KNOWLEDGE-BASED APPROACHES TO GOVERNMENT BENEFITS ANALYSIS

Marc Lauritsen
Harvard Law School
Cambridge, Massachusetts 02138
laurit@hulaw1.bitnet

Abstract

Determining the government benefits to which clients may be entitled and constructing optimal benefits plans are complex cognitive tasks requiring distinctively legal forms of expertise. Artificial systems that can perform parts of these tasks are of both immense practical value and significant research interest in legal informatics. Building such systems moreover offers insights into how lawyers reason (or should reason) about benefits maximization. This article characterizes the possibilities and problems of modelling the knowledge of the benefits practitioner, highlights aspects of that effort that seem to require artificial intelligence techniques, and sketches an analytical framework in which prior and ongoing work can be structured. A working system created by the author that exhibits some of the desired behavior is described.

Introduction

It is widely lamented in the human services field that only a few of those needing social welfare assistance actually receive all the benefits to which they are entitled. This situation derives from such factors as the complexity of the system, widespread ignorance and confusion about entitlements, scant motivation for government officials to achieve higher "take-up", and a dearth of advocates for potential recipients. As a result, many of the neediest suffer, and policy objectives underlying welfare legislation are compromised. From the perspective of many taxpayers, there is also concern that some recipients may be getting benefits to which they are not entitled, and that the distribution system is not managed as economically as it might be.

Using modern information technologies to address these problems has been more frequently discussed than achieved. Many government benefits automation challenges have yielded to conventional software development methods, but organizational and political factors have slowed progress. Efforts to build recipient-oriented benefits analysis systems, in particular, have remained insular. This article seeks to lay a foundation for steadier progress by summarizing the work that has been done and that remains to be done.

A first look at the practice

Government benefits practice is a challenging legal specialty, one that mirrors the complexity of tax and estate planning. It has mainly been pursued in the legal services context, where low-income clients are served by lawyers who are often inexperienced, frequently overworked, and always too few in number. More seasoned benefits lawyers and paralegals are generally found in "backup centers," law reform organizations, and teaching clinics. It is regrettable that thoughtful examination of the reasoning processes of benefits practitioners is largely absent from the published literature.

The analytical task emphasized in this article — determining the forms of government assistance for which a client may be eligible — makes up only part of benefits practice. In addition to this basic work, benefits lawyers deal with complex procedural issues concerning application, verification, appeal, recertification, termination, and related processes; engage in advocacy before legislative and administrative bodies; and conduct litigation over the legality, constitutionality, and proper interpretation of applicable laws, regulations, and practices.

While decent competence in government benefits analysis is the province of specialists, such issues lurk in virtually every legal services case — indeed, in the cases of all but the wealthiest individuals. The new legal specialty of elder law is very much concerned with these issues. Aspects of this work occur in telephone "hot lines," "help lines," and other human services contexts involving information, referral, and case management.

At a more abstract level, the identification process going on here is ubiquitous in law and related contexts. The government can be replaced with the nonprofit sector, the private sector, or the sociolegal world at large. Entitlements or opportunities may not be as clearly defined in the latter cases, but the matching process is comparable. Given the circumstances (actual or engineerable) of my client, what good things can he or she get? The cognitive mechanics are similar, whether those goods are welfare checks, tax shelters, research grants, business opportunities, or personal injury awards. These commonalities make research in this area both derivative from and contributory to a broad current of AI and law research.

Some of the task's complexities

Client benefit maximizing is a quintessential lawyering activity. It involves iterative processes of gathering facts, researching law, counselling clients, consulting with colleagues, and negotiating and arguing with government representatives. Even putting aside the task's affective and more advanced strategic dimensions, attempts to model its cognitive core quickly encounter substantial difficulties. These can be broken down into those that relate to the facts, the law, the client's goals, and the lawyering methods employed.

Facts

The benefits practitioner, like any lawyer, must operate in a world of "noisy" information: information that is often approximate, incomplete, inaccurate, inconsistent, and uncertain.

Some of these factual deficiencies flow from the client, who may be forgetful, misinformed, reluctant to disclose information out of pride or embarrassment, or plain dishonest. Generally there isn't time to get all the relevant facts, or the costs of ascertaining them outweigh their putative benefits. Some historical facts cannot be ascertained at all. Assumptions must often be made. Knowing when and how to make such assumptions, and keeping track of their implications for the reliability of one's assessments, is a characteristic trait of benefits experts.

Some uncertainty derives from the inherent future-orientation of benefits work. The optimality of arrangements normally depends upon what will actually happen in the future — e.g., the importance of retaining medical coverage will depend upon how well one's health holds out; a steady income stream may be preferable to a lump sum payment only if one lives long enough. Projecting what will be is rarely a factor in establishing particular entitlements, but quite often a consideration in benefits maximization.

Plasticity of circumstances is another contributor to factual complexity: in some respects client situations can be intentionally designed, as in tax practice and estate planning. Not only is the present constellation of facts important, but the facts of all worlds that are "accessible."

An adequate model will need to take into account not only the semantic content and truth value of myriad facts, but their degree of mutability and epistemic characteristics such as pedigree (source of the information), derivation, and tentativeness.

Law

The language in which entitlements to benefits are couched purports to be logical and deterministic. Like tax legislation, it is full of structural expressions and controlled vocabulary. It is equally ridden with syntactic ambiguity, open-textured terms, and plain logical errors. One almost admires the ingenuity of the regulation writers in mixing disparate logical forms — declarations of entitlements alongside of recipes for grant calculations and procedures for verification — as though they were writing the final exam for a course in knowledge representation. Much benefits jurisprudence arises from advocates' efforts to counteract official tendencies to construe this melange of speech acts in favor of government and to overlook principles of due process and equal protection.

The analytical complexity flowing from structural ambiguity and open texture is amplified by some more specific factors.

One factor is the sheer quantity of applicable law. The benefits practitioner doing a global assessment of a client's possible entitlements must attend to statutes and regulations on federal, state, and local levels in dozens of different programs. Laws previously in force, as well as those likely to be enacted in the future, may need to be consulted. A definitive benefits library could comprise thousands of volumes.

A second factor is the rapidity of change in the law. Not only do the standards, formulas, and rules change with alarming frequency, but entire conceptual structures can shift quickly, sometimes as the result of misleadingly minor changes in legislation or a regulation. Unlike real property law, where entitlements are essentially guaranteed in perpetuity, there is little assurance that future benefits regimes will be like those of today. (As this article goes to

press, the Governor of Massachusetts is considering abolishing the "General Relief" program, which has been an income source of last resort for some 40,000 people.)

Thirdly, the corpus of relevant law comprises many heterogenous components, few of which were designed to interact harmoniously. Identically named predicates (e.g., "disabled", "adjusted income") often have different meanings in different programs. One frequently needs to think in terms like "disabled for purposes of General Relief" and "Food Stamps adjusted income." Moreover, benefits problems have particularly indistinct boundaries with other legal and non-legal domains, such as family law, housing law, health law, consumer law, and social work.

Goals

Even in a world of perfect information and thoroughly disambiguated legal rules, ascertaining and optimally achieving client goals would hardly be unproblematic. Answering the basic teleological question "What is the maximand?" with "total client utility" only begins the analysis. The answer depends upon the client's unique needs, wants, values, attitudes, and possible futures. Behind the question is a whole complex of tradeoffs. There are three main dimensions along which these tradeoffs are made:

- Nature of benefit — monetary, in-kind, or service? Restrictions on use or alienability generally make benefits less valuable, but client rankings will differ.

- Beneficiaries — self, household, family, or others? Clients are self-regarding and altruistic in different degrees.

- Time — present or future receipt? People have different time/value spectra, different attitudes toward risk, security, and deferred gratification. The relative valuing of initial, continuing, and future eligibilities will depend upon these attitudes.

In addition to the inevitably subjective utility comparisons among the foregoing, goals are clouded by clients' differential valuations of the costs of receiving government benefits, such as loss of privacy, submission to behavioral regulation, hurt pride, and the ordeal of completing applications, compiling supporting documentation, and making ongoing reports. Some clients may not want unearned "handouts," and the long-range interests of many will lie in the direction of reduced reliance on governmental support. Most benefits advocates accordingly see their role as that of maximizing client choice, not client income.

Clients are not the rational economic actors of academic theory. At bottom, the tradeoffs implicit in choosing benefits goals are among incommensurables, and not susceptible to logical-mathematical analysis. But there are rational processes by which goals can be clarified and probable maxima framed with reasonable confidence. These processes comprise part of the knowledge of the benefits practitioner.

Methods

Just how do benefits practitioners manage all these different forms of complexity? The routines of interviewing, research, analysis, and counselling rest on complex knowledge structures that are unlikely to bottom out in neatly algorithmic substrata. Heuristics of matching and maximization seemingly entail a rough mix of running through internalized checklists, selecting paradigms, and following litanies of permitted transformations. (E.g., "Move X dollars into child A's account in order to render that child ineligible for Program X, thereby making the residual family eligible for Program Y"; "See how much difference it makes if child A and her cousin are considered as belonging to two different assistance units.") The process is one of abductive interplay among emerging facts, working hypotheses, and constructed theories.

The methodological knowledge driving benefits practice is conceptually distinct from the factual, legal, and teleological knowledge described above. My early impression is that experts in this area are both responsive to and benefitted by efforts to formalize their heuristics.

An ontology for the benefits world

A simple model of the benefits problem space involves a world of three kinds of entities — people, assistance units, and benefits programs. The former two can be connected by links of membership and the latter two can be connected by links of entitlement and/or receipt. (People, of course, are connected by links of relationship.) The structure and content of this world vary with time, with each moment leading to multiple possible futures that reflect both the plasticity of the present and the constraints of prior history.

People include clients and other individuals whose circumstances have implications for the clients' situation by virtue of some benefits program rule. For each person, an

array of historical and current circumstances needs to be represented as object attributes. The necessary predicates include:

- age (date of birth)
- marital status (and history)
- health status (e.g., disabled, pregnant)
- amounts and types of income, past and present
- employment history and status
- educational history and status
- current and past expenses
- assets
- place of residence
- citizenship

Client characteristics are usefully partitioned into those that are immutable — such as date of birth, race, veteran status, employment history — and those that are more or less modifiable, such as place of residence, nature and value of assets, household composition, etc. The degrees of freedom in the latter characteristics can expand the universe of potential benefits packages.

Assistance units are groups of one or more persons. They are characterized by such group predicates as the identity and number of their members, their inter-relationships, and the presence or absence of minors, elderly, disabled, or other members with special characteristics. These and most other group predicates (such as total income, assets, expenses, etc.) are derivative from those of the component individuals, but some are not so decomposable — for instance, whether or not all of the unit's members prepare food and eat together. The number of possible different assistance units into which a group of n people can combine is (discounting the empty set) $2^n - 1$.

Note that flexibility in the rules sometimes allows definition of assistance units to be an elective matter, and it is not uncommon for one physical household to comprise multiple logical assistance units.

Benefits programs are characterized in terms of such factors as:

- type — monetary, in-kind, or service
- area of need — food, housing, utilities, health, etc.
- requirements and disqualifications, framed in terms of assistance unit characteristics (including, not infrequently, what other benefits a unit receives or is eligible for)
- correlations of assistance unit characteristics with particular kinds or levels of benefits

In this scheme, a basic task of benefits analysis is to identify existing and possible entitlement links between all plausible units and programs. Optimization for any given assistance unit involves finding the set of entitlement links which, when activated, produce the greatest flow of benefits to that unit (as measured in terms of the client's utility, and net of client costs.) One can easily see how, absent powerful control strategies, the hypothetical analyst would quickly become swamped in a combinatorial sea. A household of 5 produces 31 logically possible assistance units; entitlement assessments vis-à-vis, say, 30 benefits programs requires 930 analyses before one even begins to consider cross-impacts of benefit receipts. Brute force consideration of all complete unit-program correlation sets would require review of astronomical numbers of possibilities.

Aspects of an ideal government benefits practice system

It is instructive to consider what an ideal computer-based system for support of benefits practice might look like, given technology reasonably available today. Naturally, the "system" would almost certainly not be a single software program or hardware environment, but rather a distributed combination of programs, possibly spanning several environments. (I assume that we are limiting our aspirations to a "power tool" for human use, perhaps even an "intelligent assistant", and not yet envisioning a semi-autonomous replacement for human experts. Let's put aside for now such more exotic possibilities as robotic agents that would interview clients, review welfare department records, complete applications, and conduct routine hearings.)

First of all, we want a thoroughly informed system. It should have available, through a combination of local and online services, the text of all relevant authorities, past, present, and proposed, including all figures, tables, formulas, standards, etc. Every published or unpublished text that might profitably be consulted in the course of handling a benefits matter should be available.[1] This of course includes all relevant court orders and decisions, administrative decisions, interpretative letters, planned and pending challenges, and secondary literature. Assured, prompt, and reliable access to this welter of information would alone be of enormous value.

The system should be equally comprehensive in its provision of local or site-specific information, such as (1) the names, addresses, telephone numbers and other details about government offices, social workers, human services providers, and the like; and (2) the profiles and histories of an office's clients. Such information should be as easily

navigable as that described above. Statistical and other analytical operations should additionally be supported.

Second, in interfacing with users, the system should observe several "fair information practices." Users should not have to provide information or make choices that are not needed for the assistance requested. In other words, for any given information or analysis sought, the result should be made available with as few questions as possible. Moreover, the user should be able to ascertain the informational cost of getting answers before incurring those costs.

The system should ask for information in a sequence that is most coherent, intelligible, and comfortable. Multiple views of data, and modes of entering data should be provided. It must be easy to back up and correct answers. Sometimes these goals can conflict with the above goal of efficiency.

To perform some of these functions well, the ideal system needs to be informed or knowledgeable about its own knowledge, and about the knowledge requirements of particular tasks it is called upon to perform. That is, the system needs to know not only what it does know, but what it needs to know to accomplish particular functions.

Most users will find an "electronic treatise" mode extremely useful. Building on its comprehensive textual knowledge base, the ideal system will combine the best features of hypertext delivery systems and conceptual retrieval tools for large databases. Users should be able to get almost instant access to any desired text through whatever navigational path or search query they find most convenient. Traversing the links of nested definitions and cross-references, for example, should be as easy as possible. Since these kinds of features are peripheral to the main themes of this article, though, they will not be elaborated here.

Equally important, but also outside the scope of this article, are most other features of an ideal user interface. Creative designers would undoubtedly want to experiment with such things as analog manipulation of inputs (e.g., virtual knobs or levers for specifying approximate income and asset information); hypermedia touches like audio and video reference material; and voice annotation of user data.

The system's characteristics in any given session should substantially depend on who is using it. Benefits experts have different needs in a computational assistant than do novices, generalists, lay advocates, or clients. The ideal system will offer modes of interaction appropriate to each of these user types, while drawing maximally on common underlying informational content and function.

Third, the ideal system should be broadly competent. Its basic capabilities should include:

- organizing client information for ease of access, analysis, and updating by the user;
- determining the benefits for which the client is or could be eligible;
- calculating benefits levels;
- recording case activities (and offering checklist and reminder facilities);
- assembling customized documents such as summaries of case information, histories of case activities, reports of analyses and calculations, and letters to clients and to government agencies.

Generally the system and user should go about their joint tasks (such as issue spotting) in a mixed initiative manner. Fact patterns that indicate an issue or problem should be noticed by the system, and made known to the user in a non-obtrusive way.

The system should offer explanations of both its questions to the user (explaining why they are being asked) and its conclusions (explaining how they were derived.) It should explicitly identify assumptions and simplifications that enter into its analyses. It should be capable of suggesting arguments and justifications for positions the user might want to take on controversial matters.

Knowledge inspection and revision facilities round out the list of important features. It should be possible for the user or other non-programmer to review and revise facts, criteria, rules, and strategies that are applied by the system in its work.

The imaginative benefits practitioner with some knowledge of automation possibilities could quickly add to this list of desiderata, but it should suffice for present purposes. Building a truly comprehensive benefits maximization system would be a monumental undertaking — a benefits world analogue to the human genome mapping project, with none of that project's universality or timelessness. Expending resources for a knowledge engineering project of such scope would require careful justification.

MicroMax

MicroMax[2] is an interactive government benefits questionnaire, issue analyzer, and case management tool.

It was created during the summer of 1990 under a contract with Brigham and Women's Hospital in Boston, for use in a research project studying the interplay of medical and legal problems in the urban poor. It has been used there in several dozen cases by an attorney and other staff, and a more generic version is now being tested in six human services offices and in Harvard's poverty law teaching clinics.

The user provides information about a client household by either sequencing through a dialogue of questionnaire-like screens or choosing particular topics from a menu. In either case, the information gathering "frames" are dynamic: that is, their presence and content may depend upon earlier data. The user can enter and modify all information in either mode from any point. Summaries are readily available on-screen and printed.

Automatic analysis is offered in several ways. An "issues summary" screen allows the user to track the status of some twenty issue categories, and indicates those categories in which an issue has been spotted by the system based on its programmed logic. Alternatively, the user can view a "household spreadsheet", which summarizes (in editable mode) key data about all members, and displays all issues flagged by the system given the current data. Narrative summaries of the system's assessments, and of any calculations performed, are available as hypertext. "What-if" analysis can be performed by changing data and seeing a revised list of issues spotted and benefit levels projected. Finally, the user can generate a customized synopsis of all of the system's assessments for the current household.

Several case management facilities are supplied, such as electronic worksheets for each issue, an activity log, and a "Rolodex" facility. All are linked together in an environment that permits easy movement among functions, use of a WordPerfect-like editor, storage of data between sessions, import and export of text and data, etc. I plan to integrate some of the more "intelligent" features into these case management tools (by, for instance, automatically placing entries in the user's notepad or activity log.)

MicroMax was built using CAPS/Author^{TM}, a practice system development environment that combines modular programming and other object-like features with conventional fourth-level application development tools. The programmer can make any number of instances of fifteen types of "elements", whose characteristics are defined by completion of questionnaire-like templates. The cross-referencing of elements in appropriate modes and locations constitutes the principal method of specifying a system's behavior. CAPS can usefully be thought of as an object-oriented language in which the classes have largely been pre-defined. A longer discussion of CAPS can be found in [Lauritsen 1989] and [Lauritsen 1990].

While not an expert system shell, CAPS has built-in routines for generating dialogues with users to gather data needed to perform the functions it is called on to perform. Users can work backwards by asking MicroMax for a document or assessment, and only provide the information called for. The resultant interaction is not always the efficient and coherent one idealized earlier in this article, but it is a useful supplement to the pre-crafted dialogues I have supplied. I am exploring features in CAPS that should make it possible for me to place more intelligent controls on the questioning sequences when they are automatically generated.

I separate domain knowledge from control processes by storing all substantive rules in a CAPS "table" element. Each row of that table consists of the name of a benefits program, a "computation" element which returns a true or false value, and a "document" element which returns customized text explaining eligibility and reviewing details of any calculations performed. Both the computation and document elements have internal logic that typically calls subsidiary elements of various kinds. This modularity at several different levels facilitates specification and testing of knowledge in manageable chunks.

The appendix includes illustrative screens from a MicroMax session on behalf of a fictional household.

Related work

David du Feu and others designed and implemented a system (WELBEN) for assessing eligibility of households for means-tested benefits in Scotland in 1975 as part of the Inverclyde Project [du Feu 1980]. The system consisted of 3 programs written in ANSI COBOL, running on an IBM 370: an input data validation program, a calculation program covering nearly sixty different assessments, and an output program, which produced a letter to the family summarizing probable entitlements. External data files contained the benefits rates, output statements, and other updatable information. A control table, modified by the program as it proceeded, was used to drive the calculation program.

Du Feu's 1980 article raises many fascinating administrative and policy-related questions, and describes work very similar in aspiration to my own. Sadly, the Inverclyde project — hailed as a technical success — was deemed an organizational failure and discontinued.

The ELI project of the early 1980s took a production rule approach to some aspects of British welfare law, using a complex interpretative procedure that interwove forward- and backward-chaining recursive strategies. The positivistic fervor with which this project apparently commenced was ultimately eclipsed, however, by the author's conviction that even the clearest legal rules are non-deterministic and that expert systems that attempt to axiomatize them are of little use in the legal process. [Leith 1986] supplies a trenchant critique of the naive ruled-based approach and offers some fertile suggestions for how annotation, multiple sources of expertise, rule-breaking heuristics, and similar components might combine to form a practically useful (and theoretically interesting) system.

In 1986, the Logic Programming Group at London's Imperial College undertook to represent a large portion of the United Kingdom's welfare law: specifically, legislation relating to "Supplementary Benefit." [Bench-Capon et al. 1987] Previous work on this law is described in [Hammond 1983]. The project's aim was not to build a practical system, but rather to support research on knowledge representation using logic programming formalisms. ([Bench-Capon 1987] discusses how logical models of welfare legislation can be used by policy makers to assess problems in current law and the impact of proposed revisions.) The APES system, which augments PROLOG with facilities for automatically generating dialogues with the user, was used to execute prototype formalizations.

Several points made in the Imperial College project's 1987 article are worth noting: (1) One can best proceed with a formalization of legislation that is neutral as to how it is used — that is, one that is not constructed for a particular use. (2) The level of detail to which predicates should be decomposed is governed by the extent to which subordinate predicates are used independently in the legislation. (3) A special event calculus had to be devised to deal with temporal reasoning issues in the Supplemental Benefit legislation. (4) "Deeming" provisions do not appear to yield to any simple general representational treatment.

Professor Brian Jarman, Tim Blackwell, and others at St. Mary's Hospital Medical School in London (interestingly also part of Imperial College) developed the Lisson Grove Welfare Benefits Program in 1982 to calculate claimants' entitlements to over thirty types of British social welfare benefits. Jarman, a G.P. at Lisson Grove Health Center, found that some of his patients' ill health was attributable to poverty caused by non-receipt of benefits to which they were entitled. The program is written in Basic and runs under MS-DOS. It presents a series of questions about the claimant's circumstances, announces entitlements as they are confirmed, offers online help on questions and benefits, and generates case summaries. Version 2.70 is used by over 600 organizations at approximately 2000 sites in the United Kingdom.

Harvard Law School's Project Pericles undertook to build several automated government benefits worksheets, initially as macros within the Lotus 1-2-3 spreadsheet program, and more recently in Pascal code, using the PowerScreens™ utilities. Worksheets for Aid to Families with Dependent Children, Food Stamps, and Supplementary Security Income have been completed and distributed free of charge by the School's Educational Technology Group. These worksheets — deliberately kept simple so as to run independently and easily on low-end personal computers — are limited to a single benefit program at a time, and do not produce any customized documents or narrative analyses. They do perform quite extensive calculations and offer numerous explanatory helps through a hypertext facility.

The CAPS work described above also had its roots in Project Pericles, during which Harvard became a research site for the VMS version of CAPS developed at Brigham Young University. Visiting Professor Larry Farmer and I worked with several generations of law students, some of whom built systems in the government benefits area. [Lauritsen 1991] summarizes some of this experience.

The United Seniors Health Cooperative in Washington D.C. has developed a "Benefits Outreach and Screening Software" (BOSS) package designed to help organizations screen older persons for a wide range of benefits. As in the Inverclyde project, a individualized report is produced for the client.

Peggy Jo Duffy implemented an expert system in Quintus Prolog to determine whether clients of a Philadelphia legal services program were eligible for Medical Assistance. Her thesis [Duffy 1990] focusses largely on methods for evaluating and verifying such a system. MAE (Medical Assistance Expert) uses an agenda driven forward chaining control mechanism embodying a "think, act, update" procedural strategy. Its knowledge base includes inference rules (specifying when to add new information to the dynamic database), query rules (specifying circumstances under which questions get placed on the agenda), and advice rules (specifying when to display a statement of advice). Questions are priority weighted to assist in their sequencing. If a user responds with "unknown," the system can provide a default value in order to continue, and remind the user at the end of the session which questions were answered in this manner and their repercussions.

The Anderson Consulting firm has created an expert system called MAGIC (Merced Area Global Information Communication System) that guides caseworkers in Merced County, California through interviews of prospective government benefits recipients and determines the benefits for which they are eligible.[3] This MS-DOS program was developed using Aion Corporation's ADS and is accessed on 200 workstations under Unix. Dramatic reductions in social worker time, paperwork, and error rates are claimed. Consisting of some 4,200 rules, the system reportedly cost $13.5 million and took over 135,000 hours to develop by a team of sixty consultants and twenty county employees.[4]

More generally, computerized eligibility determination systems have long been in place on the government side and innovative benefits automation efforts are underway in a number of states. In Tulare County, California, for instance, a touch-screen, multi-lingual benefits advice system called CLEAN has been developed for direct applicant use.

There is of course a larger body of reported research on planning and analysis in similar legal contexts, only some of which I will mention here.

Schlobohm's and McCarty's collaboration on estate planning [Schlobohm 1989] develops the notion of prototype plans as instruments attorneys use to deal with otherwise intractably large spaces of possible solutions. These prototypes are critiqued and "debugged" to achieve better matches with client goals. Representations of trusts and the Internal Revenue Code in McCarty's Language for Legal Discourse are given as a first step toward an expected implementation of an "EPS II" system. Two notable aspects about this effort are the perceived importance of deep representations of estate planning law, and the role of conceptual coherence in making large statutes like the Internal Review Code cognitively manageable.

Jon Bing has described [Bing 1987] how legislation such as the Norwegian Social Security Act can be shown in normalized form through arrow diagrams, and how the underlying rule structure can then be used to enhance text retrieval. His ideas about a norm-based thesaurus have promising applications to systems that need to integrate large text bases with inferencing.

Thomas Gordon's articulation of a framework for spotting issues in "interpretation spaces" and constructing legal arguments [Gordon 1989] combines nonmonotonic reasoning, truth maintenance, and natural deduction theorem proving in an elegant combination that seems powerfully suited for the less trivial dimensions of benefits analysis described earlier.

Future Directions

Much work remains to be done in applying knowledge-based techniques to government benefits analysis and in exploring the cognitive mechanics of both human and artificial experts. Several projects have confirmed that some of the basic computational tasks are quite tractable. A growing body of thought and experience is accumulating on the knowledge engineering challenges facing us in this area. These research and development advances need to be cross-fertilized and consolidated.

While at the upper end of technology currently being built and used by lawyers, MicroMax is little more than a friendlier-than-usual low-end expert system. It is an improvement over previously available tools, but I regard it as a bare start. Its main contribution may be to catalyze some rethinking of client-oriented benefits practice. Resources permitting, I plan to steer my own efforts in the direction of further object-oriented strategies that begin to exploit deeper models of benefits expertise, and try to better understand the distinct forms of knowledge needed by an intelligent assistant capable of serving experts and non-experts alike.

Notes

1. In addition to the dual availability of this information as both printable/displayable images and streams of ASCII characters, the system will need it represented in structured forms suitable for higher modes of processing. Ideally, its overt syntax would be exhaustively represented via embedded markers that capture, for example, section and paragraph divisions, cross-references, citation tags, and other meta-level information. See, e.g., [Heather 1988], [Routen 1989].

2. This name is a take-off on the "Little Max" book [Williams 1988], published now for some 13 years in Massachusetts as a compendium of government benefits information and strategies.

3. Harvey P. Newquist, "Anderson knows applications," AI Expert, December 1990, p. 64.

4. Henry Koltys, "Legal Knowledge-Based Systems in the Real World," p. 8 (1991) [unpublished manuscript on file with author].

References

Bench-Capon, T.J.M. "Support for Policy Makers: Formulating legislation with the Aid of Logical Models." Proceedings of the First International Conference on Artificial Intelligence and Law, pp. 181-189. Boston, 1987.

Bench-Capon, T.J.M. et al. "Logic Programming for Large Scale Applications in Law: A Formalization of Supplementary Benefit Legislation." Proceedings of the First International Conference on Artificial Intelligence and Law, pp. 190-198. Boston, 1987.

Bench-Capon, T.J.M. "Deep Models, Normative Reasoning and Legal Expert Systems." Proceedings of the Second International Conference on Artificial Intelligence and Law, pp. 37-45, Vancouver 1989.

Bing, Jon. "Designing Text Retrieval Systems for 'Conceptual Searching'." Proceedings of the First International Conference on Artificial Intelligence and Law, pp. 43-51. Boston, 1987.

Blackwell, Tim et al. Lisson Grove Welfare Benefits Program User Manual, Version 2.70, October 1990.

du Feu, David. Selecting Welfare Benefits by Computer. In Niblett, Computer Science and Law, pp. 183-193 (1980)

Duffy, Peggy Jo. Medical Assistance Eligibility Expert: Implementation of an Expert System with a Validation and Verification Plan. Masters thesis, Villanova University, December 1990.

Gordon, Thomas F. "Issue Spotting in a System for Searching Interpretation Spaces." Proceedings of the Second International Conference on Artificial Intelligence and Law, pp. 157-164, Vancouver 1989.

Hammond, P. "Representation of DHSS Regulations as a Logic Program." In Proceedings of the 3rd BCS Expert Systems Conference. British Computer Society, Cambridge, pp. 225-235 (1983)

Heather, Michael. "Semiotic Orders in Law." In Walter, ed., Computer Power and Legal Language, pp. 143-153 (1988)

Lauritsen, Marc. "Representing Legal Knowledge in Practice Systems." Proceedings of the Third International Conference on Logic, Informatics, and Law (Florence, November 1989)

————. "Delivering Legal Services with Computer-based Practice Systems." 23 Clearinghouse Review 1532 (April 1990)

————. "Project Pericles in Retrospect." 5 Yearbook of Law Computers and Technology, pp. 50-57 (1991)

Leith, Philip. "Clear Rules and Legal Expert Systems." In Martino, ed., Automated Analysis of Legal Texts, pp. 661-679 (1986)

Routen, Tom. "Hierarchically Organized Formalisms." Proceedings of the Second International Conference on Artificial Intelligence and Law, pp. 242-250, Vancouver 1989.

Schlobohm, Dean A. and McCarty, L. Thorne. "EPS II: Estate Planning with Prototypes." Proceedings of the Second International Conference on Artificial Intelligence and Law, pp. 1-10. Vancouver, 1989.

Williams, Lucy, ed. Little Max: Creating Maximum Benefits for Children, Elderly, Poor, and Disabled People. Massachusetts Law Reform Institute, 1988.

Appendix

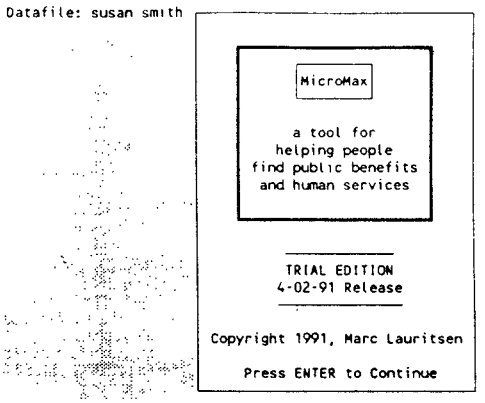

Opening Screen

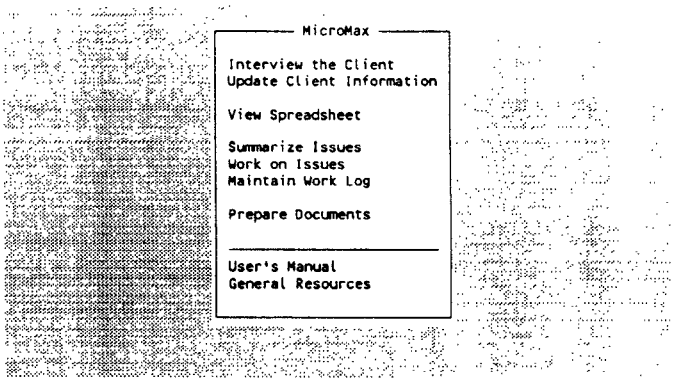

Main Menu

```
2. Basic Client Information

First name:       Susan
Middle name:      R.
Last name:        Smith
Street address:   14 Cullen Lane
City:             Waltham
State:            Massachusetts          Zip:  23344
Telephone:        454-5588
Hours to call:    early evening
Social Security #: 823-45-3334
Date of birth:    March 3, 1956          (Age:  35 )
```

Information Gathering Frame

```
Summarize Issues

Name of issue             Status  Identified  Calculated  Follow-up
Medicaid                  a  ♪    08/12/90    08/15/90    _____
GR/MA                     n
Free care                 p  ♪    _____       _____       _____
Medicare                  a       _____       _____       _____
Commonhealth              n
OASDI                     a  ♪    09/12/90    09/12/90    _____
SSI                       p  ♪    _____       _____       _____
GR                        n
AFDC                      a  ♪    08/15/90    09/03/90    _____
Food stamps               a  ♪    07/24/90    _____       _____
WIC                       a  ♪    _____       _____       _____
Child care food progr     ok ♪    _____       _____       _____
USDA TEFAP                n  ♪    _____       _____       _____
Misc. food programs       n  ♪    _____       _____       _____
Elderly nutrition         a  ♪    _____       _____       _____
Rent seems too high       n       _____       _____       _____
Family issues             a       _____       _____       _____
                                                          More ↓
```

Issues Summary

```
Update Client Information

Household Members
  Aaron (step-son)
  Ben Smith (father)
  Ellie (daughter)
  Susan R. Smith (client)
  Willie the Lion Smith (uncle)
Current Benefits (entire household)
Selected Household Expenses Grid
Benefits Terminated
Support to Family Members Not in Household
Information on Spouse or Former Spouse
Susan R. Smith's Employment History
Susan R. Smith's Monthly Expenses
Utility Needs
Rent or Own
Issues for Renters
Repair Problems
Transportation and Home Services
Debt Problems
                                          More ↓
```

Information Update Menu

```
Food stamps worksheet                                  Unanswered

Issue:          Food stamps
Status:         a
Status note:    administrative appeal pending

Date first identified:          July 24, 1990
Date calculations completed:    _____
Date of one-year followup:      _____

Working notes:

   We are taking the position that Uncle Willie does not
prepare food and eat with the family, and hence that he
should not be considered part of the household for the
purpose of calculating food stamp benefits.  This is based in
part on the recent decision in Smithers v. Brook, which in
turn construes the regulations appearing at 106 C.M.R. § 361.

MicroMax Analysis [1]                          Worklog [2]
                                                  More ↓
```

Issue-specific Worksheet

```
View Spreadsheet

 Name      Age Rel Mar Stu Vet Hea Ins Emp  GrWages   NetSelf   Unearned
x Aaron    4   sso s   n   n   g   n   u    $0.00     $0.00     $0.00
x Ben      79  f   s   n   y   g   ma  u    $0.00     $0.00     $275.00
x Ellie    1   dau s   n   n   g   ma  u    $0.00     $0.00     $0.00
x Susan    35  pat sep n   n   g   n   e    $480.96   $0.00     $390.00
x Willie   50  unc s   n   y   g   n   e    $0.00     $250.00   $0.00

                Total monthly income: $1,395.96
                Percent of poverty level: 112.94%
          (Add/delete members [1] — Abbreviations list [2])

Medicaid [3]                        Free care [4]
OASDI [5]                           SSI [6]
AFDC [7] ($348.33) [8]              Food stamps [9] ($267.10)
WIC [10]                            Child care food program [11]
USDA TEFAP [12]                     Misc. food programs [13]
Elderly nutrition [14]
```

Household Spreadsheet

```
MicroMax

Interview the Client
Update Client Information
  Prepare Documents
    Patient Interview Report
    Summary Sheet
    Issues Spotted by System
    Checklist
    Work Log
    Suggestion Memo
User's Manual
General Resources
```

Documents Menu

```
AFDC

N  The currently specified household (Aaron, Ben Smith, Ellie, Susan  ned
—  R. Smith, and Willie the Lion Smith) seems financially eligible  —
x A for AFDC.  A preliminary analysis indicates that the monthly     .00
x B grant amount should be $348.33.  This was calculated as follows: .00
x E                                                                  .00
   First test:  The household's gross countable income is $1,015.96, .00
1  which is less than the AFDC eligibility standard for a filing     .00
2  unit of 5: $1,332.00.  (Any current AFDC and Food Stamps benefits —
3  have been disregarded.)
4
5  Second test:  The household's net countable income ($371.66) is   —
6  less than the applicable needs standard ($720.00).  Net countable
7  income was figured for each individual as follows:
8
9  Aaron                                    $0.00
1
1     Earned income:                        $0.00
1     Countable unearned income:            $0.00
1                                                          More ↓
```

Hypertext Analysis for AFDC Program

```
Issues Spotted by System

Food stamps

The currently specified household (Aaron, Ben Smith, Ellie, Susan
R. Smith, and Willie the Lion Smith) seems eligible for food
stamps.  A preliminary analysis indicates that the monthly coupon
allotment would be $267.10.  This was calculated as follows:

Total earned income ($730.96)
     x 80%                          $584.77
Total unearned income:      +        665.00
Standard deduction          -       $112.00
Dep. care deduction:        -       $440.00
Medical exp. deduction:     -        $40.00
Preliminary adjusted net:           $657.77
Shelter deduction:          -       $238.11
Total net mthly income:             $419.65

Maximum allotment for
household of 5                      $393.00
30% of net income:          -       $125.89
Projected allotment:                $267.10
                                    Pg 3 Ln 21 Pos 0
```

Summary of Issues (mid-document)