The Use of Metadata for the Identification and Retrieval of Resources for K–12 Education

John S. Kendall,
Mid-continent Research for Education and Learning, USA
jkendall@mcrel.org

Abstract

Educators and publishers in the field of K–12 education have long had an interest in identifying and retrieving curriculum materials by the learning objectives that these materials serve. A review of the efforts undertaken to accomplish this goal highlights a number of continuing problems and informs a set of criteria for the development of a schema that would be useful for identifying and retrieving educational resources. One such schema and its value space is offered as a solution, which exploits a unique resource developed at Mid-continent Research for Education and Learning (McREL).
***Keywords**: metadata schema, learning objectives, education standards*

1. Standards and Resources

The advent of standards-based education (SBE) has considerably heightened interest in classroom resources. Facing increased accountability for student learning of specific knowledge and skills, educators require classroom activities, lesson plans, units, and assessments that are targeted to established content and performance standards. (For a discussion of these topics relative to SBE, see Gaddy, Dean, & Kendall, 2003[1]). Delivering such education resources targeted to specific standards cheaply and efficiently poses a significant problem. First, there is no single, national set of standards that all adopt to which the resources could be matched. Although similar academic objectives might be found in classrooms all over the U.S., in fact, internationally as well, no two states have identical ways of describing content standards, and not all school districts adopt state standards, but might construct some variation thereof. On the resource side, countless materials are available either from publishers or clearinghouses that address many of these objectives, yet these resources are marginally useful unless they can be provided to educators within the framework of their own set of standards.

A number of efforts have been undertaken to address the difficulty of pairing education materials with the standards educators must teach. These approaches all require some level of content analysis in order to establish a relationship between resources and the standards they purport to address. The efforts center around one of two approaches: maximizing the speed of the matching process, or maximizing the benefit of matches once they are made. What follows is a survey of the current work being done, what might be learned from it, and a proposal based on that learning.

2. Solutions: by Search or by Design

Education publishers, in order to persuade educators of the value of their products, have for many years attempted to show how their materials serve state-identified curriculum goals. Typically, publishers have contracted with small firms to analyze their products against state standards. The process, time consuming and inefficient, requires that the standards for each state be reviewed to determine whether and where a specific curriculum resource — a unit or lesson plan, for example — might be said to support the state standard.

Such a labor-intensive challenge was ripe for some assistance from technology. Technological solutions tend to fall into one of two major streams of work: (1) using computer processing to assist in the comparison of resources and standards, either through simple term comparisons or through more sophisticated search strategies, or (2) using a design strategy to maximize the value of experts' decisions regarding the possible matches between resources and standards.

2.1. Searching For Solutions

In the early to mid 1990s, it was common to find proposals that attempted to ease the labor burden of content mapping by resorting to computer-intensive searches. These techniques relied primarily on searching for like terms among learning objectives and the instructional materials that might support them. It soon became apparent, however, that a number of learning objectives, especially those that describe generalizations or principles, might not have unique terms or phrases in common, and thus would not be retrieved using this method. More difficult is that learning objectives might share many common terms, yet their specific requirements might vary from one state to another, and from one grade to the next. Most problematic, however, teachers do not have the luxury to sort through lists of potential matches to determine what meets the particular objective for which they are responsible. Described differently, teachers require very high precision in the accuracy of the match between materials and

objectives, even as the sheer number of potential resources appears to grow exponentially each year. Although search algorithms have become considerably more sophisticated, accurate retrieval of electronic resources remains a significant concern [2].

More sophisticated approaches develop a "look-up" table of the synonyms or synonymous phrases that are common to a given state's standards. Another approach is simply to collect all the benchmarks from all the states relative to topic ideas, so that when the cataloguer is in the topic area he or she will more likely be near the target. In the view of some, such solutions make the likelihood of missing matches too high a cost for the benefits of efficiency promised. Currently, such searching is usually considered a first-stage help that surfaces the more obvious matches of content; the standards still need close content analysis. Refinements in this area, as well as studies to determine actual efficiency, are needed. A current effort underway evaluates how effectively a machine can create matches once it has been trained on thousands of exemplary matches made by content analysts. (*StandardConnection, a* service project of the National Science Foundation's National Science Digital Library, maps NSDL Educational Objects to Content Standards [3])

2.2. Solution by Design: The Intermediary

Despite any advantage gained from more efficient searching for matches, the problem remains significant given the many state standards and the many thousands of resources to be mapped to them. In the mid 1990s, Michael Jay, currently Senior Advisor at KnowledgeQuest, began a commercial enterprise, called MediaSeek, to address the need by applying the concept of a relational database. A relational database provides a way of associating information of different kinds by identifying the feature or features that they share in common. In a system used to facilitate retrieval of academic content and resources, the feature in common is the learning objective, a description of expected student knowledge or skill. Simply put, state standards have embedded within them learning objectives, which are variously called benchmarks, indicators, or performance descriptions. Education resources (lesson plans, activities, assessments, and the like) are designed to address learning objectives. This usually means the learning objectives are either stated outright or can be inferred by studying the resource. Thus, standards and education resources can be joined by identifying the learning objectives they share in common. The relational database that links state standards and education resources has since come to be called an *intermediary*.

The intermediary, or special-purpose relational database, provides clear advantages over pairing resources one item at a time to standard after standard. Once an education resource is mapped to the learning objective it serves, and all state standards that contain the learning objective are likewise mapped to it, then a link is established between that educational resource and all state standards that serve the same learning objective. Ideally, this link is established so that the resource need not be reviewed against every standard in order to be well mapped to any standard. For example, let us say that a common learning objective is that students should be able to write a well-formed paragraph. We assign the learning objective a code, say 1.3.3, to distinguish it from among a list of other learning objectives. If an activity or lesson plan is identified as having that learning objective as its goal, it is assigned the code 1.3.3. Similarly, the standards from education agencies, such as state departments of education, school districts, and national subject-area organizations, are reviewed and assigned codes to reflect the learning objectives they describe. Thus, a standard that indicates students should be able to write a well-formed paragraph is likewise assigned the code 1.3.3. As a result, the standard that describes a learning objective has been associated with an educational resource that supports the teaching of that learning objective. The educational resource has been identified as appropriate for the standard, despite the fact that no one has had to compare that standard and resource to determine whether the fit is appropriate.

3. Metadata and the Intermediary

Readers familiar with the concept have likely recognized that the learning objectives in the intermediary function primarily as *metadata*, that is, their primary use is to describe other data, namely the expectations for students articulated in state standards or the student expectations that are the focus for educational resources. Although there are metadata schemas that describe educational resources to some degree, such as Learning Object Metadata (IEEE) [4] and the GEM extension of Dublin Core [5], such schemas do not include an element or element set that provide specific, unambiguous information regarding the learning objectives such resources might serve. This lack is likely owing to the absence of a set of national academic standards. In Great Britain, which has an established national curriculum, an element set has been developed that allows for the identification of content even to the level of numbered paragraphs within the national curriculum [6]. The need in U.S. education, or for any system that might be used internationally, is for a mechanism of content identification that does not rely on and is not restricted by the curriculum of a single country.

Over the last seven or eight years much has been learned about the use of intermediaries; thus, much has been learned that might inform thinking about how to develop an element set that describes learning objectives. The balance of this paper consists of a description of what has been learned from

work done at Mid-continent Research for Education and Learning (McREL) and elsewhere on intermediaries. From this comes a set of desiderata for such a schema, or, considered differently, a set of criteria for judging the likely success or usefulness of such a schema. Finally, a schema is proposed that appears to meet these criteria.

3.1. The Need for Publicly Accessible Metadata

The process required to identify the learning objective that captures the intent of a standard or a curriculum resource is variously called "correlating," "mapping," or simply, "tagging." In concrete terms, the work, typically undertaken by a trained subject-matter expert, entails assigning a code to standards (national or state) and educational resources (curricula or assessments) that uniquely identifies the learning objective that is within the intermediary. The materials need be mapped only once — to a learning objective — to be linked to every state standard that has likewise been mapped to that objective. Mr. Jay's enterprise made significant progress in the use of this system, both in the mapping of state standards and in mapping to a number of publishers' resources. However, the system is not in common use, principally for two reasons. First, MediaSeek's database of objectives (called the Knowledge Base) is proprietary and is not open to public scrutiny. Thus, the quality of the database — specifically, whether the academic objectives that make up the database adequately represent the content of each discipline — is unknown. Second, the Knowledge Base has been purchased by a single commercial enterprise, and is not publicly available for use. These limitations significantly impact the potential of the Knowledge Base as an indexing device for common use.

The value of a set of learning objectives to which resources and standards could be mapped became clear through the relative success of MediaSeek. Also clear, however, is that this set of objectives should be open for inspection and accessible to all. Such a criterion all but excludes any lists developed by commercial entities that hold them as a proprietary interest.

These two particular issues, that is, that the database of content is not available for public inspection and the process of development is not known, are problems that are alleviated to a great extent in a design developed at McREL. The database of content used for mapping is based on an edition of McREL's *Content Knowledge: A Compendium of Standards and Benchmarks for K-12 Education* [7], available online at www.mcrel.org/standards-benchmarks/. For each of the more than 4,100 statements of knowledge and skill in the *Compendium*, a reference is provided to every document of national significance that addresses the same content that the statement (termed a *benchmark*) describes. Thus, the means by which the objectives are developed are available to any interested party. The database itself is also publicly accessible. In addition, the current edition of the database represents the synthesis of content of more than 137 significant documents over 14 content areas. Thus, it seems very likely that a significant portion of the content likely to appear in state standards or to be the focus of educational resources is present in this database.

Thus far, we can establish a number of desiderata for this metadata schema:

- The statements that populate the metadata (analogous to the learning objectives within the intermediary) are accessible to all.
- The method and sources used to develop the items within the metadata is available to all.
- There is a high degree of likelihood that the metadata is comprehensive in its coverage of learning objectives in education.
- The set of metadata is easily extensible to address descriptions of new knowledge and skill

3.2. Problems of Varying Grain Size

The most significant problem that faces tagging approaches is the narrow range of content description that is available to tie standards to the curriculum resources. The difficulty arises because both standards and resources can and do vary significantly across states and resources in terms of the level of specificity, or grain size, at which they address or describe academic content. **Table 1** presents English language arts statements from three state standards documents:

Table 1. Sample English Language Arts Content

Source	Grade(s)	Content
New Mexico	7	narrate an account (e.g., news story, historical episode) that creates a coherent organizing structure appropriate to purpose, audience, and context and that orients and engages the reader
Virginia	6	write narratives, descriptions, and explanations
Louisiana	6–8	using narration, description, exposition, and persuasion to develop various modes of writing (e.g, poems, letters, and essays)

Note that the New Mexico content has a broad focus related to narrative writing: it includes audience, purpose, and context. The Virginia document focuses on simply writing of

different types, of which narrative writing is one. The Louisiana benchmark includes the writing of different types that Virginia lists, but adds persuasion as a type of writing. In addition, the Louisiana benchmark includes modes of writing (such as poems, letters, and essays). This degree of disparity among state standards is not uncommon. In an area like the social studies, in fact, the differing levels of granularity are common and even more divergent than reflected in the samples presented here. Unfortunately, the use of a mapping scheme, rather than resolving the issue of granularity, can further complicate it. This is because the learning objectives used to do the mapping, or form the database relation, can themselves be at a different level of granularity than the content against which they are being mapped. For example, Table 2 displays a benchmark that reflects content addressed in the above-cited state benchmarks that is also found in the McREL benchmarks:

Table 2. McREL Benchmark 1.3.7

Unique Identifier	Grades	Content
1.3.7	6–8	Writes narrative accounts, such as short stories (e.g., engages the reader by establishing a context and otherwise developing reader interest; establishes a situation, plot, persona, point of view, setting, conflict, and resolution; develops complex characters; creates an organizational structure that balances and unifies all narrative aspects of the story; uses a range of strategies and literary devices such as dialogue, tension, suspense, naming, figurative language, and specific narrative action such as movement, gestures, and expressions; reveals a specific theme)

Clearly, the benchmark addresses some of the important aspects of the content found in the state benchmarks listed earlier, but it also introduces still more content that is not necessary, and in fact, might introduce confusion if the entire benchmark is used as the intermediary link between standards and resources. Thus, the use of the benchmark as a learning objective for direct links between standards and education resources is problematic. Mapping the state standards using this benchmark as the learning objective would be better represented with a Venn diagram than a linking arrow. In an attempt to remedy this problem, the author developed a design to narrow the scope of the benchmark when it is used as a learning objective within the intermediary. This technique employs vocabulary words, derived from the benchmark, to help focus the mapping. Thus, for the benchmark identified in Table 2, the associated set of vocabulary terms in Table 3 could be used to improve the mapping.

Table 3. Vocabulary Associated with McREL 1.3.7

I.D.	Vocabulary	I.D.	Vocabulary
A	narrative	H	conflict
B	context	I	resolution
C	situation	J	character
D	plot	K	organizational structure
E	persona	L	literary device
F	point of view	M	narrative action
G	setting	N	theme

The vocabulary term is added in order to refine or narrow the set of content within the McREL benchmark that is used as a map. Thus, for the sample state content in Table 4, an appropriate mapping would be *1.3.7.ABK*.

Table 4. Sample State Content Mapped as 1.3.7.ABK

Source	Grade	Benchmark
New Mexico	7	[Students] narrate an account (e.g., news story, historical episode) that creates a coherent organizing structure appropriate to purpose, audience, and context and that orients and engages the reader

The number 1.3.7 is a unique identifier assigned to the benchmark; the letters A, B, and K are the letter codes designating the assigned vocabulary. This mapping indicates that the primary content of interest, as the McREL benchmark, is narrative (A), context (B), and organization (K). Other mappings to other McREL benchmarks would be necessary. For example, in order to address the content in the New Mexico benchmark that deals with audience and purpose, a different McREL benchmark would be used. Again, selective vocabulary would be added in an attempt to refine the content.

The process of mapping a significant number of state standards documents has made clear a number of problems with this technique. First, although it is true that the benchmarks are too broad for effective mapping, the

vocabulary terms themselves are proving to be too narrow but, paradoxically, also too ill defined. For example, consider the benchmark in Table 5.

Table 5. A Sample of Broadly Stated Content

Source	Grade	Content
Virginia	6	The student will write narratives, descriptions, and explanations.

In order to map to the content in Table 5 regarding the writing of narratives and using the benchmark and vocabulary set displayed in Tables 2 and 3, the following map code would be used: 1.3.7.A. This must be the case even though there are a number of other benchmarks in the McREL *Compendium* that also address narrative writing. This technique — that is, the selection of one benchmark to address narrative writing when a number of others are available that address aspects of narrative writing — reflects a mapping technique used at McREL in part to limit the number of tags required for a given standard or resource. It became clear during the development of the mapping process that if all references to narrative writing were tagged (i.e., mapped to a McREL benchmark plus vocabulary term), then, given the number of eligible benchmarks, many benchmarks within a state document could easily require several dozen tags each. In order to address problems like this, McREL established a tagging guide. In this guide are such rules as this: if tagging only to address the idea of "narrative writing" (as is the case for part of the Virginia benchmark above), one should always use 1.3.7.A. (The rule itself is more specific; this is a rough approximation for the sake of an example). Such a rule was adopted despite the fact that there are other benchmarks that address other types of narratives than simply short stories (which is the ostensible topic of 1.3.7). For example, there are benchmarks on biographical and autobiographical sketches, which are a type of narrative writing. This approach was also driven by practical concerns, for example, an interest in keeping the number of mappings to a manageable number per state benchmark.

The difficulty of this approach is that it requires all who tag to be aware of and follow these rules. Sometimes the rules are necessarily arbitrary, because any of a number of possible choices is equally valid. If such a decision must be made, then the critical factor is that everyone who codes the material is aware of, understands, and accepts these rules and tags accordingly. The tagging guides used at McREL are many pages long, highly detailed, and require careful attention. Each new rule invites the possibility of error, especially because arbitrary rules are difficult to remember. But without such rules, the tagging cannot be done consistently or with a reasonable level of consistency. Thus, when the use of a benchmark plus vocabulary is too narrow for some mapping, the mapping system has had to accommodate this by creating arbitrary rules intended to expand the reach of the mapping. In the example provided here, the rule is an attempt to expand the reach of a single benchmark that addresses a narrow aspect of narrative writing so that it can be used to address the larger category of narrative writing. Clearly, there should be a better way to address problems of this type.

It should be noted that the problem is exacerbated for an intermediary, such as the Knowledge Base, that is constructed of nothing but very narrowly focused descriptions of student knowledge and skill. In such a case, state standards like the examples provide in Table 1 could require a very large number of maps to address all of the content they encompass.

Another difficulty arises when vocabulary terms, though intended to narrow the focus of the benchmark, are too broad. This problem becomes clear in light of the variety of uses the mapping system might serve. At the broadest level, the content in any one state might be compared to that of another using the intermediary with a set of core learning objectives as the common field. More commonly, the goal is to provide educators with the means to retrieve appropriate curriculum resources by selecting the state standard for which they need classroom materials. A growing interest, however, is the mapping and retrieval of assessment items through this system. Such mapping would provide a powerful means for helping educators find assessment items that are directly aligned with instruction. Such an alignment, in order to be effective, must be fairly exact, and therefore at the finest level of specificity. For example, consider the state content identified in Table 6.

Table 6. Sample State Mathematics Content

Source	Grade	Content
Louisiana	4	demonstrating number sense and estimation skills, giving particular attention to common equivalent reference points (i.e., 1/4 = 25% = .25; 1/2 = 50% = .5; $1 = 100%, etc.).

The benchmark in Table 6 shares much in common with the McREL benchmark (Table 7) and its associated vocabulary (Table 8). The tagging code for the Louisiana benchmark in the method currently being described, that is, using the benchmark identifier and a selection of the associated vocabulary terms would be 2.2.2.ABCD. This code is a fairly accurate identification of the content within the Louisiana benchmark. However, if an assessment item is likewise marked as 2.2.2.ACD, a code that suggests there should be a direct link to the Louisiana benchmark, we still cannot be certain what aspect of the Louisiana content that the item would assess. The item might assess either of the

following aspects of the benchmark:

- Student converts percents to decimals
- Student knows the equivalence of common percents and decimals

Table 7. McREL Mathematics Benchmark 2.2.2

Unique Identifier	Grades	Content
2.2.2	3–5	Understands equivalent forms of basic percents, fractions, and decimals (e.g., 1/2 is equivalent to 50% is equivalent to .5) and when one form of a number might be more useful than another

Table 8. Vocabulary Associated with McREL 2.2.2

I.D.	Vocabulary	I.D.	Vocabulary
A	percent	E	whole number
B	fraction	F	reduced form
C	decimal	G	common denominator
D	equivalent forms of a number		

Just as we have seen that vocabulary is not broad enough to address some tagging difficulties (as in the case of mapping to the concept of narrative), here we see that vocabulary is not specific enough for other purposes.

3.3. Problems in Retrieval

A primary difficulty with current mapping procedures, then, is that the indexing system is at a level of granularity that is in some cases too broad, at other times not broad enough to map well to content. Although the addition of vocabulary terms (described in the section above) does help to increase the level of specificity used to map materials, there is still a level of specificity, for example that required to map assessment items, that cannot be achieved using this method. More problematic still is that such a system fails completely to reflect the fact that content can differ markedly in levels of granularity across as well as within state standards and curriculum resources.

The problem is further compounded in use, however. Because the academic objective that is used as the indexing term is fixed in its granularity, it likewise limits the user to a narrow range of generality at which content can be retrieved. For example, a given resource can not be mapped as appropriate to the topic of angles, because *angles* can only be mapped as part of particular statements of academic objectives that specify what it is about angles that students should know or be able to do. Nor can the user simply request all resources that address the topic of angles, but only those resources that address this or that aspect of angles. Simply put, specific mapping is useful for certain purposes, but it precludes any more general retrieval.

At a more technical level, this inability to map "upwards" at a more general topic level can have a significant impact on how well a standard or curriculum resource is represented and, thus, on how well or appropriately it can be joined to other standards or resources in the database. In some cases, it is appropriate to include a number of mappings to capture the breadth of a resource or standard. However, there is a practical limit, in terms of user patience and the expense of mapping and retrieval, to the number of academic objectives that can be mapped to a standard or resource. Without degrees of generality available in a mapping system, there is no solution to the problem of accurately capturing, both in the mapping and for the retrieval, the scope of some content. In light of the difficulties discussed in this section, there appear to be two additional desiderata for a useful metadata schema:

- Specific, unambiguous, & uniform in its descriptions of content
- Available at various levels of granularity

4. Proposed Schema

The problem identified thus far can be significantly ameliorated through a hierarchical content mapping system. Such a system is structured so that content can be mapped and later retrieved at *any of a number of levels of specificity*. In this system, there are *four* levels of increasing specificity available; each level subsumed by the content in the level above it. This fact has two significant consequences. First, mappings can be done more accurately, because there are four layers of granularity from which to choose when making a map. (Currently, there is a single level of granularity used, which is coded as a concatenation of a benchmark with vocabulary identifiers.) Second, once a map is made in such a system, it is nested within a conceptual hierarchy that can be used to provide, inform, and guide retrieval of content. An example from the English Language Arts is provided in Table 9, showing all levels of a hierarchical structure that can be used to index content.

In such a system, because any level can be used to map content, a significant number of problems can be resolved that have to this point required a complex rule-making system and detailed tagging guides. For example, recall the case of the Virginia benchmark in Table 5 that referred generally to narrative writing. "The student will write narratives, descriptions, and explanations."

Table 9. A Four-Tiered Structure for Identifying Content

Level	Name	I.D.	Content
I	Standard	1	Uses the general skills and strategies of the writing process
II	Topic	1	Literary/Narrative writing
III	Benchmark	3.7	Writes narrative accounts, such as short stories (e.g., engages the reader by establishing a context and otherwise developing reader interest; establishes a situation, plot, persona, point of view, setting, conflict, and resolution; develops complex characters; creates an organizational structure that balances and unifies all narrative aspects of the story; uses a range of strategies and literary devices such as dialogue, tension, suspense, naming, figurative language, and specific narrative action such as movement, gestures, and expressions; reveals a specific theme)
	Knowledge/ Skill statements	1	Writes narrative accounts
		2	Writes short stories
		3	Engages the reader by establishing a context
		4	Develops the reader's interest
		5	Establishes a situation
		6	Establishes a persona
		7	Establishes a point of view
		8	Establishes a setting
		9	Establishes a conflict
		10	Establishes a resolution
		11	Develops complex characters
		12	Creates an organizational structure that balances the narrative aspects of the story
		13	Creates an organizational structure that unifies narrative aspects of the story
		14	Uses narrative strategies
		15	Uses a range of literary devices
		16	Uses dialogue
		17	Uses tension
		18	Uses suspense
		19	Uses figurative language
		20	Uses specific narrative action
		21	Uses movement in narrative writing
		22	Uses gestures in narrative writing
		23	Uses (facial) expression in narrative writing
		24	Reveals a specific theme

Using the content in Table 9, we recognize the map that addresses narrative writing is sufficient at the topic level:
1.1: Literary/Narrative writing

The topic level organizes benchmarks beneath a standard. In this case, the topic "Literary/Narrative writing" also encompasses the benchmarks on autobiographical and biographical writing. Using this system of mapping allows us to map the content at a level of generality appropriate to that found in the state benchmark.

Conversely, if content is quite specific, the most specific statement is also available. For example, an assessment item on converting fractions to percents can be easily addressed by mapping the content to the appropriate knowledge/skill statement. The mapping for such an item, using the key provided in Table 10, would be 2.[1,2].2.2.1., "Converts fractions to percents."

Note that content can be tagged at any of a number of increasingly specific levels within this system, yet can also be retrieved at a number of increasingly more general levels. Thus, for example, the user could search for a specific assessment item, request a more general search on content related to the item (by searching the benchmark mapping code), or request content related to a topic to view a number of benchmarks. So that the retrieval is accurate to the topic level, each knowledge/skill statement is assigned only those topics that are appropriate to it. The bracketed numbers above

identify which topics have been "hard-wired" to each knowledge/skill statement.

Table 10. A Sample of the Four-Tiered Structure Using Mathematics

Level	Name	I.D.	Content
I	Standard	1	Understands and applies basic and advanced properties of the concepts of numbers
II	Topics	1	Ratio/proportion/percent
		2	Fractions
		3	Decimals
III	Benchmark	2.2	Understands equivalent forms of basic percents, fractions, and decimals (e.g., 1/2 is equivalent to 50% is equivalent to .5) and when one form of a number might be more useful than another
	Knowledge/ Skill statements	1	Converts fractions to percents
		2	Converts percents to decimals
		3	Converts fractions to decimals
		4	Knows the equivalence of common fractions and decimals
		5	Knows the equivalence of common percents and decimals
		6	Knows the equivalence of common fractions and percents

4.1. Populating the Schema

The metadata schema to be useful, however, must be populated, or have associated with it an appropriate "value space," or scheme. Value spaces, an essential complement of the metadata schema, include such things as thesauri, ontologies, and classification systems [8]. Knowledge/skill statements and their associated topics, benchmarks,standards, and vocabulary items have been developed for selected subjects in the most recent edition of McREL's online *Compendium* [9]. Currently this extension comprises over 750 four-tiered structures of the type illustrated in Tables 9 and 10. These extensions provide content descriptions for the subject areas of English language arts, mathematics, and science. Table 11 describes the current status of McREL's database, identifying the number of items available at each of the four levels of specificity.

Table 11. Number of items in the McREL database by subject

Subject	Language Arts	Mathematics	Science
Standards	8	9	16
Topics	38	54	55
Benchmarks	273	226	256
Knowledge/ Skill Statements	2669	1110	1492

4.2. The Advantages

The metadata schema and associated value space described here provide a number of noteworthy advantages, which include:

- Content can be mapped at the level of specificity at which it is described, rather than "force fit" against a differing level of specificity. For example, broad state standards can be mapped at the level of generality expressed in their standards.
- Content can be mapped at a level of specificity dictated by intended use. For example, general lessons on fractions could be mapped to the topic level, while specific assessment items could be mapped to the knowledge/skill statement level
- If the chosen level of specificity for content retrieval is not satisfactory, content can be retrieved by either "rolling upwards" to greater levels of generality or "drilling downwards" to find finer degrees of detail.
- Computer search algorithms may be made more efficient once guided by the specificity of content descriptions as well as the established hierarchy of general to more specific content.
- The database is not exclusive to a particular curriculum or set of national standards.
- The database is fully available for inspection with appropriate citations to the numerous authoritative content documents from which it was developed.

5. Summary

Over the past decade, many largely unsuccessful attempts have been made to develop a mechanism whereby educators can efficiently retrieve resources that are specifically targeted to the standards they wish to teach. The lack of a universal set of specific content descriptors has contributed to much of the difficulty. A metadata schema with its associated value space, developed from an online database of standards at McREL, has been proposed as a means for improving the identification and retrieval of education resources.

Author's note

This publication is based on work sponsored in part by the Institute of Education Sciences. The content of this publication does not necessarily reflect the views of the Institute of Education Sciences or any other agency of the U.S. government.

References

[1] Gaddy, B. B., Dean, C. B., & Kendall, J. S. (2002). *Noteworthy perspectives: Keeping the focus on learning.* Aurora, CO: Mid-continent Research for Education and Learning. Retrieved May 5, 2003 from http://www.mcrel.org/PDF/Noteworthy/5022IR_NW_Focus.pdf

[2] Lee-Smeltzer K.-H. "Finding the needle: controlled vocabularies, resource discovery, and Dublin Core." *Library Collections, Acquisitions, and Technical Services*, Summer 2000, vol. 24, no. 2, pp. 205-215(11)

[3] Sutton, Stuart A., Elizabeth D. Liddy, and John Kendall. "StandardConnection: Correlating educational resources in digital libraries to content standards." *Proceedings of the Second ACM/IEEE-CS Joint Conference on Digital Libraries* ACM, 2002: 408.

[4] Learning Technology Standards Committee (IEEE). *Position statement: 1484.12.1 - 2002 Learning Object Metadata (LOM) Standard: Maintenance/Revision.*Retrieved July 1, 2003 from http://ltsc.ieee.org/wg12/index.html

[5] Gateway to Educational Materials. *Why should metadata be created for your educational resources?* Retrieved July 1, 2003 from http://www.geminfo.org/decision.html

[6] Department for Education and Skills & Qualifications and Curriculum Authority. *National Curriculum Metadata Standard: Overview*. Retrieved July 1, 2003 from http://www.nc.uk.net/metadata/index.html

[7] Kendall, J.S. & Marzano, R.J. (2000). *Content Knowledge: A Compendium of Standards and Benchmarks for K–12 Education.* (3rd ed.). Reston, VA.: Association for Supervision and Curriculum Development. Retrieved May 5, 2003 from http://www.mcrel.org/standards-benchmarks/

[8] Sutton, S. (in press). "Principled Design of Metadata Generation Tools for Educational Resources." In *Developing Digital Libraries for K-12 Education*. M. Mardis (Ed.), ERIC Clearinghouse on Information & Technology.

[9] Kendall, J.S. & Marzano, R.J. (2003).*Content Knowledge: An Online Compendium of Standards and Benchmarks for Preschool through 12th grade: Language Arts, Mathematics, & Science.* Mid-continent Research for Education and Learning. http://www.mcrel.org/compendium/kSkillsIntro.asp