Public comment: NIST Draft Privacy Engineering Objectives and Risk Model

Christine Task
Purdue University Computer Science Department
ctask@purdue.edu

I'm a PhD candidate and my research focuses on privacy-preserving data-mining: data-analysis algorithms which include carefully calibrated noise or randomization in order to satisfy provable guarantees of individual privacy protection. Intuitively, privatized analysis ensures no single individual has a detectable impact on analysis results. Aggregated datasets (ex. histograms) can also be privatized to satisfy this guarantee, such that individual data is (with high probability) unidentifiable in the privatized data set.

I wanted to comment about a use-case which seems especially relevant to the 'Predictability' objective in your model. Learning Analytics is a recent and quickly developing field which applies data-mining and machine learning techniques to educational data. It's development was largely made possible by a 2008 weakening of FERPA (the federal law covering student privacy), which opened access to student records to "contractors, consultants, volunteers, and other outside parties providing institutional services and functions". Data can be shared, or sold, without notifying parents or students (potentially violating the Manageability objective).

On the one hand, Learning Analytics research is showing considerable promise in terms of deepening our understanding of learning and improving rates of student success. Data is being used for its generally intended purpose of improving education.

On the other hand, as you mentioned in the draft report, there may be a 'creepiness' factor. Below is a selection of abstracts from the 2014 LAK conference (http://dl.acm.org/citation.cfm?id=2567574). Many of them focus on classifying students by predicted likelihood of success/failure, using data collected from incidental interactions with course resources. It's debatable that parents and students could have predicted all of the potential uses for data gathered from their school activities.

- Every college student registers for courses from a catalog of numerous offerings each term. Selecting the courses in which to enroll, and in what combinations, can dramatically impact each student's chances for academic success. Taking inspiration from the STEM Academy, we wanted to identify the characteristics of engineering students who graduate with 3.0 or above grade point average. The overall goal of the Customized Course Advising project is to determine the optimal term-by-term course selections for all engineering students based on their incoming characteristics and previous course history and performance, paying particular attention to concurrent enrollment. We found that ACT Math, SAT Math, and Advanced Placement exam can be effective measures to measure the students' academic preparation level. Also, we found that some concurrent course-enrollment patterns are highly predictive of first-term and overall academic success.

- This paper describes a method we have developed to convert statistical predictive models into visual narratives which explain student classifications. Building off of the work done within the user experience community, we apply the concept of personas to predictive models. These personas provide familiar and memorable descriptions of the learners identified by data mining activities, and bridge the gap between the data scientist and the learning specialist.

- Assessment of reading comprehension can be costly and obtrusive. In this paper, we use inexpensive EEG to detect reading comprehension of readers in a school environment. We use EEG signals to produce above-chance predictors of student performance on end-of-sentence close questions. We also attempt (unsuccessfully) to distinguish among student mental states evoked by distracters that violate either syntactic, semantic, or contextual constraints. In total, this work investigates the practicality of classroom use of inexpensive EEG devices as an unobtrusive measure of reading comprehension.

- As providers of higher education begin to harness the power of big data analytics, one very fitting application for these new techniques is that of predicting student attrition. The ability to pinpoint students who might soon decide to drop out of a given academic program allows those in charge to not only understand the causes for this undesired outcome, but it also provides room for the development of early intervention systems. While making such inferences based on academic performance data alone is certainly possible, we claim that in many cases there is no substantial correlation between how well a student performs and his or her decision to withdraw. This is specially true when the

overall set of students has a relatively similar academic performance. To address this issue, we derive measurements of engagement from students' electronic portfolios and show how these features can be effectively used to augment the quality of predictions.

- Human body-language is one of the richest and most obscure sources of information in inter-personal communication which we aim to re-introduce into the classroom's ecosystem. In this paper we present our observations of student-to-student influence and measurements. We show parallels with previous theories and formulate a new concept for measuring the level of attention based on synchronization of student actions. We observed that the students with lower levels of attention are slower to react then focused students, a phenomenon we named "sleepers' lag".

Two questions:

Could analyses like these be a violation of Predictability, even though they direct educational data towards educational objectives?

Would the problem be resolved if analysts worked with rigorously privatized data (satisfying the guarantee I described above), rather than directly accessing raw student data? If alice is part of a privatized data set in which she *cannot* be individually identified, but that data-set is used for an application she could not predict and would find objectionable... is this still a violation of her privacy?