

2008

Unleashing the Usefulness of Educational Resources through Mining of Educational Metadata


Anne R. Diekema
Utah State University

Jennifer A. Bailey
Syracuse University

Blythe A. Bennett
Syracuse University

Holly Devaul
Digital Learning Sciences

Follow this and additional works at: http://digitalcommons.usu.edu/itls_facpub

 Part of the [Educational Assessment, Evaluation, and Research Commons](#), and the [Instructional Media Design Commons](#)

Recommended Citation

Diekema, A.R., Bailey, Jennifer A. and Bennett, Blythe A., Devaul, Holly. Unleashing the Usefulness of Educational Resources through Mining of Educational Metadata. In: Proceedings of the 2008 Open Education Conference. Logan, UT, September 24-26, 2008.

This Conference Paper is brought to you for free and open access by the Instructional Technology & Learning Sciences at DigitalCommons@USU. It has been accepted for inclusion in ITLS Faculty Publications by an authorized administrator of DigitalCommons@USU. For more information, please contact dylan.burns@usu.edu.



Unleashing the Usefulness of Educational Resources through Mining of Educational Metadata

Anne R. Diekema, Utah State University
anne.diekema@usu.edu

Jennifer A. Bailey and Blythe A. Bennett, Syracuse University
jabail01@syr.edu, babennet@syr.edu

Holly Devaul, Digital Learning Sciences
devaul@ucar.edu

Introduction

While there is a strong movement to develop new educational resources to bring students to the competencies represented by educational content standards, it is recognized that there are vast repositories of educational resources already developed that are suitable to address those competencies. However, these resources need to be indexed by national and state standards to make them accessible for teachers who are increasingly required to teach to certain educational standards (Diekema and Chen, 2005). In the early 1980s, a perceived crisis in the American education system encouraged the creation of national standards by professional subject-area organizations such as the National Council for Teachers of Mathematics. These national standards aimed to clearly define what students in certain grade levels are expected to know in core subject areas (Ratvitch, 1995). Eventually all states followed suit and published their own educational standards, often using the national standards as a guideline.

Adding state standard information of every state to each resource is a large task, especially when done completely manually. Manual standard to standard alignment efforts (e.g. Align to Achieve) proved too difficult to maintain as the standards exist in all core subject areas; on national, state and local levels; and are revised regularly. Each set of standards utilizes discrete language, differing grade bands, distinct organizational structures and different levels of specificity in the coverage of a particular standard. To remedy these problems the Center for Natural Language Processing (CNLP) at Syracuse University has created a technology (Standards Alignment Tool – SAT) for automatically aligning state standards and national standards (Diekema et al., 2007).

SAT was developed to improve the ability of teachers to locate resources in the National Science Digital Library that support standards-based instruction. A multitude of state and national standards sources exist, but typically learning resources are aligned only to one source, either a national standard or the state or local standard relevant to the resource creator or collection builder. This presents a challenge to digital library developers who wish to support a broad base of users, and to digital library users who often have information needs specific to their state or district. To extend the utility of the information available for a particular resource, SAT creates standard-to-standard equivalency relationships or alignments, utilizing algorithms from the fields of machine learning, information retrieval, and natural language processing. The intended result of this mechanism is to provide library users access to standards information about learning resources from any standards source available, regardless of which source was used in cataloging the resource. This capability can be integrated into search and display components of digital libraries and online repositories to provide on-the-fly alignments of resources to standards. It is critical that the standard-to-standard equivalencies be determined to be sufficiently robust prior to integration into any live system that is fully automatic.

The key concept in SAT is the existence of a crosswalk. Rather than creating individual alignments of each standard to every other standard, standards are only aligned to the crosswalk standard or exchange standard. The connection between different state or national standards can be made via that crosswalk. SAT uses a relatively small set of manually-determined standard to standard alignments to learn classifiers for a crosswalk. These learned classifiers are then used to align new standards to the crosswalk, thus allowing alignments between any state standard. One of the challenges facing this work has been that the manual standard to standard alignment training data is rather sparse. A large number of standard categories have no examples or very few examples, making it difficult for the system to learn. An internal evaluation of SAT has shown that while the text classification performance, to classify standards to the crosswalk, was acceptable from a classification standpoint, higher levels of accuracy are required for an automatic alignment system. (Yilmazel, 2007). These findings were echoed in a technical report by Digital Learning Sciences (DLS) whose SAT evaluation indicated that the standards alignments provided in the top ranked returns of SAT are not broadly and sufficiently reliable for immediate automatic assignment within a digital library search system that is translating National Science Education Standards to state standards.

This paper explores whether it is possible to exploit existing manual standards assignments by mining the groups of standards that have been assigned to a particular resource. In other words, rather than requiring explicit manual alignments between equivalent standards, this preliminary research is trying to use the assignment of standards as metadata to resources to determine which state standards might be equivalent. An increasing number of manual standards assignments is becoming available, possibly making this approach a viable and sustainable option. The ultimate goal of this research is to establish an automatic correlation between standards based on their shared occurrence. At the initial stage we are only considering groups of standards that were assigned to the same lesson plan. Eventually we'll take into the co-occurrence statistics of standards across the entire corpus of lesson plans.

The data

For the past year a team of catalogers made up of School Library Media librarians as well as teachers has been assigning state standards to approximately 4,000 resources on the Thinkfinity.org web portal which offers, among other things, free educational resources such as lesson plans to parents, students, and teachers in all of the core subject areas. The goal of this cataloging effort was to make it possible for Thinkfinity users to find relevant resources pertaining to a particular standard that they are working on in class. The catalogers manually assigned at most two standards per state (if available) to each lesson plan, so the maximum number of assigned standards is 100 per lesson plan.

Both state and national standards are typically organized in a hierarchical fashion, starting at the apex, with the general subject (Math), the areas within that subject next lower (e.g. Number and Operations Standard), the third level standard (understand meanings of operations and how they relate to one another) next lower, and finally the benchmark (understand the effects of adding and subtracting whole numbers) as the leaf level. The catalogers assigned the standards at the most detailed level, the leaf level. Table 1 below shows a selection of leaf level standards that were assigned to the same lesson plan.

Table 1. Leaf level standards that share the same lesson plan.

Demonstrating computational fluency for basic addition and subtraction facts with sums through 18 and differences with minuends through 18, using horizontal and vertical forms
perform basic arithmetic functions, make reasoned estimates, and select and use appropriate methods or tools for computation or estimation including mental arithmetic, paper and pencil, a calculator, and a computer;
recalling basic addition and subtraction facts, sums to 20, and corresponding subtraction facts efficiently
Concept 2: Numerical Operations Understand and apply numerical operations and their relationship to one another.: State addition facts for sums through 18 and subtraction for differences with minuends through 9 or less.
Demonstrate computational fluency (accuracy, efficiency and flexibility) in addition facts with addends through 9 and corresponding subtractions
Know the addition facts (sums to 20) and the corresponding subtraction facts and commit them to memory.
Use concrete objects to determine the answers to addition and subtraction problems (for two numbers that are each less than 10).
read and write whole numbers and know place-value concepts and numeration through their relationships to counting, ordering, and grouping;
demonstrate conceptual meanings for the four basic arithmetic operations of addition, subtraction, multiplication, and division;
Count, adding one more to the previous number and group and count by ones and tens.
Develop fact families of basic facts using the inverse relationship of addition and subtraction.
recognize inverse operations: subtraction/addition and division/multiplication;
decompose and recompose whole numbers using addition and subtraction;

Clearly, these standards that share the same lesson plan are related but some more so than others. We can recognize certain concepts that these standards share: numbers, numerical operations such as addition and subtraction. Only in the event that the standards share identical vocabulary and syntax can we safely say that two standards are completely identical. In all other cases it is necessary to quantify how similar two standards are.

On similarity or Strength of Fit and co-occurrence

The Dublin Core Metadata Initiative (dublincore.org) works towards describing resources through the use of metadata to make these resources better accessible. One of the (proposed) metadata elements describes the correlation between the objectives of educational resources and achievement standards that have been associated with the resource. One aspect of the correlation metadata is the *correlation factor* or *strength of fit* between resource and standard. Similarly we can use the correlation factor to describe the strength of fit between standard A and standard B. Dublin Core recognizes five different types of correlation: major correlation, minor overlap, broad correlation, narrow correlation, and exact correlation. The difference between these correlation types is very subtle. What is the difference between a major correlation and a broad correlation for example? A major correlation has major overlap of concepts of interest between the two items whereas a broad correlation exists when standard A covers all of the concepts of interest in standard B as well as concepts of interest not found in standard B. Can we assume that a major correlation is also a broad correlation in case the number of additional concepts is small? Philosophical questions such as these are not easy to answer and perhaps the standards metadata strength of fit vocabulary is not suitable for an automated process as the difference between the correlation types is semantically complex.

Computer and information scientists have developed various metrics to establish the similarity between two text strings. The Levenshtein distance metric for example measures how many changes or edits are needed to change one string into another. The more edits are needed, the less similar the two strings (Navarro, 2001). The problem with this and other string similarity metrics however is that they cannot tell us whether two strings are conceptually equivalent unless they are using a lot of the same words. For example, the following two standards are very similar but only have a single word in common:

Demonstrate knowledge of basic facts in four basic operations.

Recognize inverse operations: subtraction/addition and division/multiplication;

To establish whether two strings are conceptually the same requires intellectual effort and perhaps additional sources of evidence. Just the fact that two standards are assigned to the same lesson plan is not enough to establish their similarity. A lesson plan might cover two entirely different concepts that occur in the two separate standards, both assigned to that lesson plan. To weed out these spurious correlations between standards it is important to measure the co-occurrence statistics. The more often standards are assigned to the same lesson plan, the more likely it is that they are similar. Combining the string similarity metrics and the co-occurrence statistics will give us a better indication as to the probability that two standards are similar.

Conclusions and future work

An initial exploration of the data shows that automatically establishing a connection between state standards that have been assigned to the same lesson plan is quite challenging. Whether this connection can be established fully automatically based on the multiple sources of evidence suggested in the paper is currently being investigated. In addition to possibly establishing a connection between similar state standards we will need to investigate how to add strength of fit information to these correlations, making them more expressive. Being able to establish these correlations automatically by mining manually assigned metadata will make it possible for open source repositories to add standards data to their collections quickly and efficiently, increasing the accessibility of their educational materials in an affordable manner.

References

- Align to Achieve. <http://www.aligntoachieve.org/> accessed June 30, 2008.
- Center for Natural Language Processing. <http://www.cnlp.org/> accessed June 30, 2008.
- Devaul, H. & Ostwald, J. (2007). An Evaluation of the Standard Alignment Tool (SAT): Selected National to State Standards Alignments. Digital Learning Sciences Technical Report. Boulder, CO. November 21, 2007. 12p.
- Diekema, A. R., Yilmazel, O., Bailey, J., Harwell, S. C. and Liddy E. D. Standards Alignment for Metadata Assignment. (2007) In: *Proceedings of the Joint Conference of Digital Libraries (JCDL 2007)*. Vancouver, British Columbia, June 18-23.

- Diekema, A.R. and Chen, J. (2005) Experimenting with the Automatic Assignment of Educational Standards to Digital Library Content. In: *Proceedings of the Joint Conference of Digital Libraries*. Denver, Colorado, June 7-11.
- Digital Learning Sciences. <http://dlsciences.org/> accessed June 30, 2008.
- Dublin Core Metadata Initiative: Correlation Resource.
http://dublincore.org/educationwiki/Correlation_20Resource/ accessed August 22, 2008.
- Navarro, G. (2001) A guided tour to approximate string matching. *ACM Computing Surveys*, Vol. 33, No. 1, pp. 31–88.
- Ratvitch, D. *National Standards in American Education: A Citizen's Guide*. Brookings Institution Press, 1995, 242p.
- Thinkfinity. <http://www.thinkfinity.org/> accessed August 22, 2008.
- Yilmazel, O., Balasubramanian, N., Harwell, S. C., Bailey, J., Diekema, A. R., & Liddy, E. D. (2007). Text Categorization for Aligning Educational Standards. *Hawaii International Conference on System Sciences*, Waikoloa, Hawaii, January 3-6.