

# Planning for success: The social approach to building Bayesian models

Krol Kevin Mathias    Cynthia Isenhour    Alex Dekhtyar  
Judy Goldsmith        Beth Goldstein

March 14, 2007

## Abstract

We introduce a new variant of Markov decision processes called MDPs with action results, and a variant of dynamic Bayesian networks called *bowties*, for modeling the effects of stochastic actions. Bowties grew out of our work on decision-support systems for advisors in the US social welfare system. Bowties, and our elicitation process for them, are designed to elicit dynamic Bayesian network fragments from domain experts who think narratively instead of quantitatively. Our elicitation process has worked well with the welfare case managers and other domain experts, in the sense of capturing consistent and validated models.

## 1 Introduction

Planning under uncertainty with constraints is a metaphor for our times and a computational problem of interest in AI, operations research, and decision science. Within AI, planning under uncertainty is most commonly modelled using Markov decision processes (MDPs). There are well-understood solvers, hereby called *planners*, for MDPs. A major issue in using MDPs to solve real-world problems is the problem of building appropriate mathematical models of the phenomena of interest.

This paper addresses the problem of acquiring models from domain experts. We concentrate on the model-building process for the domain of the Kentucky “welfare-to-work” social welfare system. Those who guide welfare recipients (clients) through the Byzantine maze of requirements and constraints and negotiate plans with and for the clients are called *case managers*. Our domain experts have primarily been case managers, with a few welfare agency managers.

We show how the synthesis of expertise from welfare case managers, social scientists, and artificial intelligence practitioners led to a new “shape” of dynamic Bayesian network (henceforth, dynamic Bayes net). This shape, which we call a *bowtie*, captures the knowledge of the welfare-to-work system case managers about predicting client success and the understanding of the social scientists about how best to elicit that knowledge. From a bowtie model, one can compute a consistent two-phased temporal Bayes net (Boutilier, Dean, & Hanks, 1999) or some other planning-appropriate format.

The main contributions of the paper are twofold: we present a model of success-centric planning with uncertainty, and we show how qualitative research with domain experts informed the development of that model.

We are already using bowties and their associated elicitation process in other planning domains, and foresee wide application. We also expect that careful ethnographic studies of other domains will lead to further innovations in modelling and planning processes. This paper can serve as a roadmap for such studies.

## 2 Technical Definitions

Traditional AI planning assumes that actions have known, deterministic effects. However, the real world often surprises us. We may set out to buy milk, and come home with a new puppy. We might advise someone to take a class, with only statistical knowledge of how they are likely to do in that class. In particular, in the welfare system, a case manager may recommend a particular set of actions for a client to take. The case manager does not know for certain that the client will be compliant with the plan that the case manager and client agree upon. Furthermore, even if the client attempts classes, for instance, success is not predetermined.

We begin by defining standard decision-theoretic models, namely Markov decision processes and dynamic Bayes nets. We then give the formal definition of bowtie models.

The key elements of a model of decision-making under uncertainty are system states, actions and their effects, and a set of goals or utilities to guide the decision-making. A *Markov Decision Process (MDP)*  $M = \langle S, A, f, r \rangle$  consists of

- a set of states,  $S$ , (in our context, a finite set),
- a finite set of actions,  $A$ ,

- a function  $f : S \times A \times S$  that assigns probabilities to possible action outcomes ( $f(s, a, s')$  represents the probability of outcome  $s'$  when action  $a$  is taken in state  $s$ ), and
- a utility function  $r : \rightarrow R$ .

In many situations, it is easiest to represent states of an MDP in a factored form, in terms of values of *state variables*. In the welfare system, the state describes a welfare recipient in terms of his or her age, income, number of children, age of the youngest child, etc. In a medical domain, state variables might reflect the presence of various symptoms or pathogens.

Formally, a *factored MDP* has a set  $V = \{x_1, \dots, x_n\}$  of state variables with finite domains,  $dom(x_i) = \{v_1^i, \dots, v_{k_i}^i\}$ . Thus, the number of possible states is  $\prod_{j \leq n} k_j$ .

The advantage of a factored representation is that transition probabilities may be defined on each outcome variable independently. The most common representation of such probabilities is in terms of *dynamic Bayes nets*, more particularly, *two-phase temporal Bayes nets (2TBNs)*. A 2TBN consists of a directed acyclic graph with  $2n$  nodes, usually written in two columns of  $n$  nodes. (See Figure 1.) The nodes in the left column, labeled  $x_1, \dots, x_n$ , represent the state at time  $t$ . Those in the right column, labeled  $x'_1 \dots x'_n$ , represent the state at time  $t + 1$ . Edges represent dependencies: An edge from  $x_k$  to  $x'_l$  indicates that the probabilities assigned to possible values of  $x'_l$  depend on the value of  $x_k$ . Probabilities are specified by *conditional probability tables (CPTs)*.

As we argue in the rest of the paper, there is another, more intuitive representation of actions in a factored MDP. Our new representation includes an explicit state variable for the success of an action. This success node is intermediate between the predictors for that action, namely, the nodes  $x_i$  in Figure 1, and the effects, namely, the nodes  $x'_i$ . Thus, we have a *three-phase temporal BN*, where the central phase is the success node, as shown in Figure 2.

We define *MDPs with results* as factored MDPs, where each action has a specific state variable representing its outcome, or success. Actions in MDPs with results can, by definition, be represented by bowties.

The introduction of bowties was a natural outgrowth of the qualitative work on welfare case managers' decision-making. This work is described in the sections that follow. In Section 5 we show that bowties are not only natural, but have significant advantages over standard 2TBNs for modeling certain types of domain. In Section 6, we discuss the formal details of the bowties.

A *policy* for a stochastic system must specify actions to be taken at every state that may occur. The quality of a policy is measured as the expected accrued utility over time, or the probability that the agent, following the policy, will reach a goal state. We do not discuss MDP planning algorithms in this paper. A rich literature on that subject is available elsewhere. We also do not discuss here the combination of bowtie fragments to represent outcomes of simultaneous actions, although that is a natural part of building decision-support software for this application.

### 3 Introduction to the Welfare to Work Domain

In 1996 the US legislature signed the Personal Responsibility and Work Reconciliation Act (PRWORA). This act was built upon the popular assumption that welfare programs had enabled apathy, dependence, a poor work ethic, and abuse of the system among the nation's poor<sup>1</sup>. To alleviate what was perceived to be a substandard work ethic among the poor, PRWORA set a 5 year lifetime limit on welfare benefits for all recipients and mandated that welfare clients work or participate in work readiness, education, or training programs<sup>2</sup> in order to receive benefits<sup>3</sup>.

PRWORA also transferred the responsibility for welfare programs to states and shifted the emphasis of support away from government sponsored cash grants and towards the provision of support services by non-governmental organizations. As of 2004, more than one half of all TANF dollars were allocated to support services, compared to the previous welfare system (AFDC) which allocated more than 80% of all dollars to direct cash assistance (Allard, 2004).

These changes in welfare legislation significantly restructured the work of welfare (Morgen, 2001). Case managers who were once responsible for determining eligibility and processing cash assistance payments (by means of pre-existing formulas, were suddenly held accountable for informing clients about work and work-related program requirements, assisting clients in the discovery and/or definition of career goals, and helping individuals to match their interests, preferences, abilities, and goals to a long list of "countable"

---

<sup>1</sup>Supporters of PRWORA thus sought to explain poverty in terms of dysfunctional citizens, rather than in terms of economic restructuring, economic inequality, or persistent gender- and race-based inequalities.

<sup>2</sup>There are exceptions to these rules for individuals in extreme circumstances.

<sup>3</sup>Note that unpaid domestic work is not considered "work" by most states under this legislation.

activities<sup>4</sup>.

Case managers in Central Kentucky, where our research program is based, are responsible for ensuring that their clients participate in one or more of these mandated activities. In order to do so they must make complex decisions about where to refer their clients. This requires that they process information about myriad training, support, employment and educational programs available to their clients including information about prerequisites, schedules, locations and content<sup>5</sup>. In the central Kentucky city of Lexington alone, more than 200 agencies offer support services to welfare clients. Case managers develop action plans that fit client needs for these services and suggest agencies that are appropriate in terms of location, schedule, etc.

Case managers must thus be familiar with the goals, preferences, abilities, constraints and interests of their clients. To further complicate matters, the case managers participating in the research handle between 40-80 active cases. Finally, case managers must also stay abreast of changes in welfare regulations, policies and rules as shifting budget allocations dictate changes in policy execution, as well as service availability.

Decisions about where to refer clients are thus made in situations of relative uncertainty since case managers have few means and little time in which to assess the client's commitment to a goal or to accurately predict the client's ability to achieve success. With limited time to spend and precarious relations of trust with each client, case managers are often forced to operate without complete information about the client's personal life and associated constraints, often leading to client categorizations (Lipsky, 1980; Kingfisher, 1996, 2001) which, although time saving, may not serve clients well.

## 4 Planning for WtW

The increased burden on the case managers can potentially be reduced by introducing decision support/planning software into their work. Software support for case managers can take a variety of forms. In our work we concentrate on one of the most important aspects of the new responsibilities

---

<sup>4</sup>Countable activities are those which can be counted towards the fulfillment of federal participation quotas.

<sup>5</sup>Our research with case managers has been ongoing for nearly three years and includes several rounds of iterative semi-structured interviews, questionnaires, free lists, rankings, pile sorts, group interviews, electronic data elicitation and participant observation methodologies. Twenty eight case managers from three different agencies have participated in this research.

of case managers: provision of advice for their clients. In this section, we present a formalization of the welfare-to-work domain for the purpose of planning/advising.

As mentioned in Section 3, a key feature of the TANF (Temporary Aid for Needy Families) system is the focus on providing services to clients in order to make the clients employable and employed. These services can be roughly partitioned into two categories. Services in the first category are performed to *alleviate barriers* a client might have, preventing her from participation in the rest of the program. Such services may include subsidized housing, health care, child care, transportation allowances, help with basic coping strategies, etc. The services in the second category are the activities the client can be involved in, in order to (a) remain eligible for the aid and (b) become employable and employed. Such activities include volunteering opportunities, literacy training, high school equivalency or college classes, professional training, English as a second language classes, job search and interview preparation seminars and more. Case managers are entrusted with advising their clients on activities, which, in their opinion, advance the client towards the general goal of employability/employment.

We chose to model the advice provided by case managers to their clients as a *problem of planning in a stochastic domain with constraints*. Each time a case manager and client meet, they negotiate a long-term plan for the client with immediate actions, based on success so far, the client's stated goals and preferences, the client's current situation, the goals of TANF, and the case manager's assessment of the client's abilities, and likelihood of success. These contracts are renegotiated each time because goals, preferences, assessments, situations, and available services are all subject to change. Participation in a service might or might not lead to its successful completion; Success or failure will each affect the client's state in a number of ways, potentially both positive and negative.

We model current situations as *factored states*, services as *stochastic actions*, preferences as *utility functions* over possible states and actions, and regulations and limitations of clients as *constraints*. In stochastic planning, a *policy* specifies actions for all possible outcomes or states.

More specifically, we consider the factored MDP states to be formed by a number of *client characteristics*. The work described below has identified a wide range of characteristics, which include such objective attributes as client's age, education level, number of children or disability status, as well as somewhat more subjective ones, such as, client's literacy and numeracy, self-confidence or commitment level. One action in the welfare domain is client's participation in one of the services/programs, such as GED classes,

volunteering or job interview skills seminars. Such participation may affect some of the client's characteristics. For example, a job interview skills seminar affects the client's self-confidence (which can go either up or down, depending on whether or not the client feels she is ready for the rigorous job interviewing process), commitment and work-readiness.

To evaluate suggested actions, and their results, we need to determine the source of utility in the domain. Utility elicitation presents different issues in model building. Utilities in the welfare system are specific to the individual client and assigned case manager, and most be elicited during the advising process. We do not address utility elicitation in this paper.

We illustrate the case manager-client interaction on the following fictional case.

**Example 1** *A 21-year-old woman with a 4-year-old son and a 2-year-old daughter has completed 11th grade, lives in a government-subsidized apartment complex, and has been unable to seek work. The barriers to her participation in services are her lack of childcare and lack of transportation—her apartment building is not on a bus route and she doesn't have a car.*

*The case manager first addresses these barriers by providing a transportation stipend so that the client can hire transport to an approved childcare site and to an adult education center, where she will prepare for her highschool equivalency exam. The long-term goal is a clerical job, with midterm goal of enrolling the client in secretarial school. There are two options for secretarial school. One offers evening courses, which are incompatible with childcare availability. Thus, constraints dictate that she attend the other school.*

*Another option would be to send this client immediately to car mechanic training. While the training is available and convenient, and this could lead to a high-paying job, the client is unwilling to deal with the prejudice against women she expects to find in the automobile repair world. The case manager determines, therefore, that this option has a significantly lower probability of success. She chooses not to pursue this option for this client.*

In order to compare options such as secretarial school and car mechanic training, the case manager must assess the probabilities of each action's success, given the client's state, and the probable effects of both success and failure at each action. As mentioned above, we model these using dynamic Bayes nets.

There are three basic steps to building a Bayes net model for any application: Determining the key components of the domain, translating them into the components—variables, actions, dependencies, and probabilities—of the mathematical formalism, and verifying the models.

In order to accomplish the first, the social scientists in our group used open-ended interviewing techniques with welfare professionals. The second part can be grouped into three steps: (a) determining the variables and actions; (b) eliciting qualitative relationships amongst these components, and (c) determining quantitative relationships that are consistent with the elicited information.

The next few sections focus on the processes of collecting key components of the model, on eliciting dependencies, and on how the initial elicitation designed by the social scientists affected the way in which qualitative information was elicited and represented.

## 5 From Case Manager Interviews to Formal Models

In this section, we introduce the voices of the case managers and some discussion about their reasoning. We argue that traditional Bayesian elicitation modes are not well suited for model-building in this domain, and present a qualitative method that captures crucial insights from the case managers' description of their decision-making process. In Section 6, we give the formal model that grew out of the case manager interviews, and in Section 7, we describe the elicitation process we developed to model the personal elicitation designed by our social scientists.

In the welfare domain, gathering empirical data about actual cases is extremely sensitive. Requests for personal information about the nation's most vulnerable populations are not taken lightly. Issues surrounding confidentiality and privacy require the informed consent of all welfare participants before their case records are released for research. Because our research team has not yet accessed empirical case data for an adequate sample of welfare recipients, we have our data collection efforts on the elicitation of expert knowledge. We rely upon the experiential knowledge of welfare case managers, those women and men who inhabit the welfare system on a daily basis, actively making decisions about whether to, for example, refer a client to substance abuse counselling at the expense of participation rates<sup>6</sup>; recommend beauty school or an administrative assistant's training program;

---

<sup>6</sup>Case manager job performance is evaluated in part on participation rates. Because state-level and local funding are dependent on keeping a certain percentage of all cases in federally mandated "countable" activities, many case managers are hesitant to allow clients to participate in activities which are "allowable" but do not count towards quotas such as counselling.

allocate support funding to a client; or cut another's benefits for not keeping scheduled appointments.

Eliciting data from welfare case managers presents a unique set of challenges. While decision modelling requires quantifiable data, the case managers with whom we work often vehemently resist our attempts to gather generalized or abstracted data. These women and men consistently insist that it is difficult to generalize about their clients because each one is different, making it equally difficult to generalize about decision making patterns. In one attempt to elicit information about a specific activity (taking GED preparation classes), one case manager expressed her reservations with our efforts. She said,

*“I think its really difficult to think about these issues individually . . . it has to be much more holistic. I mean if you look at my list, I've got everything ranked as extremely important. Everything is extremely important and I don't think you can just rank the top five . . . In my assessments, I'm not going to just ask (my clients) for five pieces of information. It seems impossible to isolate these factors or to categorize them.” (6/15/2005)*

Of the twenty-eight case managers participating in this project, nearly half of them have five or more years experience in this capacity. They have been trained on the job to look at their clients as individuals rather than numbers.

In sharing information about their clients and case management decisions, case managers prefer to speak in narratives, imparting tacit knowledge through stories of specific clients and their unique circumstances. These perspectives and preferences had a significant impact on our data elicitation methods.

In order to gather data in a way that made sense to the case managers, we not only had to be very clear about the requirements of data modelling, stressing the need for simplification of a clearly complex decision environment, but we also had to frame questions, statements and problems in meaningful language for the case managers. This often involved asking them to consider a specific scenario consistent with cases they've worked and familiar in terms of their goals as case managers.

## **5.1 CM model of thinking : Why Bowties?**

Case managers argue that small successes have a significant impact on the client's ability to succeed in their short term and long term goals, building

commitment, self-confidence, aptitude and aspirations to levels critical for achieving self-sufficiency and securing work that pays a living wage. So, while case managers may be evaluated by state and federal bureaucracies based on their ability to keep clients in compliance with participation quotas, experienced case managers know that client success in the long term is ultimately dependent on short term success. To summarize with a cliché: Nothing succeeds like success.

Through discussions between the social scientists and computer scientists,<sup>7</sup> we have determined that we need to elicit information from the case managers in terms of the success or failure of actions. From this understanding, we then determined that the best formalization of a success-based model is the MDP with results, also called *bowtie fragments* described in Section 2

We illustrate the synthesis of case managers' knowledge and the social scientists' knowledge elicitation with a description of the pilot elicitation.

## 5.2 Elicitation of the First Bowtie

Through a series of iterative interviews, our team of social scientists, computer scientists and domain experts were able to produce a list of the client characteristics which play into a case manager's decision regarding what activity to recommend to her client (including client interests, goals, aptitude, etc.). Before abstraction and consolidation, this list of client characteristics numbered more than 150 traits with multiple values. In order to build Bayes Nets, however, it was essential that we understand which of these variables were most determinative given a set of actions in which welfare clients can participate to fulfill their work requirement.

We began our elicitation process with a pilot elicitation. We utilized a participatory group interview focused on one action fragment: "get GED". By focusing on one action fragment the research team believed we could more accurately determine the effectiveness of the method and could explore the relationship between client variables and perceptions of the likelihood of success in this action. The method was developed to gather three pieces of information: (1) client attributes of state relevant to the case manager's decision to recommend the action "get GED" via free list, (2) The relative importance of each attribute of state listed in predicting success in a GED program via Likert scales, and (3) The optimal value of each attribute of state for predicting success in a GED program via small focus groups.

---

<sup>7</sup>We are indebted to Russell Almond for translation between the two groups at the key moment.

Twenty case managers from three different agencies were first asked to independently free list the “information you need to know about your client in order to assess whether or not getting a GED is a reasonable short term goal.” We found that some case managers had significant problems with this wording, telling us that it is not part of their job to determine whether a goal is reasonable, but only to support clients in their goals. One of the most vocal case managers remarked,

*“it’s not my concern what the process is of them getting it [a GED]. That’s what the adult education workers at the GED center deals with. That’s their job. My goal is just to allow them that opportunity to get their GED done and then to tell them what else they have to do for case management to be in cooperation but I’m not going to stop them, I’m not going to determine if it is an appropriate short term goal. That’s for the client to choose, that’s for the adult education worker to determine if this is going to take three years or six months. That should not... that will not affect my case management at all.”*  
6/15/2005

In light of this resistance we rephrased our question in terms more meaningful for case managers, terms associated with the client’s chances for success. The question was rephrased as such, “what information do you need to know about your client in order to determine her/his likelihood of success in a given action?” Case managers found no difficulty responding to this question since they focus their decisions on the potential for success in a given action and place heavy emphasis on client determination as they move towards action plans designed to achieve success.

Next, case managers were asked to create a collective list, each contributing items from their personal lists to a collective catalog of client attributes. In all, the case managers free listed 37 client variables that might affect a client’s ability to succeed in a GED program. Then each case manager was asked to augment her personal list with any additional characteristics listed by the group before rating her own list on a Likert scale from extremely important to not very important. Finally, the group was broken into three smaller focus groups of 6-7 to talk about the five most important variables for determining likelihood of success in a GED program.

In order to compile the results from the pilot, the research team aggregated individual responses by number of mentions and relative importance. The five characteristics considered most important in determining outcomes

in a GED program via this method of aggregation were: (1) learning disabilities, (2) last grade completed, (3) access to childcare, (4) age of the client, and (5) client’s goals.

The focus groups helped to consolidate the client attributes into composite categories. They also reminded us that we could assume that barriers such as childcare and transportation would be immediately addressed by the agency. Finally, the focus group conversations largely validated the aggregated individual results. All three focus groups listed (1) learning disabilities; (2) educational history (in which they included highest grade completed, reading level, reason for dropping out of school); (3) and motivation (also a composite category which includes motivation, commitment, goals and resolve) in their five most important characteristics.

## 6 Bowtie Model

We have given an informal definition of the bowtie model in Section 2, and discussed the evolution of the model. In this section, we give the formal definitions. We note that the model presented for purposes of elicitation is, in fact, bowtie shaped. However, there are other, implicit dependencies between nodes that obscure the bowtie shape. In particular, when we say that, say, success at the GED exam affects self-confidence, what we mean is that *changes to an output node are based on success and on the previous value of that variable*.

Let  $D = \langle d_1, \dots, d_m \rangle$  be a list of random variables. Each variable  $d_i$  takes values from a finite ordinal domain.

**Definition 1** A bowtie model  $\mathcal{B}_A$  for an action  $A$  is a tuple  $\mathcal{B}_A = \langle \mathcal{V} \cup \mathcal{O} \cup \{s_A\}, E_{\mathcal{V}} \cup E_{\mathcal{O}}, w_i, w_o \rangle$ , where

- $\mathcal{V} \subset D$  is the set of input nodes;
- $\mathcal{O} \subset D$  is the set of output nodes;
- $s_A$  is the success node of action  $A$ ;
- $\mathcal{V}$ ,  $\mathcal{O}$  and  $s_A$  together form the set of nodes for the bowtie model;
- $E_{\mathcal{V}} = \{(x, s_A) | x \in \mathcal{V}\}$  and  $E_{\mathcal{O}} = \{(s_A, y) | y \in \mathcal{O}\}$  are the edges of the model;
- $w_i : E_{\mathcal{V}} \rightarrow \mathcal{N}$  is the input weight function;

- $w_o : E_{\mathcal{O}} \rightarrow \mathcal{I}$  is the output weight function.

The output weight function describes the direction and magnitude of the change in values of outcome nodes given the success node indicates success. (One can assume that the affects of failure have the same magnitude but the opposite direction, or one can elicit those effects separately. If the success node has multiple values, the effects will probably have to be elicited separately for each value.)

In the bowtie model, the following assumptions hold.

1. The domain for each  $d_i$  is ordinal;
2. there is causal independence between input nodes ( $V_i$ ), where  $V_i \in \mathcal{V}$ ;
3. the input weight function is always positive because of our choice of ordering of values of the input node values.

The *bowtie* model explicitly represents the outcome of an action as a central node (knot) of the bowtie.

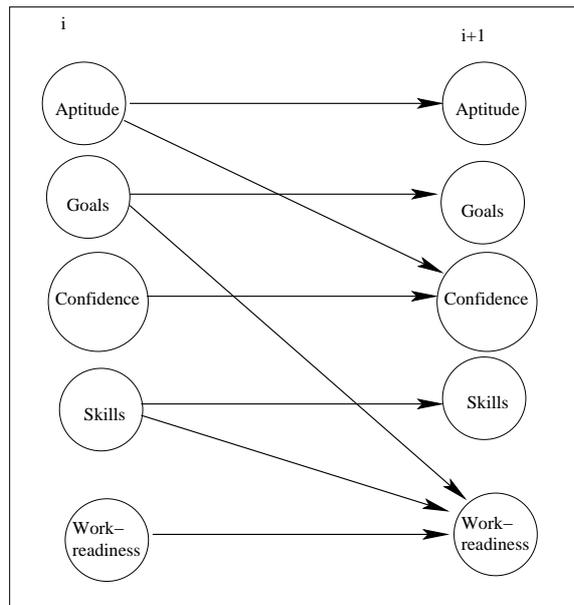


Figure 1: DBN for the action fragment “Volunteer Placement (VOP)”.

Figure 1 shows an example of a dynamic Bayesian network representing the action fragment “Volunteer Placement (VOP)”. VOP is an action from

the WtW domain, in which the client participates in volunteer work relevant to the client’s goal of getting a job.

A DBN is commonly represented as a *2-phase temporal Bayes net (2TBN)*. In the 2TBN (Figure 1), we have client attributes {APTITUDE, GOALS, CONFIDENCE, SKILLS, WORK-READINESS}, represented at times  $i$  and  $i + 1$  respectively. An edge from the attribute APTITUDE to the attribute CONFIDENCE represents the fact that the value of attribute CONFIDENCE at time  $i + 1$  can be predicted stochastically from the value of APTITUDE at time  $i$ , as well as from the value of its other parent, CONFIDENCE.

The 2TBN is a sufficient representation for an action that does not have effects that depend on the *outcome* of the action. But as we have seen in Section 5.2, in WtW, knowing the outcome of an action is extremely important in determining the expected changes in the client’s state.

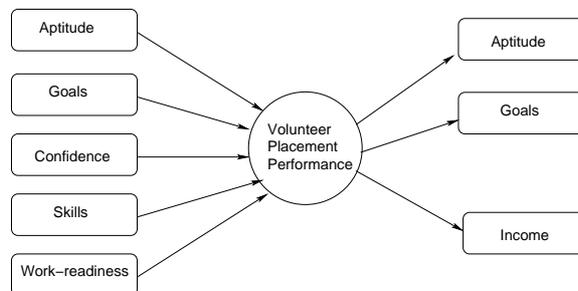


Figure 2: The elicited bowtie model for the action fragment “Volunteer Placement (VOP)”.

The bowtie model makes a convenient picture for elicitation of the action fragments in the WtW domain. As we can see in Figure 2, the bowtie model shows that the client’s aptitude, goals, confidence, skills, and work-readiness levels determine her success/failure in the action VOP; and based on the outcome of the action, the client’s aptitude, goals and income levels change stochastically. During the elicitation process for bowtie models, the experts were asked to elicit  $m$  input nodes and  $n$  output nodes. Thus, the experts were asked on the order of  $(m + n)$  questions; whereas in the case of the 2TBN elicitation, the experts are typically asked questions on the order of  $n^2$ .

Although the bowtie model in Figure 2 makes a convenient picture for elicitation, it does not represent the actual bayesian fragment. For example, from Figure 2 we only know that the client’s aptitude, goals and income

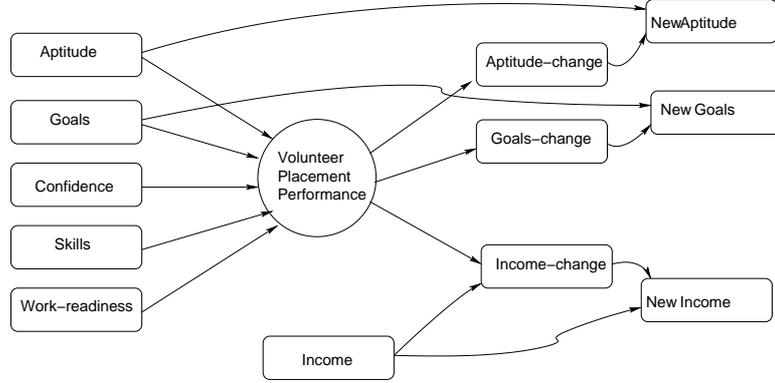


Figure 3: The actual bowtie model for the action fragment “Volunteer Placement (VOP)”.

levels **change** stochastically based on the outcome of the action VOP. Thus the elicited bowtie model (Figure 2) only represents what changes by the output nodes. This information is not sufficient to determine the final values of the output attributes. In order to do this, we introduce nodes that represent the magnitude of change in attributes and include implicit dependencies between the output attributes and their respective priors (Figure 3).

In the actual bowtie model, *each output node of an action fragment is influenced by two nodes: the node representing the change in its values and the node representing the value of this attribute prior to the action being taken.* The node representing the change in the characteristic is stochastically influenced by the outcome of the action.

Let  $D = \langle d_1, \dots, d_m \rangle$  be a list of random variables. Each variable  $d_i$  takes values from a finite ordinal domain.

**Definition 2** *An actual bowtie model  $\mathcal{AB}_A$  for an action  $A$  is a tuple  $\mathcal{AB}_A = \langle \mathcal{V} \cup \delta_{\mathcal{O}} \cup \mathcal{O} \cup \{s_A\}, E_{\mathcal{V}} \cup E_{\delta_{\mathcal{O}}} \cup E_{\mathcal{O}} \cup E', w_i, w_o \rangle$ , where*

- $\mathcal{V} \subset D$  is the set of input nodes;
- $\mathcal{O} \subset D$  is the set of output node;
- $\delta_{\mathcal{O}}$  is the set of nodes representing the change in output nodes;
- $s_A$  is the success node of action  $A$ , and  $\mathcal{V}$ ,  $\mathcal{O}$ ,  $\delta_{\mathcal{O}}$  and  $s_A$  together form the set of nodes for the bowtie model;

- $E_{\mathcal{V}} = \{(x, s_A) | x \in \mathcal{V}\}$ ,  $E_{\delta_{\mathcal{O}}} = \{(s_A, y) | y \in \delta_{\mathcal{O}}\}$ ,  $E_{\mathcal{O}} = \{(y, z) | y \in \delta_{\mathcal{O}}, z \in \mathcal{O}\}$  and  $E' = \{(x, x') | x' \in \mathcal{O}, x \text{ is the prior}^8\}$  are the edges of the model;
- $w_i : E_{\mathcal{V}} \rightarrow \mathcal{N}$  is the input weight function;
- $w_o : E_{\delta_{\mathcal{O}}} \rightarrow \mathcal{I}$  is the output weight function.

For example, in Figure 3, the client’s aptitude level after participation in action VOP is determined by the change in the client’s aptitude level, which is based on the outcome of the action VOP, and by the client’s aptitude level before participating in the action VOP.

In the two phase model, there would have to be either an assumption that the action always succeeds, or synchronous arcs to link the effects of the action, in order to model the case managers’ thinking. Synchronous arcs raise the complexity of computing with the model significantly, and ultimately lower the comprehensibility of the outcomes.

Bowties, by explicitly representing the outcome of an action, compactly represent the effects of that outcome, and explicitly link the effects to the outcome. They do so without adding synchronous arcs. The bowtie fragments offer the opportunity to develop new algorithms and data structures to improve planning and inference under uncertainty.

## 7 One HELL of an Elicitation

In this section, we describe the software used for elicitation. When we introduced software to automate bowtie elicitation, one colleague suggested calling it the High Level Elicitor. Another pronounced that, “HELL”. While many of us were sure this would offend the case managers, it turned out to amuse them. The name stuck.

### 7.1 The Software

The research team designed a software tool to replicate the pilot elicitation process for eliciting bowtie models for other actions in which welfare clients can participate.

In order to design the tool, we needed answers to the following questions.

1. What actions are recommended to clients in the WtW system?

---

<sup>8</sup> $x$  represents the random variable prior to taking the action  $A$  and  $x'$  the same variable at the next time step, after taking action  $A$

| Name                 | Abbrevn. | Name                    | Abbrevn. |
|----------------------|----------|-------------------------|----------|
| Job & Post-Sec. Edu. | JSE      | English Second Language | ESL      |
| Short Term Training  | JST      | Volunteer Placement     | VOP      |
| Voc. Rehabilitation  | RHB      | Job Readiness Class     | JRA      |
| On Job Training      | OJT      | Literacy Training       | LIT      |
| College              | COL      | Community Service       | COM      |
| Adult Basic Edu.     | ABE      | Community College       | CCO      |
| High School          | HSC      | Voc. School             | VOC      |
| Group Job Search     | GJS      |                         |          |

Figure 4: The list of 15 elicited actions.

2. What client attributes are considered by case managers before recommending an action?
3. What client attributes change based on the outcome of an action?

The social scientists interviewed case managers to answer the first question and to produce a list of client attributes relevant to the actions. We ended up with 16 actions<sup>9</sup> (Figure 4) and around 32 client attributes (Figure 5). The list of relevant client attributes and actions helped us determine the domain for client states and the set of actions to be used for planning. In addition to the steps followed in the pilot elicitation, the HELL tool was designed to elicit output nodes.

With the help of the pilot elicitation described in Section 5.2, we were able to identify an elicitation model for action fragments that aligned with the case managers' mental model. The case managers did not see a probabilistic network and were not necessarily conversant with probability theory. With the help of the social scientists, the computer scientists realized that we needed to construct "user-centered" representations in order to elicit the right information from the case managers in an efficient and convenient manner.

<sup>9</sup>The action 'Taking GED preparation classes' was elicited during the pilot experiment.

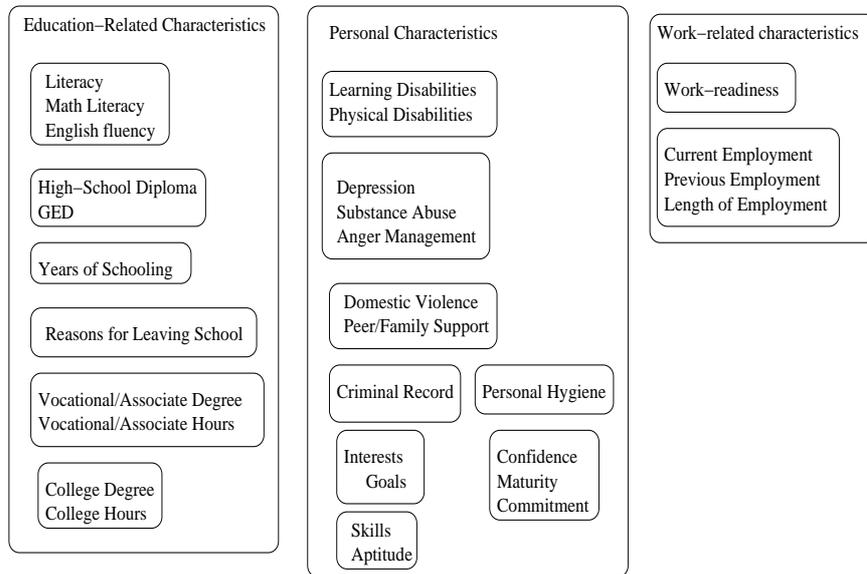


Figure 5: Client characteristics used in the experiments.

In order to make the elicitation process efficient, we decided to build a software tool, HELL, that represented the mental work flow of the case managers. The greatest challenge was to come up with the wording of questions for each stage of the elicitation process. We wanted to avoid open-ended questions since answers to those would increase the complexity in data analysis. We needed to control what data was elicited, but the case managers needed some flexibility in the elicitation process. We had to design a navigation pattern that was intuitive and also explained the user's position in the elicitation process.

We resolved usability issues in the design of the tool by applying important heuristics from human-computer interaction theory and human work flow theory. In order to make a usable elicitation tool, we focused on using the case managers' language (Mack & Nielsen, 1993) and modelling their workflow (Mayhew, 1992). The tool had to provide a step wise elicitation interface that was consistent with the case managers mental approach of recommending an action.

Since probabilistic language has little overlap with simple natural language, we did not want to elicit probabilities from case managers. Instead, we elicited qualitative information by using Likert scales for eliciting the

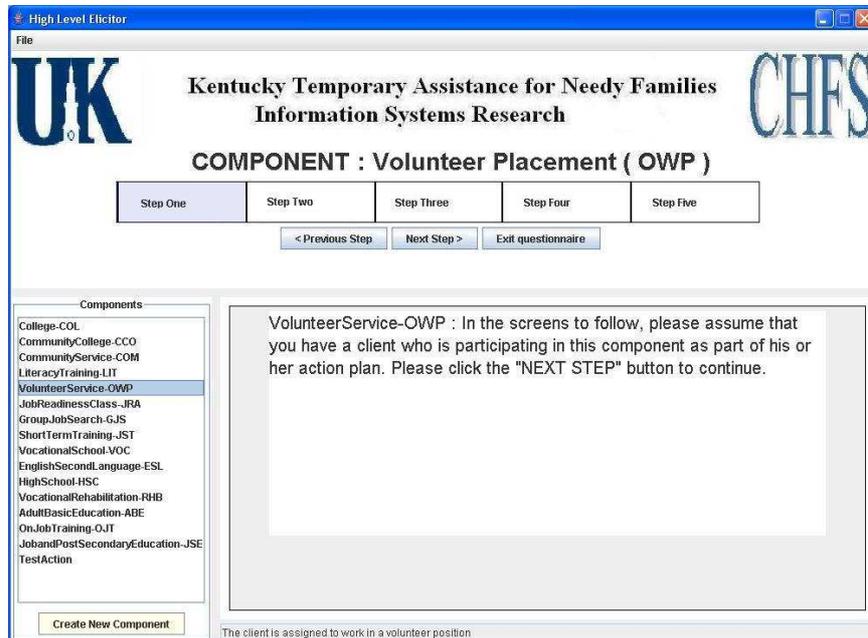


Figure 6: Introductory description for the action fragment “Volunteer Placement (VOP)”.

effect of input nodes and option choices for eliciting the effect on output nodes.

The software was  $\alpha$ -tested by the social scientists while the  $\beta$ -testing was done by former case managers. The  $\beta$ -testing helped us in collecting user feedback on usability, performance, and language used in the phrasing of instructions. All this information helped us in resolving ambiguity in instructions and task flow issues.

In the experiment, each case manager was asked to elicit bowtie models for five out of the sixteen actions. The actions were classified into two broad categories: work-related and education-related. Each case manager was provided with a userid. Actions assigned to case managers alternated between the two categories and were hardcoded based on the userid.

For a single action fragment, the elicitation proceeded as follows.

1. The first screen (Figure 6) provided details about the action to be elicited.
2. On the next screen (Figure 7), case managers were asked to choose the

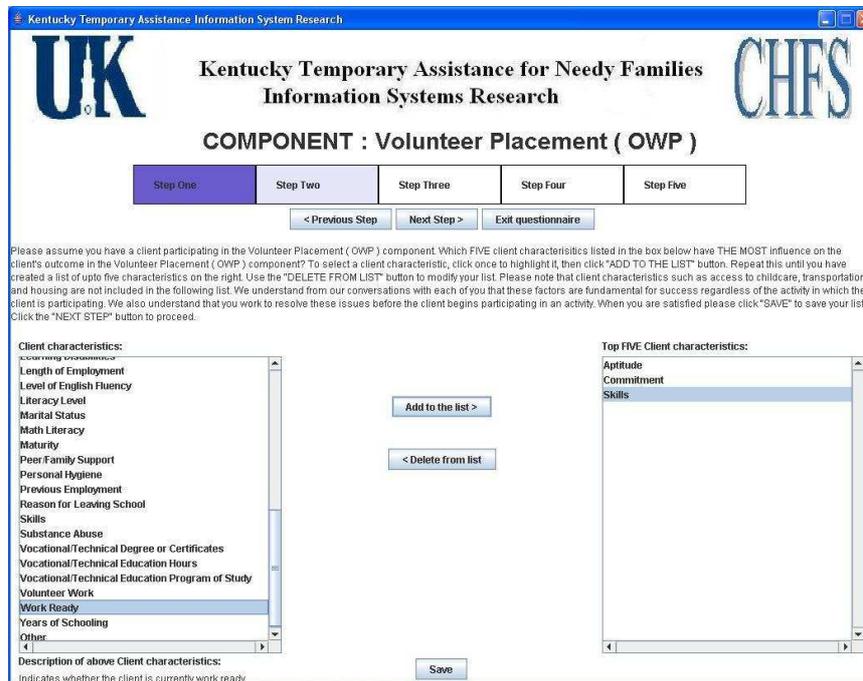


Figure 7: Step 2 of the elicitation process for the action fragment “Volunteer Placement (VOP)”.

- top five attributes that affect the outcome of the assigned action.
3. The case managers were then asked to rate these five attributes in terms of importance on a Likert scale (Figure 8).
  4. On the next screen (Figure 9), the experts were asked to assume that the client had succeeded in the assigned action. Given the success, the case managers were asked to review each of the 32 characteristics and report how they would have been affected (positively, negatively or no effect).
  5. The final screen (Figure 10) displayed a short verbal summary of the information submitted by the case managers.

Upon analyzing the data collected through the HELL elicitation, the research team felt the need for an extra piece of information in order to

Kentucky Temporary Assistance for Needy Families  
Information Systems Research

UK CHFS

COMPONENT : Volunteer Placement ( OWP )

Step One Step Two Step Three Step Four Step Five

< Previous Step Next Step > Exit questionnaire

The client characteristics that you feel have the most influence on the Volunteer Placement ( OWP ) component are listed below. We understand that all these characteristics are important. However, we also understand that some of them may be more important than the others. Using the sliding scales, please indicate the importance of each of the characteristics you listed. Please note that it is fine for one or more of the characteristics to be of equal importance. When you are finished, please click the "SAVE" button to save the data and then click the "NEXT STEP" button to proceed.

Work Ready More important Less important

Goals More important Less important

Aptitude More important Less important

Commitment More important Less important

Skills More important Less important

Save

Figure 8: Step 3 of the elicitation process for the action fragment “Volunteer Placement (VOP)”.

complete the elicitation model. The researchers realized that the information elicited about the output nodes was not sufficient. In addition to the kind of impact (positive or negative), we needed information on the level of impact. The social scientists used a spreadsheet (Figure 11) to elicit this information. The rows in the spreadsheet corresponded to client characteristics and the columns corresponded to actions. The shaded boxes indicated the client characteristics that the case manager reported as positively affected by success in an action. The case manager was asked to quantify the positive impact, using the following scale: 4 = very strong impact, 3 = strong impact, 2 = some impact and 1 = little impact. This process was repeated for the negatively impacted client characteristics.

Kentucky Temporary Assistance Information System Research

**UK** Kentucky Temporary Assistance for Needy Families Information Systems Research **CHFS**

**COMPONENT : Volunteer Placement ( OWP )**

Step One Step Two Step Three Step Four Step Five

< Previous Step Next Step > Exit questionnaire

Please assume that your client has SUCCESSFULLY completed the Volunteer Placement ( OWP ) component. Please indicate which of the client characteristics would change, positively or negatively, because of SUCCESSFUL participation in this component. If the client characteristic is not affected by the action, simply click the "NEUTRAL" button. When you are finished, please click the "SAVE" button to save the data and then click the "NEXT STEP" button to proceed.

|                             |  |  |  |
|-----------------------------|--|--|--|
| Access to Healthcare        | <input type="radio"/> positive effect            | <input checked="" type="radio"/> neutral | <input type="radio"/> negative effect            |
| Access to Housing           | <input checked="" type="radio"/> positive effect | <input type="radio"/> neutral            | <input type="radio"/> negative effect            |
| Access to Telephone Service | <input type="radio"/> positive effect            | <input type="radio"/> neutral            | <input checked="" type="radio"/> negative effect |
| Access to Transportation    | <input type="radio"/> positive effect            | <input checked="" type="radio"/> neutral | <input type="radio"/> negative effect            |
| Address                     | <input type="radio"/> positive effect            | <input type="radio"/> neutral            | <input checked="" type="radio"/> negative effect |
| Age                         | <input type="radio"/> positive effect            | <input checked="" type="radio"/> neutral | <input type="radio"/> negative effect            |
| Age of Youngest Child       | <input type="radio"/> positive effect            | <input checked="" type="radio"/> neutral | <input type="radio"/> negative effect            |
| Anger Management Issues     | <input type="radio"/> positive effect            | <input type="radio"/> neutral            | <input type="radio"/> negative effect            |

Save

Figure 9: Step 4 of the elicitation process for the action fragment “Volunteer Placement (VOP)”.

## 7.2 A Test of HELL

Our elicitation system is surprisingly efficient. We were able to elicit 90 fragments from 18 case managers in less than one hour. This gave us an average of 6 fragments per action. After the experiment, the social scientists interviewed the case managers on usability issues of the HELL tool. The responses were very encouraging. The majority of the case managers understood the intent of the elicitation and the scope of the questions.

The results of HELL reaffirmed the importance of the client’s interests, goals and commitment. When asked to list the five most important client characteristics for determining success in a given action fragment, level of COMMITMENT was included in the top five in 14 out of 15 actions and was listed as the most important characteristic in 4 actions. Similarly, CLIENT GOALS (i.e., whether the client *has* goals) was listed in the top five variables in 11 out of 15 actions and listed as one of the most important variable for

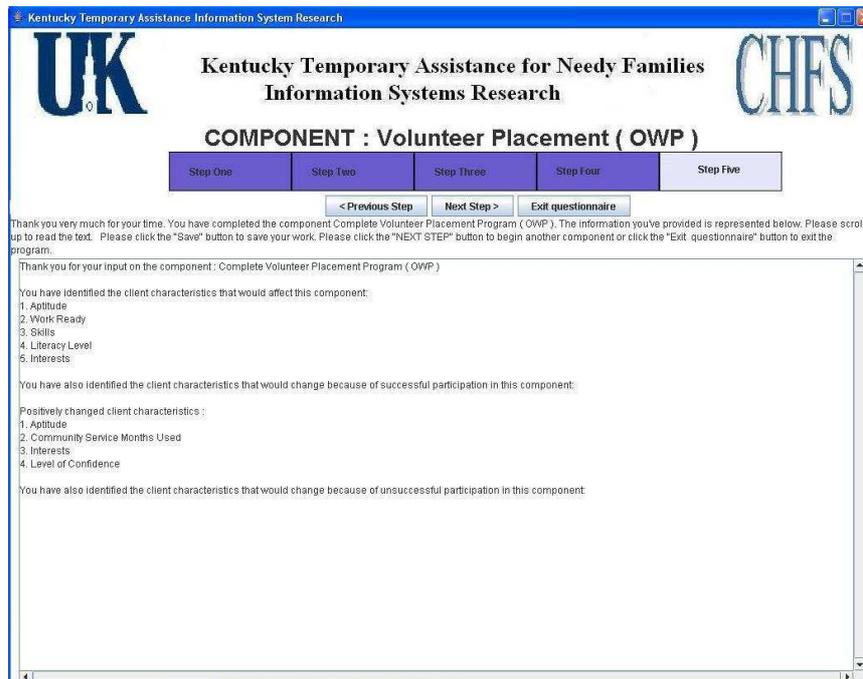


Figure 10: Textual summary description for the action fragment “Volunteer Placement (VOP)”.

4 actions. In particular, we observed significant consistency in the top four out of the five input nodes identified by different experts for each action fragment. This is due to the consistency with which experts responded to our model-building efforts. We also believe that a user-centered tool aided the case managers in the model-building process.

## 8 Other Bayes Net Elicitation Methods

There are algorithms for learning Bayes nets from data, or learning conditional probability tables from data, given the dependency graph of a Bayes net. We will not discuss those algorithms here. Because welfare recipients have such limited privacy from their case managers, data about them is highly protected. We have been unable to access statistical data.

One can elicit dependency structure separately from CPTs, but this tends to extend the elicitation process. The interaction between elicited

|  | COM | LIT | ABE | OJT | JSE |
|--|-----|-----|-----|-----|-----|
| 583 Respondent Thirteen Characteristics: |     |     |     |     |     |
| 585 Access to Clothing/Materials         |     |     |     | 4   | 4   |
| 586 Access to Dependent Care             |     |     | 4   |     |     |
| 587 Access to Healthcare                 |     |     | 4   |     |     |
| 588 Access to Housing                    |     |     | 4   |     |     |
| 589 Access to Telephone Service          |     |     | 4   | 4   |     |
| 590 Access to Transportation             |     |     | 4   | 4   |     |
| 591 Address                              |     |     |     | 3   |     |
| 592 Age                                  |     |     | 2   | 2   | 2   |
| 593 Age of Youngest Child                |     |     |     |     | 4   |
| 594 Anger Management Issues              |     |     |     | 4   | 4   |
| 595 Aptitude                             | 4   | 4   | 4   | 4   | 4   |
| 596 College Degree                       |     |     |     |     |     |
| 597 College Hours                        |     |     |     |     |     |
| 598 College Program of Study             |     |     |     |     |     |
| 599 Confidence                           | 4   | 4   | 4   | 4   | 4   |
| 600 Criminal Record                      |     |     |     | 4   | 4   |
| 601 Current Employment Status            |     |     |     |     | 2   |
| 602 Depression                           |     |     |     |     |     |
| 603 Disabilities                         |     | 4   | 4   | 4   | 4   |
| 604 Domestic Violence                    |     |     |     | 3   |     |
| 605 Goals                                |     |     |     | 3   | 3   |
| 610 Interests                            |     | 1   |     | 2   | 2   |
| 611 KTAP Months Used                     |     |     |     | 2   |     |

Figure 11: Step 5 of the elicitation process: Level of impact on output characteristics.

probabilities and dependencies provides a weak validation of the dependencies. However, many of the methods we discuss assume that the dependencies are already in place.

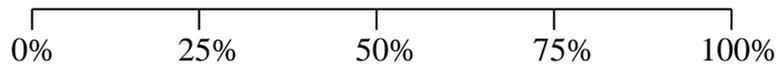


Figure 12: A numerical probability scale.

Individual probabilities can be elicited using a linear scale marked with numbers as in Figure 12 or a Likert-style scale with probability words (Renooij, 2001), or both, as in Figure 13. One can also elicit the probability of a variable having a certain value using the gamble method (Renooij, 2001), where the expert is asked to choose whether to play a lottery with a certain probability of payoff, or to gamble for the same payoff on the variable having the specified value. When the expert is indifferent between the two gambles, the probability of payoff in the lottery is assumed to be the expert's belief of the probability of the variable having that value. See Figure 14.

There are known problems with eliciting probabilities one by one, rather

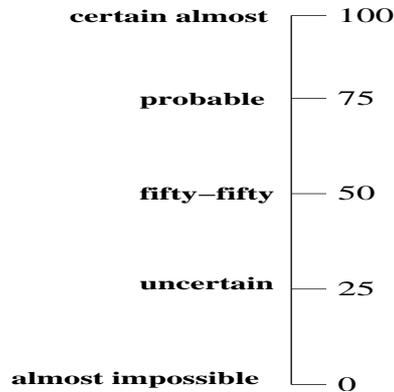


Figure 13: A scale with both numbers and probability words.

than in context (Tversky & Kahneman, 1983; Tversky & Koehler, 1994). For instance, from one probability to the next, the scale may change.

One can ask an expert to rank the relative likelihood of two values for a given variable (see Table 1 (Renooij, 2001)). Given many multi-valued variables, this method is tedious at best and yields inconsistent values at worst.

| Score | Relative likelihood                    |
|-------|--|
| 1     | A and B are equally likely.            |
| 2     | undecided between 1 and 3.             |
| 3     | A is weakly more likely than B.        |
| 4     | undecided between 3 and 5.             |
| 5     | A is strongly more likely than B.      |
| 6     | undecided between 5 and 7.             |
| 7     | A is very strongly more likely than B. |
| 8     | undecided between A and B.             |
| 9     | A is absolutely more likely than B     |

Table 1: Score table for pairwise comparison.

An entire probability distribution can be elicited at once using a set of linear scales set side-by-side for comparison (Zhao, Dekhtyar, Goldsmith, Jessup, & Li, 2004), a pie-chart style graphic called a probability wheel (Renooij, 2001), as in Figure 15, or in a spreadsheet. In any of these methods, one can allow arbitrary precision or can restrict the choices to a finite set of probability values. In the latter case, one can use pull-down menus for

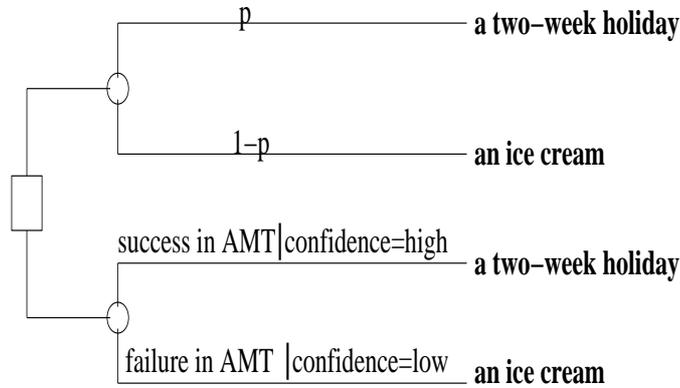


Figure 14: The gamble-like method.

more control over the input. Discrete values gloss over the lack of precision in most domain experts’ reasoning.

Simplifying assumptions about the probability distributions can greatly reduce the number of probabilities that must be elicited. For instance, if the distribution is uniform, no probabilities need be elicited; if it is linear, then only two probabilities must be. However, these are simplifications of simple distributions. In the case of a CPT with multiple parents, there are other newly canonical distributions (Pearl, 1988), such as Noisy-Or, Noisy-Max, Noisy-And, and Noisy-Min, that require a single probability for each value pair of single-parent value, child value—even if there are multiple parents. The challenge of these methods are to make clear to the domain expert the meanings of the elicited numbers in the context of the “canonical” distribution.

When Bayes nets can be broken down into smaller, independent fragments, it simplifies the elicitation process by allowing different fragments to be elicited separately (H.J., W.A.J.J., & E., 2002). The notion of eliciting small, possibly overlapping fragments is due to Laskey and Mahoney (Laskey & Mahoney, 2000). These fragments are often combined into a larger network using one of canonical distributions mentioned previously. This work on network fragments also suggests pattern elicitation, suitable both for Bayes net fragments (M., N., & L., 2000) or first-order Bayes nets (Jaeger, 1997; Koller, 1999; Laskey & Mahoney, 1998).

The canonical distributions, discrete-valued elicitation, and use of CPT patterns all implicitly admit that the elicited probabilities are imprecise. There are methods that address imprecision by eliciting upper and lower

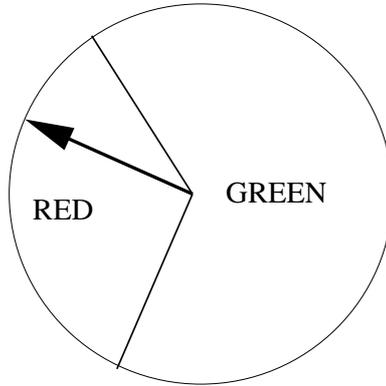


Figure 15: The probability wheel.

bounds on probabilities, either as guidelines for explicit probability elicitation (Helsper, van der Gaag, & Groenendaal, 2004; Druzdzel & van der Gaag, 1999), or to produce probability intervals.

Another approach to imprecision is the use of qualitative probabilistic networks (QPNs) (Wellman, 1990). This is the work most similar to ours. QPNs consist of dependency graphs, but have qualitative instead of quantitative representations of probabilities. Related to this is work on whether various domain experts think in numbers (Druzdzel & van der Gaag, 2000; van der Gaag, Renooij, Witteman, Aleman, & Taal, 2002, 1999); it seems that many do not.

In (Renooij & van der Gaag, 2002), an interactive process is introduced to guide experts through a quantification process. We have not yet implemented such a process for bowties, but that is a possible next step.

This paper has focused on case managers' mental models. The bowties grew out of their narratives. We are not the first to use some form of narrative in Bayes net building. Narrative approaches are used both for knowledge elicitation, for building dependency graphs (van der Gaag & Helsper, 2002), indirect elicitation of CPTs (Renooij & Witteman, 1999), and in validation of models (Popper, Lempert, & Bankes, 2005; Lempert, Bankes, & Popper, 2003). One can even see pattern elicitation as narrative as well, if one stretches the notion of a narrative. However, our story is different because it focused on listening, before we imposed a pattern on the domain experts' reasoning.

## 9 Putting the Pieces Together

This paper has described information elicitation for a particular application of Bayesian reasoning, a general-purpose formal representation of Bayes nets, and software for automating elicitation of a qualitative version of that representation. This is not, however, sufficient to build the full model of this application.

We have not described a full decision-support system. Many pieces are needed, including additional model and utility elicitation, data management, planning software, and additional user interfaces, particularly policy presentation. We have also not described qualitative-to-quantitative algorithms for the bowties.<sup>10</sup> Welfare-to-work clients usually take more than one action at a time. We have not specified how probabilities should be computed when several bowtie fragments are activated at once, and share output nodes. Furthermore, we have not presented MDP planning algorithms that take into account constraints or multiple simultaneous actions of varying durations.

What we have described is the outcome of listening to domain experts. By being flexible in our expectations, we were able to develop a new formalism that better captured the experts' reasoning. Serendipitously, this formalism led to computational simplifications in the elicitation process. Instead of considering  $\mathcal{O}(n^2)$  potential dependencies, the bowtie elicitation immediately limits the number of dependencies elicited to  $2n$ ; the maximum number of implicit and explicit edges is  $3n$ . In our elicitations, we chose to restrict the number of incoming edges to each action node, thus further simplifying the action fragment representations.

Bowties are not limited to the welfare domain. Consider the following three domains: the ubiquitous coffee robot (Boutilier et al., 1999), a generic medical example, and the New York regents exams.

In the coffee robot example, the robot can fetch coffee, deliver coffee, straighten the lab, and pick up mail. If the robot has coffee when it picks up the mail, and it succeeds in picking up mail, then it will have mail, and it might spill coffee on the mail, thus affecting the variable `ROBOTHAS-COFFEE`. If it fails to pick up the mail and did not have mail before, it cannot spill coffee on the mail. This—and any other action in this model—can be modeled by a bowtie.

In the medical model, treatments can be considered successful or failed. The effects of the treatments depend on prior states of variables (did the

---

<sup>10</sup>Our extension of HELL that produces quantitative networks is, naturally, referred to as LIMBO.

patient actually have the pathology being treated, for instance) and the success of the treatment. It is less immediate how to model information-seeking actions such as diagnostic tests, but these, too, can be modeled as bowties. In the New York public high schools, certain courses lead to state-wide examinations at the end of the year. The exams determine the student's mark in the course and whether the student passes on to the next grade level. Clearly, the student's progress can be modeled in terms of bowtie fragments.

This is not intended as a full catalogue of bowtie-amenable domains. A bowtie-amenable domain is one in which actions have distinct results that affect other variables in the system. By presenting action results as explicit variables in the model, bowties give knowledge engineers an intuitive and appealing elicitation process. Thus, the major contribution of this paper is to the knowledge engineering/information elicitation phase of Bayesian model building.

## References

- Allard, S. (2004). The Changing Urban Geography of Poverty and Service Programs.. The Brookings Institution Survey Series. Washington, DC.
- Boutilier, C., Dean, T., & Hanks, S. (1999). Decision-Theoretic Planning: Structural Assumptions and Computational Leverage. *Journal of Artificial Intelligence Research (JAIR)*, 11, 1–94.
- Druzdzal, M., & van der Gaag, L. (2000). Building Probabilistic Networks: Where Do the Numbers Come From?. *IEEE Transactions on Knowledge and Data Engineering*, 12(4), 481–486.
- Druzdzal, M. J., & van der Gaag, L. C. (1999). Elicitation of probabilities for Belief Networks: Combining qualitative and quantitative information.. In *Proceedings of the 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pp. 141–148 San Francisco, CA. Morgan Kaufmann Publishers.
- Helsper, E. M., van der Gaag, L. C., & Groenendaal, F. (2004). Designing a Procedure for the Acquisition of Probability Constraints for Bayesian Networks.. In *EKAW*, pp. 280–292.
- H.J., K., W.A.J.J., W., & E., T. B. (2002). Decision support for medical diagnosis.. In J.Meij (Ed.), *Dealing with the data flood. Mining data,*

text and multimedia. The Hague: STT/Bewetong (Study centre for Technology Trends).

- Jaeger, M. (1997). Relational Bayesian Networks.. In *UAI*, pp. 266–273.
- Kingfisher, C. (1996). *Women in the American Welfare Trap*. University of Pennsylvania Press, Philadelphia.
- Kingfisher, C. (2001). *The New Poverty Studies: The Ethnography of Power, Politics and Impoverished People in the United States.*, chap. Producing Disunity: The Constraints and Incitements of Welfare Work, pp. 273–292. New York University Press, New York.
- Koller, D. (1999). Probabilistic Relational Models. In *ILP '99: Proceedings of the 9th International Workshop on Inductive Logic Programming*, pp. 3–13 London, UK. Springer-Verlag.
- Laskey, K. B., & Mahoney, S. M. (2000). Network Engineering for Agile Belief Network Models.. *IEEE Trans. Knowl. Data Eng.*, 12(4), 487–498.
- Laskey, K., & Mahoney, S. (1998). Extensible Multi-Entity Models: A Framework for Incremental Model Construction. In *Valencia Conference on Bayesian Statistics* Valencia, Spain. Springer-Verlag.
- Lempert, R. J., Bankes, S. C., & Popper, S. W. (2003). *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-term Policy Analysis*. RAND Corporation.
- Lipsky, M. (1980). *Street-Level Bureaucracy: Dilemmas of the Individual in Public Services.*. New York: Russell Sage Foundation.
- M., N., N., F., & L., N. (2000). Building Large-scale Bayesian Networks.. *The Knowledge Engineering Review*, 15(3), 257–284.
- Mack, R., & Nielsen, J. (1993). Usability inspection methods: report on a workshop held at CHI'92, Monterey, CA. *SIGCHI Bull.*, 25(1), 28–33.
- Mayhew, D. J. (1992). *Principles and guidelines in software user interface design*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Morgen, S. (2001). The Agency of Welfare Workers: Negotiation Devolution, Privatization and the Meaning of Self-Sufficiency.. *American Anthropologist*, 103(3), 747–761.

- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.
- Popper, S. W., Lempert, R. J., & Bankes, S. C. (2005). Shaping the Future. *Scientific American*.
- Renooij, S. (2001). Probability elicitation for belief networks: Issues to consider.. *The Knowledge Engineering Review*, 16(3), 255–269.
- Renooij, S., & Witteman, C. (1999). Talking probabilities: Communicating probabilistic information with words and numbers.. *International Journal of Approximate Reasoning*, 22(3), 169–194.
- Renooij, S., & van der Gaag, L. C. (2002). From Qualitative to Quantitative Probabilistic Networks.. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, pp. 422–429.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgement.. *Psychological Review*, 90(4), 293–315.
- Tversky, A., & Koehler, D. J. (1994). Support theory: a nonextensional representation of subjective probability.. *Psychological Review*, 101(4), 547–567.
- van der Gaag, L., Renooij, S., Witteman, C., Aleman, B., & Taal, B. (2002). Probabilities for a probabilistic network: A case-study in oesophageal cancer. *Artificial Intelligence in Medicine*, 25, 123–148.
- van der Gaag, L., Renooij, S., Witteman, C. L. M., Aleman, B. M. P., & Taal, B. (1999). How to Elicit Many Probabilities. In *Proceedings of the 15th Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)*, pp. 647–659 San Francisco, CA. Morgan Kaufmann.
- van der Gaag, L. C., & Helsper, E. M. (2002). Experiences with Modelling Issues in Building Probabilistic Networks. In *EKAW '02: Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web*, pp. 21–26 London, UK. Springer-Verlag.
- Wellman, M. P. (1990). Fundamental concepts of qualitative probabilistic networks. *Artif. Intell.*, 44(3), 257–303.

Zhao, W., Dekhtyar, A., Goldsmith, J., Jessup, E., & Li, J. (2004). Building Bayes Nets with Semistructured Probabilistic DBMS. *GI EMISA Forum*, 24(1), 29–30.