Seven Fundamental Questions
for the
Science of Library Classification

Robert M. Losee
School of Information and Library Science
U. of North Carolina
Chapel Hill, NC 27599-3360
USA

losee@unc.edu

Abstract: For document ordering and classification to advance toward optimal manual and automated systems, it is necessary for a science of document or library classification to be developed. Seven questions are posed that the author feels must be addressed, if not answered, before optimal systems can be developed. Suggestions are made as to the form that answers to these questions might take.

As classification continues to move forward into a largely computer-based information age, it becomes increasingly possible to incorporate aspects of theoretically optimal classification procedures into libraries and other information systems. Classification for purposes here is considered the assignment of a value to an entity and the ordering and organization of these entities by these values.   Current classification systems are unlikely to be optimal; if they are optimal, proof of this would be of great benefit. Optimal systems must be developed consistent with methodologically rigorous or scientific principles, leading to a *science of classification.* This science can make explicit the fundamental parameters of library classification, as well as explain relationships between the components of an optimal classification system. This science may rapidly become more formal than current classification systems and, indeed, somewhat more quantitative. This will be largely due to the relative ease with which elegant ordering systems may be developed based on models from within the mathematical sciences.

An understanding of these facets of library classification may lead to an understanding of (1) the functional operation of present systems, (2) the effects of modification to these existing systems, and (3) how completely new systems might perform. Another role a science of classification may play is the answering of practitioners' questions. However, this is not the primary purpose of this or other sciences and practitioners should neither accept nor reject the basics of a science of

classification primarily because of the way it addresses the problems they currently see as being central to the discipline. For a science to progress, it must do so on its own terms; to focus primarily on the needs and concerns of practitioners is to blunt the progress of the discipline.

The science of library classification (Moravcsik, 1986) has been developing at a slow pace. A clear statement of what classification is and how it can be improved is of critical importance to the further development of classification systems in both traditional libraries and in paperless systems. A classification system developed from rigorous foundations has many advantages over more *ad hoc* classification systems. For example, one can understand why a theoretically based system performs as it does, if it is grounded on scientific principles, providing information professionals with the capability to *explain* what is occurring. The ability of scientific models to *predict* will allow library professionals, particularly technical services staff, to project into the future the collocation capabilities and retrieval performance of a classification system, given a collection organized consistent with scientifically based classification principles. The current state of classification, largely existing as an art and a set of philosophical constructs imposed on a knowledge base, does not allow this sort of prediction or explanation to be made.

Optimal enumerative classification systems can be developed that support patron browsing through user-oriented classification systems. While faceted and synthetic classification systems may eventually prove to be excellent bases for classification systems, enumerative systems may be

most easily described and optimized by scientifically based systems. These scientifically based systems may be based on objective criteria and optimized for particular groups of classification system users. Unlike other classification warrants (Beghtol, 1986), this *functional* warrant suggests that a particular classification system function might be optimized, in this case, for the end users' preferences. Variations of this functional warrant allow for the optimization of other combinations of classification system components, including both the users' and the librarians' goals and needs, unlike current systems which do not explicitly or implicitly attempt to optimize.

Reclassification can be accomplished through the development of a new document ordering based on information not available when the initial classification was made, such as changes in frequencies of specific terms in recently published documents or bibliographic records, or because of changes in user preferences. Procedures that could automatically and optimally reclassify an entire OPAC and produce new spine labels might be useful in reclassifying existing book collections, if the significant effort expended in reclassification was justified by the increased access to the collection. As the number of online full-text systems increases, the ability to reclassify electronically stored documents will increasingly require *no* physical changes in the system, such as relabelling books or editing catalog records. Reclassification will be performed easily by a computer program, based upon a rigorously defined set of criteria for a classification system.

There are seven basic questions that need to be answered or further explored if a science of classification is to develop so that patrons and librarians can rely on the classification system to perform in the most efficient way possible, with a degree of reliability not found in current systems. Designed to articulate with the methodologies of other scientific disciplines studying the organization and accessing of materials, these questions are both explicit and objective. The questions, directions that might be pursued in seeking their solutions, and examples of answers and the form that answers might take, as well as some of their variants and related problems, are as follows:

1. What should be the form of subject-indicating representations?

What best represents a document's subject if the representation is to be incorporated into a classification number assigned to a document? These representations take the form of combinations of letters, digits, and punctuation marks in Library of Congress (LC), Dewey, and many other classification systems. If library classifiers are intended to group documents by subject for browsing and a classification number is to represent directly or indirectly the subject of the document, then understanding how a subject-indicating variable is used in a call number is essential to understanding a classification system. Let us accept that bibliographic materials may not be completely describable by a simple single statement; the subject of an intellectual product has many facets or

characteristics and thus a representation of what it is about must include the values of several subject-indicating variables. Note that determining the subject of a document is not strictly part of the science of classification, although determining the subject of a document is a part of the day-to-day work of a practicing classifier (Foskett, 1982; Kwasnik, 1991; Lancaster, 1986). It is studying the *representations* of these subjects and their subsequent *ordering* and *organization* that is the domain of the science of classification.

The subject indicating variables can hold fixed length representations, such as a fixed number of letters, integers (whole numbers), or binary numbers, or be variable length or infinite representations, such as real numbers (numbers with decimal points) which, for example, could represent 1/3 using an infinite string of digits, 0.333333.... Because a real number can be infinite in length, its use as a representation to be printed on a spine label or stored in a computer record is excluded. A limited range of integers, for example, 0 to 9, could represent a set of 10 possible levels of subject-aboutness in a finite sized representation. A binary representation simply represents whether a document is about a subject or not. Representations can always be finite in size given a finite set of document-subjects.

Systems such as the Library of Congress and the Dewey Decimal Classification Systems (and their variants) implicitly limit the set of features that a document might exhibit, and thus may be represented using a finite representation system. However, in many cases, they do not

comfortably represent documents being equally about two different subjects or having two equally important attributes; their inherent hierarchical nature arbitrarily places one subject over another in terms of importance. By choosing one initial letter for the LCC for a document, one is excluding a set of other possibilities. While this may be seen as a strength, it may be a weakness when documents are interdisciplinary in character.

Allowing a document classification to have any possible combination of feature values can result in very long classification "numbers." These numbers may be decreased in size through the use of data compression (Storer, 1988), which reduces the size of a representation without information loss. Given data compression techniques, the representations provided for a given document by any of a number of different possible classification systems may all be reduced to the same average size, assuming that all the classification systems represent the same universe of possible document-subjects.

1a. Which subjects should be used in a classification and representation scheme?

If inclusion of all potential subjects is a requirement of the system, the question becomes "Given a very large set of possible subject-indicating variables, which subset or alternative set of variables can most efficiently represent the subject content of all documents?" More geometrically, one might ask which *feature space* is best. One answer to this question is to

choose a set of features that are *statistically independent* and thus provide no information about each other, minimizing redundancy. For example, features such as *cataloging* and *classification* are related and not statistically independent in most library catalogs. To choose a variable that is in part dependent on another variable would be to select variables that do not carry as much information as is possible, thus forcing the system to provide larger representations than necessary.

Independent variables may be computed through procedures such as factor or principal components analysis (Borko, 1985; Deerwester et al., 1990). These procedures begin with data, such as natural language text, and compute a set of independent features, features that do not correlate with each other or provide information about each other. These features have the least amount of redundancy possible.

A conceptually different procedure, dimension reduction, explicitly reduces the number of features in a data set. Two examples of dimension reduction are provided by the field of cartography. Terrain maps are capable of showing three-dimensional terrain in two dimensions. Similarly, maps using the Mercator projection show a three dimensional, round planet on a two dimensional surface. Both of these methods distort and leave something to be desired, but they do diminish the number of dimensions or features represented by the descriptive mechanism, which may be valued in a representation scheme. Dimension reduction in classification involves removing factors of lesser importance and leaving those of greater significance. More specifically, one should continue to

reduce dimensions as long as little or no information is lost. This is done by almost all human-based systems, which find it easy to distinguish the wheat from the chaff, but it is far less easily done with automated systems.

Classification scholars have suggested varying numbers of dimensions as being best for representation and ordering. For example, empirical support exists for Ranganathan's PMEST model which uses 5 basic categories (personality, matter, energy, space, time) (Hemalata, 1992). Whether these models are optimal, close to optimal, or adequate, are empirical questions whose answers requires knowledge of the optimal classification system, given a set of desired constraints.

1b. What representational factors may be easily manipulated and remembered by human beings?

When selecting a feature set upon which to base a classification system, the classification scientist may wish to choose those features that are easiest for humans to remember and manipulate, all other factors being equal. Relatively little is known about the comparative cognitive load placed on a user by different representational systems, although numerous studies are described in the cognitive literature about memory and the ability to remember different types of representations (Dahlberg, 1992; Garcia-Marco, 1993).

The ease of use of a classification system is a factor only when the classification system is a *visible classification system,* a system that needs

to be examined and used directly by the patron. *Invisible classification systems* place similar documents near each other without displaying the classification number. These classification numbers are for the internal use of the computer or information system only. Invisible classification systems are likely to become increasingly important as computer-based systems provide retrieval through one mechanism, e.g., Boolean searching, and document organization to support browsing a set of documents through another mechanism. One might use Boolean queries to locate an initial interesting document in a database; neighboring similar documents might then be examined. This notion of classification is more limited than existing library and book-based visible classification systems, which function both as browsing and finding tools, with users conducting known item searches, moving from the catalog to the book stacks with transcribed call numbers.

2. How should distance and dissimilarity values between individual document features be combined?

When making decisions about document classification, it is necessary to consider factors such as how alike two books are and how close or far apart they should be placed in a collection. Yet, given a universe of different subject-variables, it is not obvious how one should combine the distances (dissimilarity values) that might be computed for each feature or variable to provide a composite document-similarity measure.

Two approaches to measuring similarities between multiple featured variables have been used widely in information theory (Losee, 1990). The first is to merely count the number of features by which two representations differ. This is often referred to as the Hamming distance, as well as being called the Manhattan or city block distance. The latter names are derived from the idea that if each possible representation is taken as a street-corner in a city, the distance between representations may be measured as the distance moving from one intersection to another by using streets and not cutting across lots, and where it is understood that each block has a length of 1. The other method of measuring distance, based on familiar Euclidean geometric considerations, computes distance "as the crow flies," that is, where one is not limited to traveling on streets. This is similar to incorporating statistical dependence information between features. These methods assume that one can combine distances associated with different variables.

3. What average between-document or shelf distance should be minimized?

The distance between documents is a critical factor in a classification system. If documents are to be browsed, it is necessary for similar documents to be a short distance from each other. It also may be desirable to explicitly place less similar documents at a greater distance. Thus, the study of document distance is a critical part of placing similar

documents "close" together. Attention to distance may take many forms, the two extremes of which are examined here.

Ignoring the choice of taking no account of distance in a theory, the least attention one might pay to distance is to only examine adjacent documents, that is, only consider distances of 1. In essence, a collection may be classified by randomly selecting a document, placing next to it the document with the greatest degree of similarity, placing next to this document the document with the greatest degree of similarity to it, and so forth. This method is probably relatively stable, given the addition or deletion of documents, especially in large collections.

At the opposite extreme is taking a global view, taking into account similarities and distances between all possible pairs in a collection, adjacent or not. This method is somewhat sensitive to small changes; adding a single document can cause a ripple effect, modifying the order throughout the entire collection.

3a. Should distance be treated as a linear quantity?

How individuals react to distances in the library is an empirical question that remains largely unexplored. Patrons appear to browse through a small area on a classified shelf, examining books located within a few centimeters to a few meters apart, at most. It is not obvious that a book located 200 meters from the start of browsing is twice as costly to retrieve as one located 100 meters from the starting point, that is, that the

distance or effort function should be treated as linear. This is in part due to retrieval procedures when finding documents at some distance from each other. These differ from procedures used when browsing through sets of very close documents. Documents of interest on the same shelf are likely to be located through browsing, while documents several ranges away will need to be located through other procedures, such as through a known-item searched using information provided by a catalog record.

4. *What is the cost of placing similar or dissimilar documents a certain distance apart?*

Several factors potentially useful in developing classification systems are objectively determinable. The distance between documents is one such objective factor. The informational similarity between documents is likewise objective if information is objectively measured as is done by Shannon and others (Boyce et al., 1990, Losee, 1990). Given these objective quantities, how might a classification system be "customized" for a particular collection, assuming that classification systems should not be the same for all collections? A collection is most easily customized by noting that the difference between library collections indicates an expected difference in the information needs of the patrons; these different needs place different economic constraints on the classification system, suggesting that developing an economics of classification may result in improved adaptive classification systems.

The value of grouping certain documents together or placing two documents next to each other may be empirically determined and will be different in different informational environments. These costs may be further broken down into the costs associated with ordering (in some sense) documents based on a single feature. Once single feature costs have been determined, they may be combined, allowing for the estimation of the costs of placing documents at a certain distance with a given degree of dissimilarity. In addition, there are cognitive factors associated with documents of different types being adjacent, or being placed at a particular distance.

5. In what order should documents be placed?

An ordering applied to a set of documents may be understood as a path through a space consisting of all possible documents or all possible document-subject combinations. If a library's books were randomly placed, a classification system could be viewed as a path from each book to a randomly selected "next" book, providing an order in which they could be placed on a shelf. This path *is* the ordering principle within the classification system, and understanding the nature of the classification path is fundamental to understanding the nature of the classification system itself.

To examine paths through a collection, let us assume that each of n possible independent features is represented or is not represented in a

document, with each value denoted by a 1 or a 0, respectively. The set of possible documents' subjects can be represented as the set of corners or vertices of an *n*-dimensional cube (Cosgrove, 1991; Losee, 1992). A classification system might use the binary Gray code (Gilbert, 1958) to supply an ordering of vertices and a means of representing what the document is about.

A simple binary coded classification system with only one feature in the system would be represented by a line with two ends, one end representing a 1 or the presence of the feature and the other end representing a 0 or the lack of the feature. Similarly, a system with two features would be represented by a square, with each of the four possible values (00, 01, 10, 11) for the two features represented by one of the four corners of the square. A classification system's ordering in this instance might be a path around the square.

A set of *n* features is said to produce an "*n*-cube." A classification system may thus also be seen as a *path connecting all vertices on the n-cube.* Each corner represents a particular set of coordinates. Thus, each path through this cube represents a counting sequence for the set of all possible coordinates. Classification may thus be viewed both as a path and as a counting sequence. This counting sequence provides the familiar filing order for books.

In many circumstances, multiple paths may satisfy the requirements that constrain a classification system. Consider, for example, a traditional table with four legs and thus with six planes

composed of the vertices at the ends of the table legs (on the surface of the 3-dimensional cube): the top, bottom, front, back, left and right planes (diagonal planes are ignored here). We assume the distance minimizing criteria proposed earlier, in which the primary concern is to place each document next to another very similar document. This criteria may be met by starting at one corner of the top of the table and moving through all vertices on the table top. When this circuit has been completed, move straight to the floor and then make a circuit of the vertices on the bottom, where the table legs meet the floor. The criteria might also be met by moving around the edges of the front of the table, then, when this is completed, moving to the back and then making a circuit of the back plane of the table.

These, as well as other possible circuits, all meet the criteria established earlier. One of these systems may be arbitrarily accepted, if all perform equally, or other criteria may be introduced to further limit the number of acceptable paths or classification systems.

6. Should paths and classification numbers stay constant for all collection variations, e.g., different collection sizes and concentrations, or should they be adaptive?

Given that classification may be viewed as a path along the edges of the *n*-cube, should the varying degrees of density or unequal densities modify the best path? In other words, should different kinds of libraries or different sized libraries use different classification systems? If the

classification system (1) is invisible or (2) the patron only uses this one collection or (3) the library or system manager follows the lead of retail stores who organize their goods in the manner they feel is best for their specific needs, then it is best to customize the document ordering for the particular group of clients.

6a. Should the path stay constant for all different collection emphases?

Consider a classification system for four possible documents. The documents are conceptually arranged clockwise around a square of the four vertices in the order A, B, C, D, with A in the upper left. Using the Hamming distance to measure both subject dissimilarity and the physical distance apart, the physical distance and subject-dissimilarities between adjacent documents are 1 and the distance and dissimilarities between diagonally opposed documents are 2. If all four documents are present, then the orders {A, B, C, D}, {B, C, D, A}, {C, D, A, B}, and {D, A, B, C} are (equally) optimal under many classification schemes and each ordering would be rationally acceptable. However, if the possibility is decreased or removed of having a document B, because of a shift in collection emphasis, it becomes clear that an ordering including the sequence {A, C, D} is far from optimal, because A is adjacent to a subject-opposite document, C. A better arrangement with B removed is clearly {A, D, C} or {C, D, A}. Therefore, the classification path should not stay constant for all different collection emphases.

6b. Should the path stay constant for all different collection sizes?

The path should stay constant as a collection grows in size only if the growth is uniform, that is, the probability that a document will have a given feature remains constant as the collection grows. If there is a shift as the collection grows, as there almost always is in practice, then the argument presented in the response to Question 6a above would hold, making it necessary to reclassify as a collection grows, if optimal classification is to be achieved.

6c. In an adaptive classification system, what should be optimized, and how?

Classification systems used in most libraries are static, having been developed for a particular situation. Classification systems may be dynamic and may adapt to the needs of individuals (Losee, 1997). Adaptation may be based on a number of criteria, such as geometric considerations of the closeness of documents, or economic considerations of minimizing costs to one or several users. While these systems are currently rare, they will rapidly increase in frequency as more paperless systems attempt to provide better service by organizing documents in such a way that the document organization is optimized for each individual.

7. How should classification performance be measured?

The development of a classifier performance measure is essential to the theory and development of a science of library classification. A measure should be developed based on whatever measurement characteristics the user desires to study with a measure. This idea of a measure treats a measure as being related to a model; a measure is developed so that it accurately analyzes the characteristics that the model and its assumptions dictate are important. One might arbitrarily decide, for example, that the measure should have a value of 1 if perfect classification occurs and a value of 0 if random classification occurs. This scale might be easily changed so that it would provide a different range of values, for example, 0 to 100, or change the meaning of the top and bottom of the range, changing the bottom value of 0 to occur only with the worst possible classification performance.

One possible classification measure is the *expected browsing distance*, the average physical (and possibly cognitive) distance that the user has to move, whether on computer screens or through a library's stacks (Losee, 1997). By directly addressing the measurable cost to the user and providing a measure whose value can be predicted analytically (Losee, 1998), we have a general measure of browsing performance whose characteristics can be understood and manipulated.

7a. What is the best possible classification performance?

An interesting question is whether the "best" system should be a very good practical system or one that might never be fully achieved but is perceived to be the best imaginable. For example, an ideal but never to be achieved system would place adjacent to each other all the documents each user might find of interest. For any realistic set of users, this is almost unachievable. Should this be used as the "best" performance obtainable?

7b. What is the worst possible classification performance?

A measure of classification performance may require knowledge of the performance of the worst possible classification system. Accepting the n-cube model described above, the worst possible classification system is one where paths constantly move almost directly across the n-cube from one corner to an almost-opposite corner. Always moving to the opposite corner is impossible, as one would constantly move back and forth between the same pair of corners. This worst-possible system may be implemented in a classification system through use of the anti-Gray code proposed by Hamming (1986).

Conclusions

Library classification, the ordering of materials to facilitate patron browsing, can be studied objectively. Classification systems can be optimized, providing maximal browsing support for the user. Developing these "best" systems to support both traditional book-based libraries and paperless, full-text electronic systems requires answers to the seven fundamental questions examined here. In addition, information professionals should strive for the development of library systems consistent with models that allow us to *explain* what happens and to *predict* future performance.

These questions have been proposed in such a way that answers may be easily found, given research in this field or borrowing results from other scientific disciplines faced with similar problems (librarians are not alone!). In particular, the assumption that subject-bearing features are binary has allowed partial answers to some of these questions to be framed in terms of the n-cube, an easily analyzed theoretical construct that generalizes from the familiar three-dimensional cube. The author believes that studies of this model will allow for the further development and improvement of objective classification procedures.

Empirical research is needed to examine further human behavior and preferences when using classified materials. For example, questions about user reactions to distances as well as uncertainty about costs of placing documents at a certain distance can only be studied through

empirical research. Given both this further empirical work and theoretical analyses of linear classification systems, the science of library classification can be expected to move forward at an increased pace.

References

Beghtol, C.: (1986). Semantic Validity: Concepts of Warrant in Bibliographic Classification Systems. *Library Resources and Technical Services* 30 (2), 109-125.

Borko, H. (1985). Research in Computer Based Classification Systems. In L. M. Chan, P. A. Richmond, and E. Svenonius (Eds.). *Theory of Subject Analysis: A Sourcebook*. Littleton, Col.: Libraries Unlimited, 287-305.

Boyce, B., Douglass, J. S., Rabalais, J., Shiflett, L., and Wallace, D. P. (1990). Measurement of Subject Scatter in the Superintendent of Documents Classification. *Government Publications Review* 17, 333-339.

Cosgrove, S. J. (1991). Item Classification Using the N-Cube's Hierarchical Knowledge Representation Schema. In C. McDonald and J. Weckert. *Libraries and Expert Systems*, Los Angeles: Taylor Graham, 88-98.

Dahlberg, I. (1992). Cognitive Paradigms in Knowledge Organization. *Knowledge Organization*, 19(3), 125, 145.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by Latent Semantic Analysis. *Journal of the American Society for Information Science* 41 (6), 391-407.

Foskett, A. C. (1982). *The Subject Approach to Information.* 4th ed. Hamden, CT: Linnet Books.

Garcia-Marco, F.J.. & Estaban-Navarro, M.A.. (1993). On Some Contributions of the Cognitive Sciences and Epistemology to a Theory of Classification. *Knowledge Organization*, 20(3), 126-132.

Gilbert, E. N.(1985). Gray Codes and Paths on the N-cube. *Bell System Technical Journal,* 37, 815-826.

Hamming, R. (1986). *Coding and Information Theory*, Second Edition. Englewood Cliffs, N.J.: Prentice-Hall.

Hemalata, I. (1992). Subject Representation and Entropy. *International Classification.* 19 (1), 15-18.

Kwasnik, B. (1991). The Importance of Factors that are Not Document Attributes in the Organization of Personal Documents. *Journal of Documentation,* 47 (4), 389-398.

Lancaster, F. W. (1986). *Vocabulary Control for Information Retrieval.*
Arlington, VA: Information Resources Press.

Losee, R. M. (1997). Browsing Document Collections: Automatically
Organizing Digital Libraries and Hypermedia using the Gray Code
Information Processing & Management, 33 (2), 175-192.

Losee, R. M. (1992). A Gray Code Based Ordering for Documents on
Shelves: Classification for Browsing and Retrieval. *Journal of the American
Society for Information Science* 43 (4), 312-322.

Losee, R. M. (1990). *The Science of Information.* NY: Academic Press.

Losee, R. M. (1998). *Text Retrieval and Filtering: Analytic Models of
Performance.* Boston: Kluwer.

Moravcsik, M.J. (1986). The Classification of Science and the Science of
Classification. *Scientometrics* 10 (3-4), 179-197.

Storer, J. A. (1988). *Data Compression: Methods and Theory.* Rockville,
MD: Computer Science Press.