EXTENDING EXPERT SYSTEM TECHNIQUES WITH INFORMATION RETRIEVAL METHODS TO SELECT FOSTER HOMES

Edward A. Fox (fox@vtopus.cs.vt.edu)
Virginia Tech, Blacksburg VA 24061

Sheila G. Winett
Roanoke College, Salem VA 24153

Abstract

Information retrieval concepts and methods have been incorporated during the design, development, and testing of FOCES, a FOster Care Expert System built upon the GUESS expert system shell, which can select placements for children almost as well as trained social workers. The difficult task of knowledge acquisition and representation for this human services application has been greatly simplified through the use of p-norm queries; the problem of selecting suitable homes has been addressed through the use of matching methods; and the evaluation of system performance has been gauged in part through the use of retrieval effectiveness measures. Thus, it is possible to develop expert systems for classes of applications prevalent in government social agencies, and information retrieval techniques can greatly simplify such efforts.

1. Introduction

Approximately one half million children are now in foster care in the United States of America [29]. Though the long term goal of this human service is to identify permanent homes for these children, many must remain in foster care for a number of years, and may be forced to "bounce" between homes during that period [18]. In developing an expert system for foster care it was therefore decided that initial placement in the proper homes was an important short term objective [33]. However, as in many social science domains, it is extremely difficult to identify situations where there is a good person-environment fit [21].

After many interviews with practicing social workers and investigation of available literature, it became clear that conventional expert system development methods of knowledge acquisition and knowledge representation [11] were inadequate for the task at hand. It was hypothesized that representation and processing methods used in experimental information retrieval systems would simplify the overall effort [34], and that proving the utility of this approach would be helpful to other practitioners. It might ultimately then become easier to develop expert systems for the social services and other difficult domains [27].

Section 2 describes the foster care task domain and the challenges involved in applying expert system techniques to the social services. Section 3 gives an overview of the knowledge representation approach, which employs forms supported by the GUESS expert system shell [15] as well as special Prolog structures for "p-norm" queries [25]. Section 4 describes the data transformations and matching operations which allow the home selection process to be viewed as an information retrieval problem [34]. Section 5 discusses the results of a small-scale evaluation of the system's effectiveness. Section 6 summarizes this project, and suggests future extensions.

2. Problem Domain

Foster care services are provided by designated social services agencies to a child and his/her family to prevent removal of the child from the home. However, when parents are unwilling or unable to use these preventive services to change conditions in the family home, placement of the child in a foster home may be sought by the local social services agency empowered to address child welfare problems in the community in which the family resides [32]. When such placement is necessary, the social worker who has been assigned to investigate the situation — the individual directly responsible for protecting the child's health and welfare — selects the home. That choice is made in consultation with other social work professionals and appropriate supervisory personnel. The decision is constrained by state and federal laws and administrative regulations, heuristics of "good social work practice" and local exigencies.

Despite the pervasiveness of foster care, some key domain issues remain unresolved. As a result of these unanswered questions the developers of FOCES were required to rely on unconventional means to accomplish the knowledge acquisition tasks traditionally considered requisite for satisfactory expert system development, such as defining a benchmark for measuring system success. While permanency for a child may be the ultimate goal of foster care programs, its achievement or non-achievement is, in most cases, wholly independent of what a child may experience while in placement. Therefore, the primary developer of FOCES had to construct her own definition of a successful placement: one that continues without any unplanned interruption caused by the foster parent requesting the child's removal.

To achieve its stated goal, FOCES had to be capable of discovering those patterns of child-foster home interactions that consistently result in a stable placement, i.e., lead to a good person-environment fit [6]. Unfortunately, there are serious methodological and philosophical problems associated with defining what constitutes a good "person-environment fit" and what factors influence the achievement of such a fit. Definitive rules have yet to emerge, and defining person-environment fit remains an open social sciences research question [21] [28]. This lack of consistent domain knowledge greatly influenced the match strategy adopted by FOCES.

The specific problem addressed by FOCES was to provide assistance to social workers in the Roanoke City Department of Social Services. Consequently, the design of FOCES was also influenced by the nature of the information available from that agency. Long, free-form narrative records, complemented by a limited amount of information

coded on forms for use in Virginia's automated information system, are maintained by the agency about its foster home resources and clients. To systematize the gathering of data from these records, two data collection forms were developed, one about the child and one about the foster home. The child form focused on current circumstances, and only included historical data about the child's prior placements. The child's special needs and notable behaviors were identified using codes taken from an agency form and verbatim relevant case record comments. The data collection form for the foster home included information on family demographics and composition, physical capacities of the home, caregiver preferences/prohibitions, past experience/training, and strengths and weaknesses of the home.

3. Representation

How knowledge about children and foster homes should be represented has been a critical concern for the designers of FOCES. Several factors, including the descriptive, free-form nature of the "real world" data, the plan to use the General pUrpose Expert System Shell (hereafter referred to as GUESS) as a shell for building the system, and the decision to adapt information retrieval techniques as a guide for classifying children and homes and finding appropriate matches between the two enitities, influenced the selection of knowledge representations. To satisfy the requirements imposed by each of these factors, multiple knowledge representations were used by FOCES as can be seen in Figure 1.

Figure 1. Knowledge Representations

	Data Collection Forms	GUESS Data Repre-sentations	Vector Repre-sentations	Classifi-cation Lists
CHILD$_i$	F_i^C	D_i^C	V_i^C	L_i^C
HOME$_j$	F_j^H	D_j^H	V_j^H	L_j^H

In all phases of processing the knowledge representations for children were similar to those for homes. Initially, some important real-world data about each such entity was recorded from available sources on a specially designed data collection form. Next, using the human-computer interaction features of the GUESS expert system shell, relevant data were entered and stored in the form of Prolog facts. Tables of these facts were then transformed to allow further processing. In particular, a single vector was constructed for each real-world entity. Using p-norm queries, each child vector was classified according to a set of prototypes. Finally, each home was classified as to how well it could accommodate each child prototype.

3.1 Data Collection and Forms

The first knowledge representations used by FOCES were two data collection forms designed to capture information from child and foster home case records. These were developed after the primary developer had spent approximately six months using "traditional" knowledge acquisition mechanisms, such as interviews with experts and searches in the domain literature, and still did not have sufficient information to proceed with system development.

Social workers are often required to deal almost exclusively with the "squeakiest wheel," i.e., the most severe emergencies. The necessity to work in "crisis mode" diminishes one's capacity to step back and analyze daily activities. As a result, one's capacity to serve as an "expert" is considerably reduced. An additional complication in the foster care domain is the long-term, nationwide shortage of foster care homes [13]. As a result, social workers rarely have the luxury of selecting from among many available homes.

The answers given to interview questions aimed at eliciting expert knowledge did not provide a sufficient base for building rules for matching. Nor did it seem that many more rules could be obtained, even with more elaborate collection methods [10]. As a result, the primary developer had to conduct an extensive investigation process to find the specifics needed to transform the general attributes identified by the social workers, such as "child behaviors," "special needs of children," and the like, into suitable matching factors. For example, the records had to be scanned to identify which child behaviors were "noteworthy," which preferences of foster home caregivers were routinely cited, and the like.

Approximately 50 child case records and a like number of home case records were carefully read, and detailed notes were taken on each of the attributes represented on the data collection form. The results of this effort provided the basis for formulating the criteria that were subsequently used to select suitable homes for a child. Looking at the "child behaviors" attribute as an example, the result of this data collection process was to identify eleven notable behaviors that seemed to be characteristic of different types of children. The presence of these behaviors in a child's case record was one factor that subsequently served as a basis for distinguishing whether a child better fit into one prototype or another for matching purposes.

3.2 GUESS Tables

Following the data collection process, information encoded on the data collection forms had to be represented in GUESS tables, as FOCES is implemented using GUESS and HC Prolog. The use of Prolog for constructing expert systems has been suggested [7], in spite of a general lack of support capabilities [3], and its application at Virginia Tech is especially appropriate since the HC (Horn Clause) version has been in use for several years. GUESS, a shell for developing expert systems, was also developed at Virginia Tech and has been demonstrated on several problem domains [15]. A later version of GUESS has since extended that effort, using the Edinburgh Prolog syntax [16].

Each GUESS table represents a single attribute of a child or home. For example, there is a table to represent the attribute identified on the child data collection form as birthdate of a child and a second to represent the attribute known as reason for placement. The former contains the unique identifier given to each child in the FOCES database and that child's birthdate. The latter also has each child's unique identifier. It lists the reasons that placement in a foster care home is being sought. As can be seen in Table 1 below, there is also a table to represent the attribute identified as race of caregiver on the home data collection form and another to represent the age range of children that the caregiver would prefer to have placed in her home.

Table 1. Partial List of GUESS Tables

Names of Child Tables

race of children
birthdates of children
primary reason for placement
gender of child
special needs of children

Names of Home Tables

race of caregiver
family members in foster home
gender of children approved for
age of children approved for
special needs of children approved for
number of children preferred
gender of children preferred
age of children preferred
special needs of children preferred

In all there are nine GUESS tables to represent child attributes and seventeen GUESS tables to represent home attributes. For a single child, therefore, the information maintained about him/her is distributed across a series of tables, each representing a single child attribute. The same is true of information maintained about each foster home. Although the developers of FOCES found GUESS tables sufficiently flexible to represent all of the data collected, GUESS does support other representation schemes such as trees and networks. In order to use the information retrieval techniques being tested, however, it was necessary to adapt additional knowledge representations to further augment the capabilities of GUESS.

3.3 Vectors

In many experimental information retrieval systems, vectors are used to describe documents [23] [31]. Vector representations of information about children and homes were constructed using the data stored in the GUESS tables. There is an almost 1:1 mapping between the information stored in the GUESS tables and the vector representations that were used. Whereas each table was about a particular attribute of children or homes, and had an entry for each such individual, vectors were used to describe all attributes of a particular child or home. Another key difference between the two schemes lies in the assignment of a real valued membership function to every attribute/value pair in the vector representation. The membership function is in the range of 0.0 – 1.0 and signifies the degree to which an individual child or home possesses a particular attribute. This adaptation of fuzzy set theory methods [14] to vector processing is particularly valuable. Finally, then, a vector can be viewed as a collection of tuples of the form:

(<child/home-id>, <attribute>, <value>,
<membership-fn-value>)

Tables 2 and 3 illustrate the tabular and vector representation schemes and make clear the relationship between the two. Table 2 is about children, and shows how the birthdate attribute is recorded in HC Prolog lists. It should be noted that in the vector representation "cfact" stands for "child fact", and ages are given in months. Further, membership function values are shown using integers in the range 0-10, and are later normalized to the unit interval.

Table 2. Representations of Child Birthdates

A. GUESS Representation

TABLE "birthdates of children"

"000001"	"09/00/77"
"000002"	"09/00/73"
"000003"	"09/00/71"
...	

B. Tuple Form of Vector Representation

cfact "000001" age 107 10
cfact "000002" age 155 10
cfact "000003" age 179 10
...

Table 3 is about homes, and gives related information about ages, measured in years. Specifically, the range of ages that a home has been approved for is shown. In the GUESS table, the range is shown as a list. In the vector form, there are different attributes indicating the youngest and the oldest ages that a home can accommodate.

Table 3. Representations of Approved Age Ranges for Homes

A. GUESS Representation

TABLE "age of children approved for"
"086631" (00 18)
"088382" (01 14)
"088315" (00 18)
...

B. Tuple Form of Vector Representation

hfact "086631" youngest 0 10
hfact "086631" oldest 18 10
hfact "088382" youngest 1 10
hfact "088382" oldest 14 10
hfact "088315" youngest 0 10
hfact "088315" oldest 18 10
...

3.4 P-Norm Queries about Prototypes

Since it is ultimately necessary to identify good homes for children, there must be some way to relate the two types of entities. The use of stereotypes [22] and prototypes [1] is a well-established practice in artificial intelligence, and relates easily to the classification and indexing operations required in many information retrieval systems. A set of 16 prototypes was therefore identified, each to describe a particular constellation of attributes thought to characterize a type of child or home. Each child could then be described in terms of how well he/she fits each prototype, and each home could be described in terms of how well each of the prototypes could be accommodated.

There are prototypes based on the age of a child, for example. Others are based on a child's race or on the primary need patterns associated with different types of children, such as medically needy children, dependent/neglected children, and the like. The complete list is given below in Table 4.

Table 4. List of Prototypes

1. infant (0-1 yr)
2. toddler (1-3 yrs)
3. preschooler (3-5 yrs)
4. schoolaged child (5 -13 yrs)
5. teenager (>= 13 yrs)
6. female
7. male
8. white child
9. black child
10. child of "other" race
11. dependent/neglected child
12. sexually or physically abused child
13. medically needy child
14. emotionally impaired child
15. intellectually impaired child
16. child with behavior problems

During the data classification phase of processing, the vector representations of children and homes are used to determine how similar each child's and home's attributes are to those identified as belonging to each child or home prototype class. Thus, the complete vector description of a particular child can be used to determine how well each prototype characterizes that child. This data reduction or classification process is similar to that carried out in SDI (selective dissemination of information) systems when a user profile, in the form of a query, is related to a newly received document. Consequently, it was decided that the most powerful query language available for information retrieval systems should be used to describe each prototype.

Boolean forms have already been shown to aid in expert system development [12]. Extending Boolean queries to allow relative weights on each clause component has been studied by a number of information retrieval researchers [4] [5] [20]. Additional work has shown that "softening" the interpretation of the Boolean operators themselves [19] can give further expressive power and lead to more effective retrieval [8] [25] [26]. These later studies found the use of "p-norm" extended Boolean queries to be particularly powerful, so Prolog routines for handling them were developed at Virginia Tech for an expert information retrieval system [9].

Table 5 shows, in HC Prolog notation, p-norm queries which can be applied to home and child vectors, respectively, in order to determine how well prototype 11 characterizes those homes or children. These p-norm queries, therefore, replace a large number of production rules, and provide a declarative representation of the meaning of each prototype. They are, then, a novel knowledge representation form used for the first time in this manner in FOCES.

The general form of p-norm query used here is:

(<relative-weight> <and/or-operator>
<operator-softness/strictness> <clause>... <clause>)

where the softness-strictness values are in the range of 1 (softest, corresponding to taking a weighted average) to infinity (strictest, corresponding to fuzzy set definitions of **min** and **max**). Each clause can be another p-norm query, or a ground clause. Ground clauses are of the form

(<attribute> <comparator> <attribute-value>
<relative-weight>)

and so evaluate to a (<truth-value>,<weight>) pair that can be handled by the p-norm interpretation routines. For a particular child or home, the entire query is then evaluated by recursive application of p-norm formulas (see [25]), from the inside out, to finally yield a single membership-function value that describes how well child or home is characterized by the corresponding prototype.

Table 5. Examples of Home and Child P-Norm Queries for Prototype 11

```
((homeproto 11
; PART A: home —
;    PLACE A DEPENDENT/NEGLECTED CHILD
    (1 and 1
        (4 and 2
            (approvedneeds = 25  5) (approvedneeds = 26  5)
            (approvedneeds = 29  5) (approvedneeds = 98  5)
            ) ; end approved need clause
        (2 and 2
            (strength = 01  4) (strength = 02  4)
            (strength = 03  5) (strength = 04  3)
            (strength = 05  3) (strength = 06  4)
            (strength = 07  2) (strength = 08  6)
            (strength = 09  5) (strength = 10  4)
            (strength = 11  3)
            ); end  strength clause
        (1 and 2
            (preferredneeds = 25  5) (preferredneeds = 26  5)
            (preferredneeds = 29  5) (preferredneeds = 98  5)
            ) ; end preferredneed clause
        )
    )); end homeproto 11
((childproto 11
; PART B: child —
;    A DEPENDENT/NEGLECTED CHILD
    (1 and 1
        (4 or 2
            (reasonforplacement = 01  5)
            (reasonforplacement = 04  5)
            (reasonforplacement = 05  5)
            (reasonforplacement = 06  5)
            (reasonforplacement = 07  5)
            (reasonforplacement = 99  5)
            ) ; end reason for placement clause
        (2 and 2
            (specialneeds = 25  5) (specialneeds = 26  5)
            (specialneeds = 29  5) (specialneeds = 98  5)
            ) ; child special needs clause
        (1 and 2
            (childbehavior = 03  2) (childbehavior = 04  3)
            (childbehavior = 05  2) (childbehavior = 06  2)
            (childbehavior = 09  2) (childbehavior = 10  2)
            (childbehavior = 14  2) (childbehavior = 98  3)
            ) ; end child behavior clause
        )
    )) ; end child prototype 11
```

P-norm queries are extremely flexible, although becoming facile in writing them requires a bit of practice. In the retrieval environment it is sometimes possible to automatically construct p-norm queries [24]. Thus, if further need arises to develop queries for new prototypes, it may be feasible to develop semi-automatic procedures to assist in this process.

After all 16 p-norm home queries have been applied to the vector for a given home, the home is then described by 16 fuzzy truth values, each in the range 0.0 – 1.0, which characterize how well each prototype can be accommodated in

the home. This home classification list is paralleled by a similar list of 16 values for any given child. Classification lists are then the final description of children and homes used in FOCES, and allow direct comparison between those children and home representations.

4. Computation

The operation of FOCES has been described indirectly in the last section in terms of how knowledge is represented. To view this more directly it is appropriate to view the flow of data as is shown in Figure 2.

Figure 2. Data Transformations

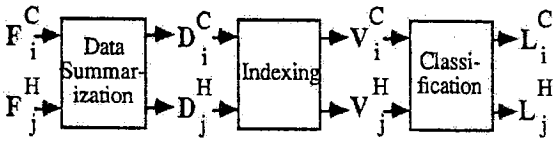

Initially, a social worker should collect data on either a child or home form, represented by the **F** entries above. The data is then interactively entered into FOCES, and stored in tables, represented by the **D** entries above. These steps can be viewed as background activities, to provide FOCES with the necessary information to carry out its more specialized selection/advising task.

In its current mode, FOCES will analyze a batch of children and homes, to carry out experimental comparisons. First, the GUESS tables are all transformed into corresponding vectors, thus carrying out the indexing mapping from **D** to **V** forms that is shown above. Then, p-norm queries about children (Q^C) are used to produce child classification lists (L^C) from child vectors (V^C). Similarly, p-norm queries about homes (Q^H) are used to obtain home classification lists (L^H) from home vectors (V^H).

After the classification lists for children and homes have been produced, these lists can simply be correlated. In FOCES, cosine correlation is used to obtain a similarity between each child and each home. For each child, all homes are ranked, and a list of homes is printed in descending order of similarity value. Homes near the top of the list should be most suitable for a given child. It is hoped that a social worker using FOCES in a real-life situation would view this ranked list as a set of suggestions to be carefully critiqued.

5. Evaluation

In order to evaluate the performance of FOCES an experiment was conducted involving two social workers in the Roanoke Department of Social Services. A database of fifteen homes and seventeen children was available, taken from actual case records, with identifying information removed to protect the privacy of all parties involved. A reasonable balance existed between males and females, and between blacks and caucasians, that reflected the actual distributions. Similarly there was variation in age, reason for entering foster care, personal characteristics, and behavior patterns. Many of the homes were approved to care for children of all ages, so matching could not be done in a simplistic fashion. The amount of available data about homes and children varied widely, so selection had to be made based only on partial information.

The two social workers rated each home for each child on a scale of 1 (excellent choice) to 5 (very poor), and so determined a partial ordering. FOCES ranked the set of homes for each child. A set of thresholds were identified to divide up the FOCES match values, to allow comparisons with the ordinal scales used by the social workers.

The social workers agreed with each other, and with FOCES, most often regarding homes that were clearly poor choices. Thus, agreement between the social workers reached almost 90% in some cases, but when homes rated 5 were excluded, the highest rate of agreement was less than 30%. The rate of agreement between FOCES and the ratings given by at least one of the social workers was at least as good as the inter-social worker rating reliability in 70% of the cases. However, this rate of agreement only ranged between 13 and 40% when homes rated 5 were excluded. Apparently, FOCES did a reasonable job at identifying which homes would be appropriate for a child.

Since the normal use of the output produced by FOCES would be for a social worker to look down the list to find all the good homes for a child, it was observed that the measures of recall and precision used in information retrieval systems would be of value in analyzing the experimental results. Instead of working with relevant documents, however, it was necessary to find appropriate homes. Based on the rankings of the two "experts" a home was judged appropriate for a given child if either social worker rated it 1 or 2. At any given point in examining the ranked list for a child, then, recall indicates how many of the acceptable homes have been identified, and precision gives the ratio of acceptable homes found to the total number examined.

Figure 3 shows these results in graphical form. Precision stays quite high, never dropping below .67. Even at a recall of .5, the precision is .84. Thus, FOCES tends to list mostly good homes first. Indeed, in almost all of the cases, the first home on the list was viewed as acceptable by at least one of the social workers.

To give additional insight, it was observed that in eight of the 17 cases, both of the social workers gave a rating of 1 or 2 to the home that was top-ranked by FOCES. Roughly 20% of FOCES' top three choices had been rated 1 or 2 by both social workers. Finally, in ten cases, the home ranked as second best by FOCES was also among those receiving a rating of 1 or 2 by at least one social worker.

Though a more definitive test is needed with a larger collection of data and in a more realistic or closer to operational setting, it does seem clear that FOCES has performed almost as well as the "experts" in this initial evaluation.

Figure 3. Summary of Recall – Precision Results

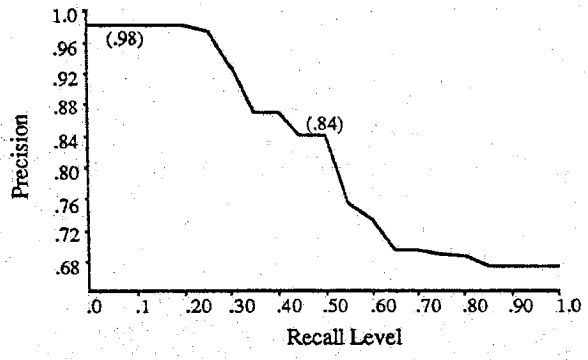

6. Summary and Future Work

Selecting a foster care home for a child is a difficult but important task carried out by social workers, where expert system methods could be of significant assistance. With the help of the Roanoke City Department of Social Services, a two person project team with experience including social work, computer science, and information retrieval, developed FOCES, the FOster Care Expert System. Because of the complexity of the problem domain, it was necessary to supplement conventional knowledge acquisition, knowledge representation, and processing methods with advanced information retrieval techniques. In particular, p-norm extended Boolean queries were used to classify homes and children according to 16 different prototypes, and matching and ranking steps allowed the best homes to be highlighted for a given child.

A preliminary evaluation with data about 17 children and 15 homes was carried out with the aid of two practicing social workers. Each social worker used the data available to FOCES to rank homes as to suitability for each of the children. The agreement between FOCES' ratings and the ratings given by at least one of the social workers was at least as good as the inter-social worker rating reliability in 70% of the cases. In 16 of the 17 cases, the home ranked by FOCES as most appropriate for the child was rated highly by at least one social worker. Using suitably adapted definitions of recall and precision, FOCES performed exceptionally well [35].

Further work with FOCES can proceed in a number of directions. An expert system could be integrated with the current information system of the Department of Social Services, to share existing data and to allow social workers to benefit from advice produced. Post processing of the ranked output list could eliminate homes, if rules were given by the social workers for special conditions not now covered. Additional prototypes could be added, or extra child or home attrributes might be identified, to facilitate more refined classification and matching. Finally, more comprehensive testing could lead to further tuning of the system.

Of greater importance, however, might be the adaptation of methods learned during the development of FOCES to other artificial intelligence applications. Now that a great deal of interest exists in the use of artificial intelligence methods for information systems (eg., [36], [30], [2], [17], [9]) it may be especially appropriate for the artificial intelligence community to consider these and other ways to use techniques studied by information retrieval researchers.

Indeed, there are many situations where objects must be classified, or some type of matching is required, and precise rules are very difficult to identify. In diagnosis problems, for the sake of efficiency it may be useful to rapidly screen data or identify sets of likely candidates that are "close" to prototypical cases. Whenever standard expert systems methods are applied to these classes of problems, then, success may require extensions like those employed in FOCES.

Bibliography

[1] Aiken, Janice S. Prototypes and Production Rules: A Knowledge Representation for Computer Consultation, Stanford Univ. Computer Science Dept. Report No. STAN-CS-80-814, Palo Alto CA, 1980.

[2] Biswas, Gautam, Viswanath Subramanian and James C. Bezdek. A Knowledge Based System Approach to Document Retrieval. In Proc. *CAIA-85*, 455-460.

[3] Bobrow, D.G. If Prolog is the Answer, What is the Question? or What it Takes to Support AI Programming Paradigms. In *IEEE Trans. on Software Engineering*, 11(11): 1401-1408, November 1985.

[4] Bookstein, A. Fuzzy Requests: An Approach to Weighted Boolean Searches. *J. Am. Soc. Inf. Sci.*, 31(4):240-247, July 1980.

[5] Buell, D.A. and D.H. Kraft. A Model for a Weighted Retrieval System. *J. Am. Soc. Inf. Sci.*, 32: 211-216, 1981.

[6] Caplan, Robert D. Person-Environment Fit: Past, Present and Future. In *Stress Research: Issues for the Eighties*, ed. Cooper, Cary L., John Wiley & Son, Ltd., 1983.

[7] Clark, K. and F. McCabe. PROLOG: A Language for Implementing Expert Systems. In *Machine Intelligence*, Vol. 10, ed. Hayes, J. and D. Michie, Ellis and Horwood Ltd., Chichester England, 1982.

[8] Fox, E.A. Extending the Boolean and Vector Space Models of Information Retrieval with P-Norm Queries and Multiple Concept Types. Dissertation, Cornell University, University Microfilms Int., Ann Arbor MI, Aug. 1983.

[9] Fox, E.A. and R.K. France. Architecture of an Expert System for Composite Document Analysis, Representation and Retrieval. *Int. J. of Approximate Reasoning*, 87(2): in press.

[10] Gammack, John C. and Richard M. Young, Psychological Techniques for Eliciting Expert Knowledge, In *Research and Development in Expert Systems*, ed. Bramer, 105-112, 1985.

[11] Hayes-Roth, F., Waterman, D.A. and Lenat, D.B., eds. *Building Expert Systems*, Reading, MA: Addison-Wesley, 1983.

[12] Helly, John Joseph, Jr. A Distributed Expert System of Space Shuttle Flight Control, UCLA Report No. CSD-840038, Los Angeles CA, 1984.

[13] Kadushin, Alfred. *Child Welfare Services*, MacMillan Publishing Co., Inc., NY, 1980.

[14] Kandel, Abraham. *Fuzzy Mathematical Techniques with Applications*. Reading, MA: Addison-Wesley, 1986.

[15] Lee, N.S. and J.W. Roach. GUESS/1: A General-Purpose Expert System Shell. In *Proc. Expert Systems in Gov. Symp.*, Oct. 1985.

[16] Lee, N.S. Programming with P-Shell. *IEEE Expert*, 1(2):50-64, Summer 1986.

[17] McCune, Brian P. et al. RUBRIC: A System for Rule-Based Information Retrieval. *IEEE Trans. on Software Eng.*, SE-11(9): 939-945, Sept. 1985.

[18] Olsen, Lenore J. Predicting the Permanency Status of Children in Foster Care. *Social Work Research and Abstracts*, 18(1): 9-20, NASW: Washington D.C., 1982.

[19] Paice, C.D. Soft Evaluation of Boolean Search Queries in Information Retrieval. *Inf. Tech.: Res. Dev. Applications*, 3(1):33-42, 1984.

[20] Radecki, T. Reducing the Perils of Merging Boolean and Weighted Retrieval Systems. *J. Doc.*, 38(3): 207-211, Sept. 1982.

[21] Rappaport, Julian. *Community Psychology*, Holt, Rinehart & Winston, NY, 1977.

[22] Rich, Elaine. User Modeling via Stereotypes, *Cognitive Science*, 3:329-354, 1979.

[23] Salton, G. and M.J. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, New York, 1983.

[24] Salton, G., C. Buckley and E. Fox. Automatic Query Formulations in Information Retrieval, *J. Am. Soc. Inf. Sci.*, 34(4):262-280, July 1983.

[25] Salton, G., E. Fox, and Wu. H. Extended Boolean Information Retrieval. *Commun. ACM*, 26(11):1022-1036, Nov. 1983.
[26] Salton, G., E. Fox, and E. Voorhees. Advanced Feedback Methods in Information Retrieval, *J. Am. Soc. Inf. Sci.*, 36(3):200-210, May 1985.
[27] Schoech, Dick, Hal Jennings, Lawrence L. Schkade, and Chrison Hooper-Russell, Expert Systems: AI for Professional Decisions, *Computers in Human Services*, 1(1): 81-116, Spring 1985.
[28] Segal, Steven P., Jim Baumahl, and L. Scott Nelson. Seeking Person-Environment Fit in Community Care Placement, In *Issues in Community Residential Cure*, ed. Budson, R., Jossey Bass, San Francisco CA, in press.
[29] Shyne, Ann W., and Anita G. Schroeder. *National Study of Social Services to Children and their Families*, U.S. Dept. of Health, Educ. and Welfare, Washington D.C., August 1978.
[30] Tong, R.M. et al. A Rule-Based Approach to Information Retrieval: Some Results and Comments. *Proc. AAAI-83*, 411-415, 1983.
[31] Van Rijsbergen, C.J. *Information Retrieval: Second Edition*. Butterworths, London, 1979.
[32] Virginia Department of Social Services, Policy Manual, Vol. VII, Section III, Chapter B, July 1984.
[33] Winett, S. and E.A. Fox. Information Retrieval Techniques in an Expert System for Foster Care. In *Proc. Expert Systems in Gov. Symp.*, 391-396, Oct. 1985.
[34] Winett, S. and E.A. Fox. Using Information Retrieval Techniques in an Expert System. In *Proc. Second Int. Conf. on Artificial Intelligence Applications*, Dec. 11-13, 1985, 230-235.
[35] Winett, S. Using Information Retrieval Methods in an Expert System for Foster Care Placement, M.S. Thesis, Virginia Tech Dept. of Comp. Sci., January 1987.
[36] Yip, Man-Kam. An Expert System for Document Retrieval. M.S. Thesis. M.I.T., 1981.