *A typology of information and communication technology users*. Pew Internet & American Life Project, 2007.
- [15] M. G. Kelly and A. McAnear. *National Educational Technology Standards for Teachers: Preparing Teachers To Use Technology*. International Society for Technology in Education (ISTE), Eugene, OR, 2002.
- [16] K. E. Maull, M. G. Saldivar, and T. Sumner. Online curriculum planning behavior of teachers. In *Proceedings of the 3rd International Conference on Educational Data Mining*, pages 121–130, Pittsburgh, PA, June 2010.
- [17] G. A. Moore. *Crossing the chasm: Marketing and selling technology products to mainstream customers*. HarperCollins: First Collins Business Essentials, New York, NY, USA, 2006.
- [18] M. C. Pennington. Cycles of innovation in the adoption of information technology: a view for language teaching. *Computer Assisted Language Learning*, 17(1):7–33, 2004.
- [19] S. Ram and H. Jung. The conceptualization and measurement of product usage. *Journal of the Academy of Marketing Science*, 18(1):67–76, 1990.
- [20] E. M. Rogers. *Diffusion of Innovations, 5th Edition*. The Free Press, New York, NY, USA, 2003.
- [21] C. E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423, 1948.
- [22] C. Shih and A. Venkatesh. Beyond adoption: Development and application of a Use-Diffusion model. *The Journal of Marketing*, 68(1):59–72, January 2004.
- [23] B. Smerdon. *Teachers’ tools for the 21st century: A report on teachers’ use of technology*. US Dept. of Education, Office of Educational Research and Improvement, 2000.
- [24] E. T. Straub. Understanding technology adoption: Theory and future directions for informal learning. *Review of Educational Research*, 79(2):625–649, 2009.
- [25] T. Sumner and CCS Team. Customizing science instruction with educational digital libraries. In *Proceedings of the 10th annual joint conference on Digital libraries*, JCDL ’10, pages 353–356, New York, NY, USA, 2010. ACM.
- [26] M. Turner, B. Kitchenham, P. Brereton, S. Charters, and D. Budgen. Does the technology acceptance model predict actual use? A systematic literature review. *Information and Software Technology*, 52(5):463–479, 2010.
- [27] M. P. Venkatraman. The impact of innovativeness and innovation type on adoption. *Journal of Retailing*, 67(1):51–67, 1991.
- [28] J. Weatherley. A web service framework for embedding discovery services in distributed library interfaces. In *Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, JCDL ’05, pages 42–43, New York, NY, USA, 2005. ACM.
- [29] B. Wilson and J. A. Wood. Teacher evaluation: A national dilemma. *Journal of Personnel Evaluation in Education*, 10(1):75–82, 1996.
- [30] B. Xu, M. Recker, and S. Hsi. Data deluge: Opportunities for research in educational digital libraries. In *Internet Issues: Blogging, the Digital Divide and Digital Libraries*. Nova Science Pub Inc, September 2010.