Adaptive decision support for planning under hard and soft constraints

Alex Dekhtyar
University of Kentucky
Dept. of Computer Science
dekhtyar@cs.uky.edu

Raphael Finkel
University of Kentucky
Dept. of Computer Science
raphael@cs.uky.edu

Judy Goldsmith
University of Kentucky
Dept. of Computer Science
goldsmit@cs.uky.edu

Beth Goldstein
University of Kentucky
Dept. of Ed. Policy Studies
bethg@uky.edu

Cynthia Isenhour
University of Kentucky
Dept. of Anthropology
Cynthia.Isenhour@uky.edu

Abstract

We describe the "Welfare to Work" scenario, and the software we are designing to support case managers' planning for their clients.

The Changing World of Welfare to Work

President Clinton signed the revised welfare legislation, "Personal Responsibility Work Opportunity Reconciliation Act (PRWORA)" in 1996. This legislation stipulates a set of supports and regulations for welfare recipients that aim to move those recipients into the paid labor force. Through federal block grants to the states, "Welfare to Work" (WtW) recipients may access such services as financial support, health and mental health services, child care, transportation, and literacy and job-skills training. Each recipient, or *client*, may receive a total of sixty months of services. These 60 months need not be contiguous; clients strategize how to go in and out of the programs to establish economic self-sufficiency and/or maintain later WtW eligibility. The key decision making to access benefits and services occurs during discussions between the WtW clients and program case managers, who play the role of advisors and regulators.

The programs are affected by a plethora of frequently changing mandates, laws, rules and regulations. For instance, agencies must maintain a certain proportion of their clients' time in federally defined "countable" activities, namely those leading directly to employment. These requirements are often at odds with the needs and preferences of the clients. Services recommended by a particular agency or a particular case manager depend on local regulations, preferences, and availability of resources and services.

The typical case manager has a nearly unmanageable case load, and must thus rely on various shortcuts in advising. These seem to include using broad templates for client needs and ignoring certain data elicited from the clients. Furthermore, case managers have a hard time keeping up with changing regulations and availability of services. Thus, automated decision support tools can make their jobs more manageable.

We are working on two components of decision-support software: The model-building component and decision support itself. The underlying paradigm for our decision-support component is planning with constraints under uncertainty.

Given the 60-month limit on services, it is crucial that advice from case managers or software take into account longer-term plans, rather than immediate gratification. The constraints on planning arise from the client's preferences and her limitations, and from availability of services. For instance, a low-literacy client cannot be expected to succeed in college courses, nor can a client without a car arrive at a site unreachable by public transportation. Other constraints arise from regulations, such as a limit to how many months of volunteer work a client may use to satisfy the work requirement. Uncertainty arises whenever case managers make judgments about the likelihood of a client's success in a specific activity. Factors such as dependents' unstable health, the client's physical and mental health, and transportation problems affect client participation in advising and success in planned activities.

The Welfare to Work world is in constant flux. Legislative and administrative bodies change regulations in response to legal, political and budgetary considerations. Service availability changes from day to day, whether services are sponsored by government agencies or by private organizations and charities. And client preferences shift as clients learn more, and as their family, physical, mental and economic conditions change.

We present an overview of current and planned software for decision support for Welfare to Work programs in Kentucky. In Section 2, we briefly describe our approach to building decision-support software for the Kentucky Temporary Assistance Program (Kentucky's WtW program). We then give a fictional case study, and use it to illustrate needed software. Finally, we briefly outline key challenges we are encountering in this process.

Planning with Uncertainty and Constraints

As mentioned in the introduction, we approach decision support in a Welfare to Work setting, as well as other settings that involve advising as a key component of decision making, as a problem of planning with uncertain information in the presence of constraints.

Figure 1 illustrates the architecture of the decision-support system for WtW we are currently building. In this diagram, white boxes represent software components, and dark boxes represent data used for inference and planning.

There are three major stages in this design: (a) model building, (b) model refinement, otherwise known as *knowledge-based model construction* or *KBMC* (Breese 1992) and (c) planning.

Model Building

A case manager bases her advice to a client on the case manager's estimates of the client's likely success on a given job-training path. Such estimates are based on characteristics of the client (such as literacy level, work preparedness, and time-management skills); the case manager's knowledge about available services (such as mental-health care and job-training classes) and regulations; the client's current record with the program; and the case manager's experience with this and other clients.

Our first task in building the decision-support system for WtW is to determine and formalize how the case managers make decisions. We model the WtW domain using dynamic Bayes nets. Some of the random variables in the WtW domain represent the client's participation and success in specific WtW activities and services. Other variables represent characteristics of the client, such as mental and physical health, social skills, job skills, and education level. In addition, we represent nonstochastic factors that affect client performance in the program (such as vehicle ownership and the number of children). We collect information on two levels: *qualitative*, for building an initial graph structure of the Bayesian network representing the domain, and *quantitative*, for the actual conditional-probability tables.

We have gathered qualitative information via in-person interviews with domain experts (case managers and other WtW personnel). We are developing knowledge-engineering tools for building models based on obtained information (Zhao *et al.* 2004). We can obtain quantitative information for the models both from domain experts and from analyzing existing longitudinal data. We built PET, a Probability Elicitation Tool (Zhao *et al.* 2004), for eliciting context-dependent probability tables from domain experts, and are working on tools to analyze and extract statistical information from existing data.

We store conditional-probability tables and associated context information obtained from experts and data in a special-purpose probabilistic database management system (Zhao, Dekhtyar, & Goldsmith 2004). All of our software that works with probabilities is interoperative with this database management system.

In addition to building the probabilistic model of the WtW domain, we also need to formalize the rules and regulations of the WtW program as well as the availability of services. We obtain this information from two major sources: WtW program documentation and in-person interviews with domain experts. We represent the information internally as a collection of constraints. We use *smodels* (Niemelä & Simons 2000), an answer-set formalism based on stable-model semantics, to model hard constraints. We are in the process of building a richer formal language than smodels for representing WtW program constraints.

Model refinement

The problem of finding optimal plans for factored stochastic domains is intractable (Mundhenk *et al.* 2000). Existing planning heuristics work well on small instances, but the generic domain model built in the first stage of the project may be large enough that general planning is intractable.

However, we don't need to compute plans for the full model. A good planner personalizes plans to the client. This observation leads us to investigate model refinement. Our goal is to construct smaller, *situation-specific* domain models on demand. A situation-

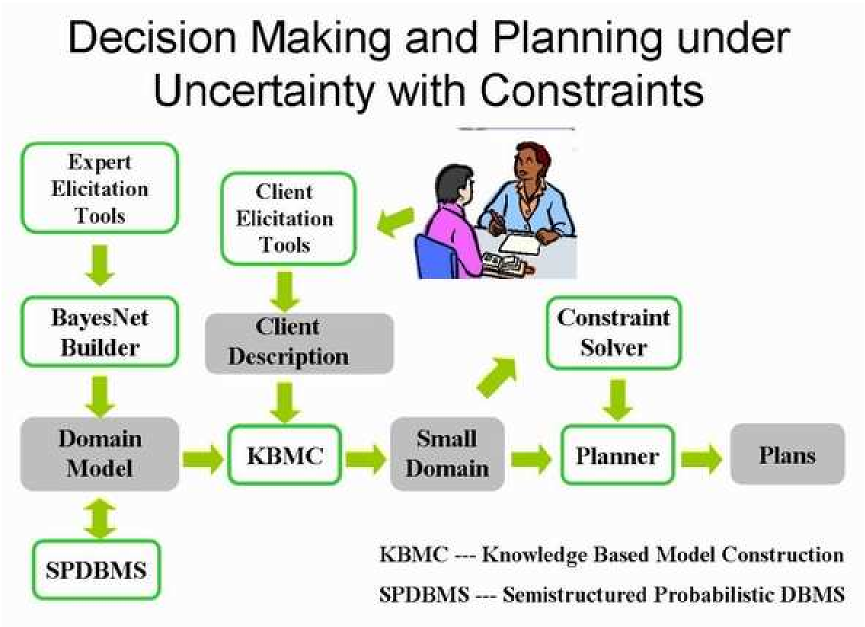

Figure 1: Architecture of decision-support system for WtW

specific model takes into account three types of information: service availability, client preferences and client characteristics. Service availability is highly dynamic: Shelters and classes fill up, classes are canceled or rescheduled. This information can be maintained in an easily-updateable database accessible to the case managers. Client preferences should be elicited during an advising session with the case manager. We are currently developing POET, a flexible Programmable Online Elicitation Tool for user preferences (Royalty *et al.* 2002; Williams *et al.* 2004). Client characteristics are elicited and observed by case managers and other professional assessors.

A process called Knowledge-Based Model Construction (KBMC) (Breese 1992) combines client preferences and other situation-specific information to construct a smaller situation-specific domain model. We are developing a toolkit of KBMC algorithms for domains represented with Bayes nets.

Planning

The situation-specific domain produced by KBMC software serves as the input to the WtW planner. The planner must integrate traditional decision-theoretic planning in uncertain domains with constraint satisfaction to make sure that the produced "good" plans are feasible, adhere to regulations, and satisfy the client's preferences to the best of the planner's ability.

Client preferences, as distinguished from client demands or needs, are represented as *soft constraints*. Combining soft constraints with smodel representations of hard constraints allows us to use constraint solvers to produce recommendations that satisfy all hard constraints and optimize the soft constraints. At present, we are able to generate one-step advice and are actively developing algorithms that combine constraint solvers and decision-theoretic planners.

Example

The software we propose to build will work together with case managers to suggest plans for WtW clients. We illustrate this process on the following example.[1]

In the month of October, a case manager (CM) receives a case of a 23 year-old woman, a mother of two children ages six and three. The client has a 10th-grade

[1]The example used here is fictional but realistic. It is based on one of the synthetic case histories created by Beth Goldstein and Cindy Isenhour in the process of elicitation of knowledge from the case managers.

education and no GED (General Educational Development, a high-school equivalent diploma), has no work history or specialized training and has occasional bouts of depression. (We factor the space of possible client states, with factors such as "education level", "number of children," and "personal hygiene".) The CM identifies needs for childcare and personal health care and a lack of resources for personal hygiene and clothes as employment barriers for the client.

The CM then advises the client that, to be eligible for the program benefits, she must fulfill a work requirement. The client expresses a preference for work in health care or child care. Based on this information, the CM determines her own actions and the actions she will advise the client to take. The CM's actions are: (i) refer the client to children's services for childcare; (ii) issue a one-time stipend for clothing and personal-hygiene items; (iii) refer the client to take a career-assessment test; and (iv) refer the client to a depression-counseling program.

The CM can advise a number of possible training options. Lack of GED suggests that the client start by enrolling in GED classes with the goal of completing a GED within 6 months. The client can also pursue a career in health care: Nurse Assistant, Certified Medical Assistant (CMA) or Registered Nurse (RN). Each comes with a different timeline and requirements. Other career paths are available, such as beautician, administrative assistant, and salesperson.

The advice the CM gives to the client is guided by her assessment of the client's strengths and preferences and by the constraints imposed by the rules of the program and current realities. Program constraints dictate that the client be enrolled in work-related activities and specify what those may be. Because the client is older than 20, her enrollment in GED classes counts for only half of the work requirement, leaving her responsible for another 10 hours of "work" per week. Current realities concern the availability of certain options: CMA classes at a local technical college start in August and in January, so the client must wait three months before enrolling and must meanwhile engage in another activity. All these constraints must be incorporated into the planning process, whether by a CM or a computer.

The CM's advice depends on her assessment of the client's likelihood of success in specific tasks. Without a high-school diploma, the RN career path is closed for the client. A high-school diploma also may increase the likelihood of success in study towards CMA and preparation for an administrative-assistant career. In our case, the CM's initial advice comes in the form of the following plan:

October – December: Volunteer full-time at a local hospital.

January – May: Take GED classes; volunteer part-time.

June – July: Take the GED examination; volunteer full-time.

August – December: *If* the GED is obtained, enroll in the CMA program at a local technical college. *Otherwise*, enroll in Nursing Assistant training, which does not require a GED or diploma.

The CM has implicitly performed KBMC by narrowing down the range of options with each new piece of information about the client. She eliminates administrative-assistant training and training at four-year colleges because these activities require a high-school diploma. She determines that the client has no interest in a career in sales and discards this option as well.

The proposed plan is subject to revision. Say the career-assessment test reveals that the client has poor reading skills, reducing the CM's assessment of the client's likelihood of success in GED study. The CMA program is not reading-intensive, so the CM believes the client has a higher chance to succeed if she foregoes the GED and starts CMA program directly in January. She changes her plan for the client accordingly. Perhaps in April she finds out from the instructor at the college that the client is likely to fail CMA training. Followup discussion with the client identifies that the client is afraid of becoming infected in the course of performing CMA duties. The plan now needs a drastic revision: The CM discards health-care related careers and revives previously discarded options of beautician and salesperson.

Towards Automation

The example from Section illustrates how we apply our proposed and actual software.

Before advising, we will have built a representation of client states with variables for factors including facts such as age, education level, etc., and assessables such as self-esteem, mental health, personal hygiene, wardrobe, etc. We will have a set of atomic actions, such as "take GED training," or "buy business attire", and a model of which factors each atomic action effects.

In addition, we will have probabilities associated with the effects of actions, based on relevant preconditions. For instance, the action "interview for office job" might have possible effect OFFICEJOB with preconditions (parents, in the dynamic Bayes net model) that include variables GED, SKILLS, HYGIENE, BUSINESSATTIRE, SELFESTEEM. With a low value for the variable SELFESTEEM, the action would have low probability of success. Thus, any plan that has a goal OFFICEJOB would have to address self esteem in any state describing a client with low self esteem. Probabilities will be elicited from WtW personnel using some variant of PET (Zhao *et al.* 2004), and will be extracted from data.

The first phase of advising is gathering information. Clients and CMs complete a vast array of overlapping forms about the client. We would like to build a central database of client information that can be accessed automatically to fill in those forms.

In addition to facts such as education level and children's ages and assessed values from tests, we will elicit client preferences, such as a preference for childcare or healthcare work. We will also elicit CM preferences about local agencies, good career options for clients based on the local economy, etc. A version of POET will be used for this elicitation.

The facts and assessments form the "state" of the client in our Bayes net representation of a Markov decision process. The client's and CM's preferences are used to compute a utility function for that client.

A KBMC algorithm is then used to eliminate unreachable states (such as the client being 19 years old, in the above example) and actions (such as attending college). This step makes use of constraint solvers. We plan to use Answer Set Programming (ASP) formalisms such as smodels (Niemelä & Simons 2000) and high-level constraint-representation languages that translate into ASP representations (Finkel, Marek, & Truszczynski 2004). Since not all assessments are immediately available, the KBMC tool makes use of network fragments (Laskey & Mahoney 1997) to allow conditional probabilities to be evaluated. We apply an information-fusion algorithm to define conditional probabilities for possible transitions based on whatever contextual information is available (such as the locations the client can reach for depression counseling).

We then apply a planning algorithm in combination with constraint solvers to generate and evaluate a plan that satisfies the hard constraints (if possible!) and maximizes the utility function within those constraints. Later, when the state or utility function changes, for instance when assessments are recorded or the client changes her preferences, we can update the plan. We are therefore investigating "tweakable" planning heuristics.

Challenges

There are two distinct sets of challenges for this project. One set consists of technical challenges in developing heuristics for highly intractable problems. The other set arises in social contexts, in developing and testing software (Goldsmith, Goldstein, & Mazur 2004), building software appropriate for the client base, and in obtaining approvals to work with highly at-risk populations.

CS Challenges

Planning in factored systems is intractable (Mundhenk *et al.* 2000), as is constraint solving. Handling vast quantities of probability data is a new area, although one in which this group is a leader (Zhao, Dekhtyar, & Goldsmith 2004).

The process of collecting data from diverse sources about changes in regulations and services is both a social and a technical challenge. Turning that information into formal constraints is more of a technical challenge.

The database of conditional-probability distributions for the domain may contain distributions for the same nodes under different conditions, or contexts. However, data will not necessarily be available for all contexts. When situation-specific networks are built, the key problem is to determine the most appropriate distributions, and the manner in which to combine them. This calculation is known as probability fusion.

Human-subject Challenges

Human-subject and ethical challenges arise throughout the processes of building models and deploying decision-support software. For example, domain experts do not think in mathematical formalisms; we must convert their knowledge into our probabilistic framework. Many Welfare to Work clients have low literacy skills; we must develop human-computer interfaces that take these limitations into account. Other challenges arise from privacy considerations in eliciting and storing personal information from clients; we must satisfy Institutional Review Board protection of participants in our testing and development.

Any decision-support software runs the risk of being used to replace human decision makers. This risk is a danger both to the decision-makers' employment and to the quality of service they provide. No matter how good the developers' intentions are, software is just a tool and may be used in unintended ways. No software will be able to come up with the outside-the-box solutions that good human advisors can produce. Our intention is to provide decision support, rather than decision making.

Acknowledgments

We acknowledge the contributions of the students working on this project, particularly Peg Wimmer. This work was partially supported by NSF grant ITR-0325063. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency, or any social welfare agencies.

References

Breese, J. S. 1992. Construction of belief and decision networks. *Computational Intelligence* 8(4):624–648.

Finkel, R.; Marek, W.; and Truszczynski, M. 2004. Constraint Lingo: Towards high-level constraint programming. *Software — Practice and Experience* 34(15):1481–1504.

Goldsmith, J.; Goldstein, B.; and Mazur, J. 2004. Social construction of technology in the welfare to work project. In *Proc. Workshop on Bayesian Applications at UAI '04*.

Laskey, K. B., and Mahoney, S. M. 1997. Network fragments: Representing knowledge for constructing probabilistic models. In *Proc. UAI '97*, 334–341.

Mundhenk, M.; Goldsmith, J.; Lusena, C.; and Allender, E. 2000. Complexity results for finite-horizon Markov decision process problems. *Journal of the ACM* 47(4):681–720.

Niemelä, I., and Simons, P. 2000. Extending the smodels system with cardinality and weight constraints. In Minker, J., ed., *Logic-Based Artificial Intelligence*. Kluwer Academic Publishers. 491–521.

Royalty, J.; Holland, R.; Goldsmith, J.; and Dekhtyar, A. 2002. POET, the online preference elicitation tool. In *Proc. AAAI Workshop on Preferences in AI and CP: A Symbolic Approach*.

Williams, D.; Dekhtyar, A.; Finkel, R.; and Goldsmith, J. 2004. Towards reliable preference elicitation. Manuscript.

Zhao, W.; Dekhtyar, A.; Goldsmith, J.; Jessup, E.; and Li, J. 2004. Building Bayes nets with semistructured probabilistic DBMS. *GI-EMISA Forum* 24(1):29–30.

Zhao, W.; Dekhtyar, A.; and Goldsmith, J. 2004. A framework for management of semistructured probabilistic data. *Journal of Intelligent Information Systems*. Accepted.