Proceedings of the Twentieth Innovative Applications of Artificial Intelligence Conference (2008)

The Law of Choice and the Decision Not to Decide

David C. Wilson† and **Suzanne Leland**‡ and **Kenneth Godwin**‡ and
Andrew Baxter‡ and **Ashley Levy**‡ and **Jamie Smart**‡ and
Nadia Najjar† and **Jayakrishnan Andaparambil**†
†Department of Software & Information Systems
‡Department of Political Science
University of North Carolina at Charlotte
9201 University City Blvd., Charlotte NC, 28223, USA
{davils,smleland,godwink,abaxter1,aclevy,jdsmart,nanajjar,jhandapa}@uncc.edu

Abstract

Public school choice at the primary and secondary levels is one of the key elements of the U.S. No Child Left Behind Act of 2001 (NCLB). If a school does not meet assessment goals for two consecutive years, by law the district must offer students the opportunity to transfer to a school that is meeting its goals. Making a choice with such potential impact on a child's future is clearly monumental, yet astonishingly few parents take advantage of the opportunity. Our research has shown that a significant part of the problem arises from issues in information access and information overload, particularly for low socioeconomic status families. We have developed an online application, called *SmartChoice*, to provide parents with school recommendations for individual students based on parents' preferences and students' needs, interests, abilities, and talents. The application employs content-based recommender systems techniques to generate a personalized list of schools that have a demonstrated track record with children who share a child's interests, skills, and needs. The first version of the application is deployed and live for focus group participants who have been using it for the January and March/April 2008 Charlotte-Mecklenburg school choice periods.

Introduction

Families whose children attend schools that are not making adequate yearly progress (AYP) under the guidelines of the 2001 No Child Left Behind (NCLB) Act are granted the legal right to make an important decision—whether to send their child to a different school. When this occurs, school districts must give students the option of moving to a school that is meeting its AYP goals. However, across the country, less than 6% of eligible students take advantage of this provision of the law (GAO 2004; Brown 2004; Howell 2006). Part of the theory of action behind NCLB, and public school choice more generally, is that giving parents information about school performance and allowing students to leave low-performing schools will improve student outcomes and force poorly performing schools to improve. However, this provision of NCLB has largely failed to be effective because parents must (1) be able to identify which schools will improve their child's performance and (2) actively make choices about their child's education. Research has shown that poor tactical implementations of NCLB policy can effectively strip parents, particularly low-income parents, of their legal options and rights, discouraging and limiting practical participation, and ultimately creating a culture of disincentive to exercise choice rights under current policy.

Our applied research program examines whether a comprehensive community program, collectively referred to as the *SmartChoice Program*, centered on a computer-assisted decision support system, *SmartChoice*, can help to overcome this implementation problem in current education policy. The decision system provides personalized recommendations for individual students based on parents' preferences and students' needs, interests, abilities, and talents. In general, the *SmartChoice Program* studies the extent to which (1) incentives for participation and (2) personal assistance in the decision process can increase choice behavior among low-income and minority families who occupy the majority of NCLB designated schools in our region. In the Applications of Artificial Intelligence context, *SmartChoice* employs core user modeling, personalization, and content-based recommender systems techniques to accomplish its goals.

We deployed the first version of *SmartChoice* in January 2008, and it is currently live and available for use by over 50 real end-user participants. Initially, the *SmartChoice* user population has been limited to participants selected for our pilot study. Participants have been using the system as support for the January 2008 Charlotte-Mecklenburg magnet school choice period and the subsequent March/April 2008 general school choice period. We expect *SmartChoice* access for participants to remain available beyond the study period and for the foreseeable future. While this is a deployed application, it has been in use for only a short period, and choice evaluation results of the pilot study will not be available until ongoing follow-up studies are completed with current users over the next six months to a year. Thus we view the current *SmartChoice* deployment as a closed beta that serves as a focus for discussing the application of recommender systems and related techniques to the domain problem of school choice.

In order to motivate the application, this paper begins with an introduction to the school choice domain. It goes on to detail the need for recommender support in school choice, and it identifies some common erroneous user choice behaviors as potential design considerations. The paper provides a brief overview of recommender systems in general, and it describes the modeling and recommendation approaches employed in *SmartChoice*. The paper concludes with a description of the ongoing initial deployment, discussion, and future directions.

The Law of Choice

The No Child Left Behind Act (NCLB 2001) is just one example of the type of instrument that is likely, in various incarnations, to mark education policy in the coming years—one that combines additional resources and new rights for low-income students. Supporters and opponents of No Child Left Behind agree that the legislation marked a revolutionary change in the federal government's role in education (Meier & Wood 2004; Kim & Sunderman 2004; McDermott & Jensen 2005). In one stroke NCLB changed the requirements for continued funding of Title I schools (those schools with high percentages of students at or near the federal poverty line). After the passage of NCLB, to continue the funding stream, districts must use outcome-based evaluations and show that individual schools have met Adequate Yearly Progress (AYP) goals. If a school does not meet AYP goals for all ethnic and income groups for two consecutive years, the district must offer students in the school the opportunity to transfer to a school that is meeting its goals.

Research on the implementation of NCLB has discovered several obstacles to effective reform. School districts have been slow to alert parents of their right to choose and typically give parents less than three weeks to investigate their options, make a choice, and apply for the school transfer. In addition, districts often fail to give parents good options to their current school (Howell 2006).

The most important barriers to choice, however, are procedural complexity, parents' lack of information and their unwillingness to choose. Despite being notified by their school district, many low-income parents are unaware of their right to choose and of the fact that their school failed to meet its AYP goals for two consecutive years (Howell 2006). Those parents who are aware of the options NCLB guarantees, often are confused about the choice process and believe they lack the capacity to choose wisely (Howell 2006; Leland, Godwin, & Baxter 2007). This often leads to a conscious decision not to exercise school choice options.

The Decision Not to Decide

In order for choice provisions targeted at low-income parents to raise the achievement of their children, at least two conditions must hold: (1) low-income parents must participate in the choice process; and (2) they must be able to evaluate their child's likely success at a prospective school.

Academic research on NCLB's choice provisions reveals that few NCLB parents have exercised the right to choose an alternate school. Kim and Sunderman (2004) studied ten urban school districts and discovered that fewer than 3% of eligible students requested a transfer to a different school. These findings were corroborated by a national study conducted by the GAO, which estimated that less than 1% of eligible parents had participated in the choice provision (GAO 2004). Howell (2006) found that only 3.8% of NCLB-eligible students in Massachusetts applied for a transfer. Brown (2004) supplies the highest nationwide estimate at 5.6 percent of eligible students.

Howell (2006) also revealed that although parents with children enrolled in NCLB schools claimed that they were familiar with its provisions, the majority were unaware that their child was enrolled in an NCLB-sanctioned school. Howell concluded that if advocates of NCLB are to boost participation, the options provided to parents must be in a format that encourages them to participate in the choice process (p. 174).

Our own research underscores Howell's finding. In 2006, our project team held focus groups with parents whose children attend NCLB-sanctioned schools to investigate why they did not choose an alternative to their assigned school. We found that the principal reasons were: (1) they were unaware of their choices; (2) they feared that they would make an incorrect choice and reduce their childs educational opportunities; (3) they found the information difficult to use and it made them feel confused and ill-informed about their available options; and (4) they did not trust their school system (Leland, Godwin, & Baxter 2007).

Need for Recommendation

A key aspect of the argument for the early advocates of public school choice was that it would force all schools to improve because a high percentage of parents would be informed about school quality. They would become well-informed because having the right to select their children's schools would give parents the incentive to become informed (Coons & Sugarman 1978; Chubb & Moe 1990). The difficulty with this argument is that parents have very little information about schools, even in districts with expansive public school choice (Ambler 1994; Schneider *et al.* 1997; Godwin & Kemerer 2002; Dunk & Dickman 2003).

The information low-income families have about schools tends to come from two easily seen cues—(1) the composition of its student body and (2) the school's average test score. Some parents believe that schools with fewer minorities or with more high-income families are better schools for their child (Leland, Godwin, & Baxter 2007). While these relationships may hold in the aggregate, for individual students they often are misleading. For example, research by Southworth and Mickelson (2007) shows that schools which are predominantly Anglo and high-income often place African American students in tracks with low-performing students and that this reduces their academic achievement. Thus a minority parent who chooses a predominantly white school may harm rather than help their child.

The other, largely unhelpful, cue that parents use to evaluate schools is the average end-of-grade test scores. Our focus groups found that almost all parents believe that peer effects are important in a child's academic growth and that their child will learn more the higher the average test score at a school. But some research suggests that this view is often incorrect.

Hoxby and Weingarth (2005) demonstrate that choosing a school based on the aggregated test score can be unwise because peer effects are neither uniform nor linear. On average, having higher achieving peers is better than having lower achieving peers. But Hoxby and Weingarth discovered that students whose achievement scores are substantially above or below the modal score of their peers do not learn as much as students whose scores are closer to the mode. Thus, selecting a school on the basis of higher test scores could harm rather help a student.

In another study of the impact of NCLB choice on student outcomes, Hastings and Weinstein (2007) find that for the students who enrolled in an alternate school rather than stay in a NCLB-sanctioned school, their achievement scores increased if the new school's average test score was greater than one school-level standard deviation higher than their NCLB school. In line with the Hoxby and Weingarth findings, however, the research found that the effect of moving to a school with higher test scores was greatest for students whose prior test scores were substantially above average.

Paul Hill writes that parents not only need to be informed about choice and the options available to them, they "need rich information about what different schools offer and how effective those schools are with children who have different interests and learning styles" (2007). To match a student with a school, the parent needs to know how well the available schools have educated students with similar interests, backgrounds, and skills to their child. But unless the parent is an econometrician with the appropriate dataset, they cannot obtain the information they need (McCaffrey *et al.* 2003).

Overall, previous research concerning school choice by low-income parents indicates that what low-income parents need and want is an individualized, easy-to-understand recommendation that comes to them in a non-numeric way (Leland, Godwin, & Baxter 2007; Hill 2007)

Recommender Systems

So called recommender systems have gained considerable interest since the 1990's as a means of helping users to deal with ever-increasing problems of information overload (Resnick & Varian 1997). Algorithmic developments to date have given rise to a variety of different basic recommendation techniques and strategies (Burke 2002). *Content-based* techniques, for example, rely on the availability of descriptive meta-data that captures the essence of the items available for recommendation. So, a movie recommender might make use of descriptions including genre, actor, and summary plot information as the foundation for similarity assessment techniques to match a target user's profile interests (Rosenstein & Lochbaum 2000; Basu, Hirsh, & Cohen 1998; Soboroff & Nicholas 1999; Smyth & Cotter 2001).

Collaborative Filtering techniques provide an alternative strategy that relies on ratings-based user profiles instead of descriptive meta-data (Schafer *et al.* 2007; Konstan *et al.* 1997; Smyth & Cotter 2001; Terveen *et al.* 1997). Suitable items for recommendation are identified not because their description matches them with a target user, but rather because these items have been liked by users who are "similar" to the target user in terms of how they have rated other items. So, a collaborative filtering movie recommender "knows" nothing about a movie's genre or actors, but it knows that other users have liked this movie and that these users have demonstrated a similar taste to the target user, having liked and disliked many of the same movies in the past.

Demographic techniques make recommendations based on demographic classes, by analyzing personal attributes of the user (Krulwich 1997). Likely items for recommendation are identified because clusters of users with similar personal attributes have demonstrated similar needs or tastes. So, a demographic movie recommender could, for example, focus on the language spoken or geographic region as a basis for recommendation. Thus in general, content-based methods rely on item-item (Sarwar *et al.* 2001) and item-user (Sarwar *et al.* 2000) similarities whereas collaborative filtering and demographic methods rely on user-user similarities (Konstan *et al.* 1997; Krulwich 1997).

A wide variety of recommender systems have been deployed online for commercial, educational, and community-building purposes (Montaner, López, & de la Rosa 2003; Burke 2002; Schafer, Konstan, & Riedl 1999; Herlocker, Konstan, & Riedl 2000). In our work, we employ primarily content-based techniques to match users and students with schools that are expected to provide improvements in academic achievement. Content-based recommender systems often employ classification learning algorithms as the foundation for recommendation. Typical models underlying content-based recommendation include Decision Trees, Nearest Neighbor approaches, Relevance Feedback, Probabilistic methods, and Linear Classifiers (Pazzani & Billsus 2007). We adopt a Linear Classifier approach for modeling student achievement that is based on existing educational and behavioral science models (Tekwe *et al.* 2004). The stakes are fairly high in real-world adoption of school choice recommendations, so we were particularly interested in beginning with an approach to achievement modeling that has a solid grounding in the domain. The two main components in content-based recommendation are (1) to measure the relevance of candidate items to the target user, and (2) to select the most relevant items for presentation as a recommendation set. The following sections provide an overview of these aspects in the *SmartChoice* recommendation approach.

Student Achievement Model

Using achievement, demographic, and enrollment data for approximately 74,000 students in Charlotte-Mecklenburg Schools (CMS) in 3rd through 8th grades from 2004-2006, we estimated the effect of attending a particular school on a particular student's growth in his or her composite math and reading standardized end of grade exams. In order to model the expected performance improvement relationship between a given student and potential schools, we employed an ordinary least squares regression analysis with school fixed effects (Tekwe *et al.* 2004). Some of the characteristics we accounted for include the student's ethnicity, sex, socioeconomic status, prior test scores, learning disabilities, and parents' education level, as well as the estimated effect of repeating a grade or being in a new school.

Unlike some models that estimate a school's "value-added" effect on student achievement, our model intentionally did not attempt to control for school-level variables (e.g., percentage of students who are free lunch eligible) in estimating the school's effect on achievement, for what is important to parents is what their child will actually experience at the school given precisely the variables that models used for accountability reasons often attempt to isolate from the estimate of the school's effectiveness (e.g., (Ballou, Sanders, & Wright 2004; McCaffrey *et al.* 2004; Raudenbush 2004; Rubin, Stuart, & Zanutto 2004)).

To address possible year-to-year fluctuations in these school-level factors that would influence our estimates of a school's effect, we averaged our estimates over the last three years. Furthermore, given the segregation of Charlotte-Mecklenburg area neighborhoods and a student assignment policy based primarily on a student's residence, we were concerned that we might extrapolate effects to schools that had enrolled few if any students with a specific combination of characteristics, (e.g., black males not eligible for free lunch). To address our concern, we estimated separate models for 18 such combinations so that the confidence intervals for school effects in such cases would be wide enough to reflect the degree of uncertainty for a particular type of child.

We trained and then tested our models on distinct sub samples, selecting the models that maximized the accuracy of our predictions. We found that, on average, students who changed schools to one that our algorithm predicted would improve their scores experienced statistically and substantively greater improvements in test scores than both students who changed to schools that our algorithm predicted would not be appropriate for them and students who did not change schools.

Recommendation Strategy

As one can imagine, making good recommendations in this domain is of paramount importance. Thus in making recommendations, we adopt a conservative analysis of the expected performance improvement results for a particular student to obtain a set of "recommendable" schools. Our analysis filters examine the point estimates and confidence intervals from the achievement models. Here we discuss two of the filters employed, along with our baseline clustering approach. The fundamental criterion for our recommendable schools is that they should not allow a reduction in expected academic performance improvement.

Recommendable Schools

As a baseline, we conduct a meta-analysis of the results, examining the model rather than the domain. Results that have a poor statistical grounding with our dataset are discarded out of hand. For example we do not consider schools that the model does not have enough data to make a prediction (naturally) or schools that have significant zero crossings in the confidence interval.

At the domain level, because a student's baseline option is to remain at their currently scheduled school for the next year, we analyze the recommendations relative to a student's expected performance improvement at their scheduled school. Schools that have an confidence interval falling completely below the scheduled school are not included. In addition, schools with only modest relative overlap on the scheduled school's lower confidence interval are considered to have a low probability of matching the baseline, and they are also not included. If, for some reason, a student does not have a currently scheduled school, the analysis is made relative to the grand mean.

Clustering Recommendations

In order to make the best possible recommendations, we would like to determine whether one school shows an unequivocal performance improvement relative to others. Taking confidence intervals into account, this amounts to clustering across a partially-ordered set. We have adopted a clustering approach, which groups recommendable schools into clusters, where (1) each cluster has a definite rank in relation to the other clusters, and (2) within each cluster, a candidate school has a definite rank relative to the other candidates.

In clustering, we select the recommendable candidate with the least upper bound on its confidence interval as the baseline. Candidates with significant confidence interval overlap to the baseline are grouped into the same cluster. Then the next unassigned candidate with the least upper bound on its confidence interval is selected, and the process is repeated until no candidates remain. In essence, all entries in a "better" cluster must have a least lower bound that is greater than the worst least upper bound in the "worse" cluster(s).

In this way, we know that all entries in a "better" cluster are definitely and significantly better than at least some of the entries in a "worse" cluster. And while some of the entries in the "worse" cluster may have confidence intervals that potentially put them above some of the entries in a "better" cluster, they could just as easily fall in range of the definitely worse category. Such entries are considered tainted by the latter possibility and relegated to the lower-ranked cluster.

Candidate Ranking

Within each cluster, a relative ranking of the candidates is determined. We use the size of a confidence interval as a proxy for the degree of uncertainty in a given point estimate. Point estimates are first normalized to the overall cluster interval. We then normalize the point estimates to reflect the uncertainty of their associated intervals in relation both to (1) the least uncertainty estimate within the cluster and (2) the overall uncertainty measure of the cluster itself. This gives us a relative measure of the uncertainty of a particular estimate in comparison to other estimates within the cluster. The resulting score is used to rank candidates within the cluster.

Making Recommendations

Given the size of a recommendation list k, the default recommendation mode is to select the best ranked candidates from the best ranked clusters until the recommendation set is full. We have developed functionality to allow more flexible navigation of the recommendation space, focusing for example on special programs, but it was decided not to deploy this functionality for the pilot study. We note that it is entirely possible to receive an empty recommendation list, either in the case that a student's current school is already the best of the best, or the underlying data is inconclusive. In these situations, we provide appropriate notification to the user. Often, recommender systems would employ a heuristic to avoid users feeling like their interaction has not been worthwhile (e.g., random or serendipitous selection). In the *SmartChoice* domain, however, such heuristics are not appropriate. Instead, we provide users with a reference list, where they can select from any available school. Thus users can still access supporting school information, but since it is their own selection, it does not carry the authority of a recommendation.

Initial Deployment

SmartChoice is deployed in a commercial web-hosting environment using a typical Linux, Apache, MySQL, and PHP (LAMP) baseline platform, and the application employs a custom model-view-controller architecture. Since we deal with potentially sensitive user data, particular development attention has been paid to addressing web application security issues, and all user-data interactions take place over secure HTTP connections.

The research site for *SmartChoice* is the Charlotte-Mecklenburg School District, an urban, predominantly minority school district in Charlotte, North Carolina. In 2007, children in approximately 8,200 families in CMS (98% are people of color, 92% low-income) attended low-performing schools that qualified for NCLBs choice provision.

We fielded focus-group usability tests in Fall of 2007, which resulted in significant improvements to interface and interaction characteristics. Our team includes members with significant domain expertise both in the general theory and practice of school choice, as well as with CMS in particular. Both the student achievement model and its application in the *SmartChoice* recommendation approach were tested extensively and validated by our domain experts prior to deployment.

Since the beginning of January 2008, we have been making *SmartChoice* available to a closed pilot-study group of over 50 parents with children in 3rd-7th grade. Our pilot-study is taking place in the West Boulevard Corridor of Charlotte, where about half of the NCLB schools in the district are located, during the 2008 school choice application periods for Charlotte-Mecklenburg schools. There have been two primary choice periods. In January 2008, parents made choices for the magnet school assignment lottery. Results of the magnet school assignment selections became available in mid-March and parents can currently apply for general school reassignment until early April 2008.

During the January choice period, we hosted weekly sessions where participants (1) registered for the program and site access, (2) went through a full initial *SmartChoice* recommendation session (with significant, moderate, or no counseling support, depending on study group assignment), and (3) provided feedback on the experience in an exit interview for the initial session. Sessions were hosted at the West Boulevard Public Library, and we provided dedicated computing / printing setups for participants. Once registered, participants have standard web access to *SmartChoice*. Here we present a short overview of *SmartChoice* application interaction.

Example Interaction

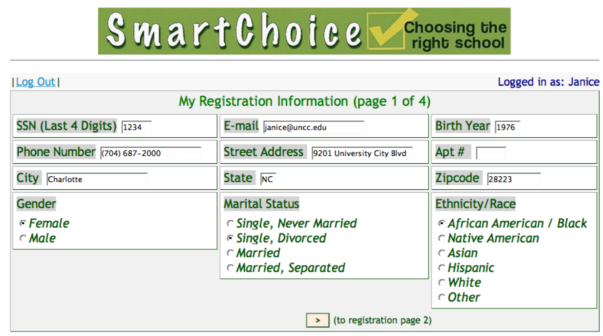

Figure 1: *SmartChoice* user registration.

Take as a hypothetical parent Janice who has a son Allen. Janice would first register with the system, as shown in Figure 1. She might then indicate that she is looking for a school for her fifth-grade, Latino son who enjoys math but currently requires programs to assist him in acquiring English. Example student information screens are shown in Figures 2 and 3.

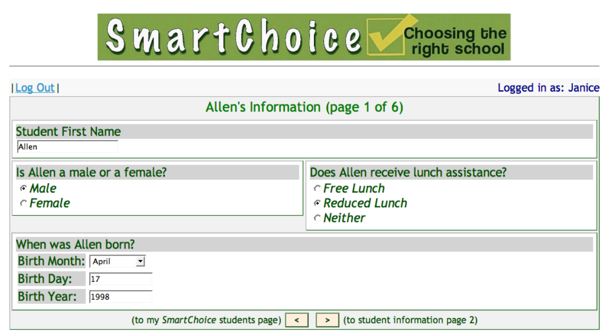

Figure 2: *SmartChoice* student information I.

In the recommendation list, *SmartChoice* provides basic preference feedback, showing estimated distance to the users address, along with visual cues for whether a school: (1) is a Title I school (those schools with high percentages of students at or near the federal poverty line), (2) has a magnet program, (3) will provide transportation to Allen, or (4)

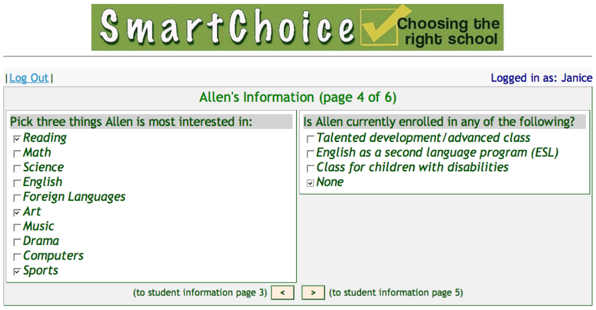

Figure 3: *SmartChoice* student information II.

has a specific selected type of magnet program. An example school results list is shown in Figure 4. Janice can select

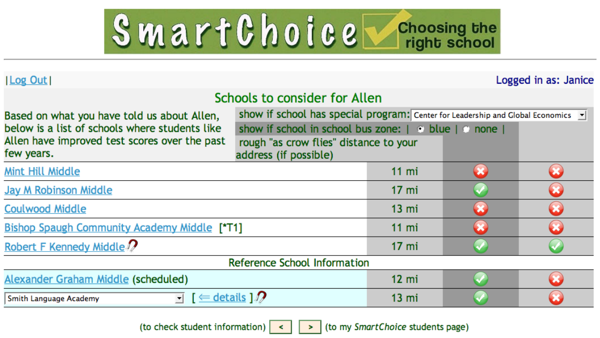

Figure 4: *SmartChoice* school recommendations.

a particular school from the list to see additional details, as shown in Figure 5.

Preliminary Results

The January sessions included a total of 54 participants. We have compiled initial qualitative results from the exit interviews for how satisfied participants were with *SmartChoice* support from their initial session. As shown in Table 1, a majority of participants were very satisfied with the support provided.

Responses	Satisfaction
30	Very Satisfied
20	Satisfied
2	Dissatisfied
0	Very Dissatisfied
2	*no answer*

Table 1: Satisfaction with *SmartChoice* support.

The most telling results will arise only after following up with the participants over the next six months to a year, in order to determine the impact on making choices and potential student outcomes. Initial followup calls show promising results. With 20 responses to date, 60% elected to choose a new school, and 62% of those received their first choice. While the data is only suggestive at the moment, in light of historical choice rates at fewer than 7% exiting NCLB home schools, it is also fairly encouraging.

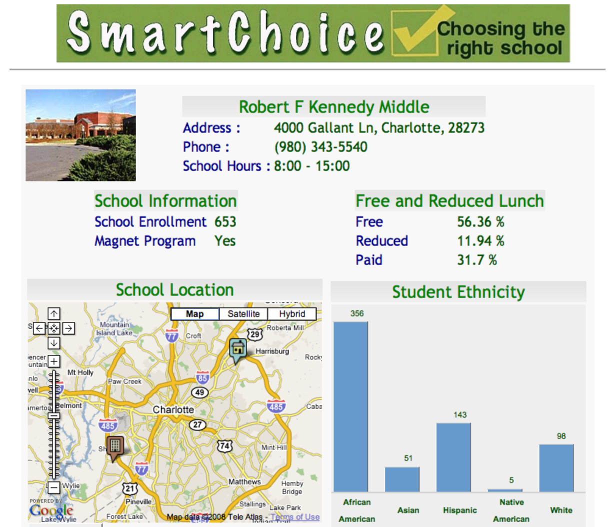

Figure 5: *SmartChoice* school detail.

Discussion and Future Work

In general, while we are employing recommender systems techniques, we believe that it is important not to frame system recommendations as such. Rather, we present them as "schools to consider." We recognize that the domain space is enormously complex (e.g., with peer effects, teacher / school / curriculum churn, and many other factors), and that *SmartChoice* focuses on a limited, but important aspect of that space. We are up to the challenge, however, and we are actively working to secure funding to build *SmartChoice* out to a full generally available deployment. Beyond making *SmartChoice* generally available, there are myriad refinements to be made. We are considering latent growth curve modeling of student achievement (Singer & Willett 2003; Bollen & Curran 2006) to increase the precision of modeling estimates. We know that it is important to present recommendations in a way that conveys meaning balanced with complexity (Herlocker, Konstan, & Riedl 2000). There is a significant proportion of parents in our NCLB areas who face serious digital divide (Haythornthwaite 2007) and digital literacy (Hargittai 2005) issues, and both recommender and online usability concerns are a major factor in our design considerations. We are also hoping to incorporate a stronger, but straightforward interface for exploration of the recommendation space, such as in the FindMe systems (Burke, Hammond, & Young 1997).

Ultimately, our goal is to provide a straightforward and solid foothold as a starting point in the face of information overload within a highly complex space. We do so in the context of an overall community program to help parents take advantage of the law of choice and to at least consider their decision options. In this way, we hope to facilitate the development of more active choosers, and, hopefully, better futures for our kids.

References

Ambler, J. 1994. Who benefits from educational choice: Some evidence from europe. *Journal of Policy Analysis and Management* 13:454–476.

Ballou, D.; Sanders, W.; and Wright, P. 2004. Controlling for student background in value-added assessment of teachers. *Journal of Educational and Behavioral Statistics* 29(1).

Basu, C.; Hirsh, H.; and Cohen, W. W. 1998. Recommendation as Classification: Using Social and Content-Based Information in Recommendation. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence, Madison, Wisconsin, USA, July 1998*, 714–720. AAAI Press.

Bollen, K., and Curran, P. 2006. *Latent curve models: A structural equation perspective*. Wiley-Interscience.

Brown, C. 2004. Choosing better schools: A report on student transfers under the no child left behind act. Citizens'Commission on Civil Rights. Washington, DC.

Burke, R.; Hammond, K.; and Young, B. 1997. The findme approach to assisted browsing. *IEEE Expert* 12(4):32–40.

Burke, R. 2002. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction* 12(4):331–370.

Chubb, J., and Moe, T. 1990. *Politics, Markets, and America's Schools*. Washington, D.C.: Brookings Institution.

Coons, J., and Sugarman, S. 1978. *Education by choice: the case for family control*. Berkeley: University of California Press.

Dunk, E. V., and Dickman, A. 2003. *School choice and the question of accountability: the Milwaukee experience*. New Haven: Yale University Press.

2004. No child left behind act: Education needs to provide additional technical assistance and conduct implementation studies for school choice provision. Report to the Secretary of Education GAO-05-7, U.S. Government Accountability Office, Washington, D.C.

Godwin, K., and Kemerer, F. 2002. *School choice tradeoffs: liberty, equity, and diversity*. Austin: University of Texas Press, 1 edition.

Hargittai, E. 2005. Survey measures of web-oriented digital literacy. *Social Science Computer Review* 23(3):371–379.

Hastings, J., and Weinstein, J. M. 2007. No child left behind: Estimating the impact on choices and student outcomes. Technical Report 13009, National Bureau of Economic Research Working Paper Series.

Haythornthwaite, C. 2007. Social facilitators and inhibitors to online fluency. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*.

Herlocker, J. L.; Konstan, J. A.; and Riedl, J. 2000. Explaining collaborative filtering recommendations. In *CSCW '00: Proceedings of the 2000 ACM conference on Computer supported cooperative work*, 241–250. ACM.

Hill, P. T. 2007. NCLB school choice and children in poverty. In Gamoran, A., ed., *Standards-Based Reform and the Poverty Gap*. Washington, DC: Brookings.

Howell, W. 2006. Switching schools? a closer look at parents' initial interest in and knowledge about the choice provisions of no child left behind. *Peabody Journal of Education* 81(1):140–179.

Hoxby, C., and Weingarth, G. 2005. Taking race out of the equation: School reassignment and the structure of peer effects. Technical report, Department of Economics, Harvard University, Cambridge, MA.

Kim, J., and Sunderman, G. 2004. Does nclb provide good choices for students in low-performing schools? Technical report, The Civil Rights Project at Harvard University, Cambridge, MA.

Konstan, J. A.; Miller, B. N.; Maltz, D.; Herlocker, J. L.; Gordon, L. R.; and Riedl, J. 1997. Grouplens: Applying Collaborative Filtering to Usenet News. *Communications of the ACM* 40(3):77–87.

Krulwich, B. 1997. LIFESTYLE FINDER: Intelligent user profiling using large-scale demographic data. *AI Magazine* 18(2):37–45.

Leland, S.; Godwin, K.; and Baxter, A. 2007. Examining school choice through the eyes of parents. In *Annual Meetings of the Midwest Political Science Association*.

McCaffrey, D.; Lockwood, J.; Koretz, D.; and Hamilton, L. 2003. Evaluating value-added models for teacher accountability. Technical report, RAND Corporation, Santa Monica, CA.

McCaffrey, D.; Lockwood, J.; Koretz, D.; Louis, T.; and Hamilton, L. 2004. Models for value-added modeling of teacher effects. *Journal of Educational and Behavioral Statistics* 29(1).

McDermott, K. A., and Jensen, L. S. 2005. Dubious sovereignty: Federal conditions of aid and the no child left behind act. *Peabody Journal of Education* 80:39–56.

Meier, D., and Wood, G. H. 2004. *Many Children Left Behind: How the No Child Left Behind Act Is Damaging our Schools*. Boston: Beacon Press.

Montaner, M.; López, B.; and de la Rosa, J. L. 2003. A taxonomy of recommender agents on the internet. *Artificial Intelligence Review*.

2001. No child left behind act of 2001. 2001-PL107-110.

Pazzani, M. J., and Billsus, D. 2007. Content-based recommendation systems. In Brusilovsky, P.; Kobsa, A.; and Nejdl, W., eds., *The Adaptive Web: Methods and Strategies of Web Personalization*. Berlin: Springer-Verlag. 325–341.

Raudenbush, S. 2004. What are value-added models estimating and what does this imply for statistical practice? *Journal of Educational and Behavioral Statistics* 29(1).

Resnick, P., and Varian, H. R. 1997. Recommender Systems. *Communications of the ACM* 40(3):56–58.

Rosenstein, M., and Lochbaum, C. 2000. Recommending from Content: Preliminary Results from an E-Commerce Experiment. In *CHI '00 Extended Abstracts on Human Factors in Computer Systems*, 291–292. ACM Press.

Rubin, D.; Stuart, E.; and Zanutto, E. 2004. A potential outcomes view of value-added assessment in education. *Journal of Educational and Behavioral Statistics* 29(1).

Sarwar, B.; Karypis, G.; Konstan, J.; and Riedl, J. 2000. Analysis of Recommendation Algorithms for E-Commerce. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, 158–167. ACM Press.

Sarwar, B.; Karypis, G.; Konstan, J.; and Reidl, J. 2001. Item-based Collaborative Filtering Recommendation Algorithms. In *Proceedings of the Tenth International Conference on World Wide Web*, 285–295. ACM Press.

Schafer, J. B.; Frankowski, D.; Herlocker, J.; and Sen, S. 2007. Collaborative filtering recommender systems. In Brusilovsky, P.; Kobsa, A.; and Nejdl, W., eds., *The Adaptive Web: Methods and Strategies of Web Personalization*. Berlin: Springer-Verlag. 291–324.

Schafer, J. B.; Konstan, J.; and Riedl, J. 1999. Recommender systems in e-commerce. In *EC '99: Proceedings of the 1st ACM conference on Electronic commerce*, 158–166. ACM.

Schneider, M.; Teske, P.; Roch, C.; and Marschall, M. 1997. Networks to nowhere: Segregation and stratification in networks of information about schools. *American Journal of Political Science* 41:1201–1223.

Singer, J., and Willett, J. 2003. *Applied longitudinal data analysis: Modeling change and event occurrence*. Oxford University Press.

Smyth, B., and Cotter, P. 2001. Personalized Electronic Programme Guides. *Artificial Intelligence Magazine* 22(2):89–98.

Soboroff, I., and Nicholas, C. 1999. Combining Content and Collaboration in Text Filtering. In *Proceedings of the IJCAI-99 Workshop on Machine Learning for Information Filtering*.

Southworth, S., and Mickelson, R. 2007. The interactive effects of race, gender and school composition on college track placement. *Social Forces* 86(2).

Tekwe, C.; Carter, R. L.; Ma, C.-X.; Algina, J.; Lucas, M. E.; Roth, J.; Ariet, M.; Fisher, T.; and Resnick, M. B. 2004. An empirical comparison of statistical models for value-added assessment of school performance. *Journal of Educational and Behavioral Statistics* 29(1):11–36.

Terveen, L.; Hill, W.; Amento, B.; McDonald, D.; and Creter, J. 1997. PHOAKS: A System for Sharing Recommendations. *Communications of the ACM* 40(3):59–62.