SIMULATION METHOD FOR SOLVING HYBRID INFLUENCE DIAGRAMS IN DECISION MAKING

Xi Chen
Enlu Zhou

Department of Industrial & Enterprise Systems Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA

ABSTRACT

Influence diagrams (IDs) are powerful tools for representing and solving complex decision making problems. This paper presents a simulation-based approach for solving decision making problems formulated by hybrid IDs, which involve both discrete and continuous decision and chance variables. In the proposed method, Monte-Carlo simulation is applied in both approximating the expected conditional utility and solving the optimal decision strategies. The forward Monte-Carlo method is presented for expectation calculation, and it does not require Bayesian inference as in the standard "roll-back" method. The cross-entropy method in optimization is introduced to solve the optimal strategies. The decision variables are treated as random variables, and the decision strategies are solved by recursively updating the probability density of the decision variables. Finally, we present the simulation results of a bidding problem as an illustration.

1 INTRODUCTION

In decision making problems under uncertainty, the goal is to find an optimal alternative which maximizes the expected utility of the outcome. IDs (Howard and Matheson 2005) are graphical representation of the decision problems. It describes the probability dependence among the chance variables, specifies the information for each of the decisions, and provides a computational framework for solving the best decision strategy. Hybrid IDs are IDs involving both discrete and continuous chance and decision variables. Solving the optimal decision of hybrid IDs not only includes the computation of the expectation of the utility associated with a mixed joint distribution of discrete and continuous random variables, but also includes the optimization problem on the continuous support. Both the computation of expected utility and the optimization problem may be intractable, especially when the probability model is complex, the expected utility as an objective function is complicated, and a sequence of decisions is included. In addition, in most of the decision making problems, the interested marginal probability and the conditional probability in the standard "roll-back" method for solving IDs are not given, and Bayesian inference is required. Bayesian inference increases the computation, and may also be intractable in the hybrid IDs. Therefore, we propose a simulation method to provide approximate optimal strategies of the decision problem.

In the review of literature, most of the past research considered discrete IDs which contain only discrete chance and decision variables, and provided exact solutions, such as (Shachter 1986, Tatman and Shachter 1990, Jensen, Jensen, and Dittmer 1994, Zhang 1998). For the problems with continuous chance variables, some papers deal with the particular probability, the multivariate Gaussian, and solve the continuous IDs using exact methods, such as (Shachter and Kenley 1989, Madsen and Jensen 2005). Several papers studied approximation methods of solving IDs with continuous chance variables. Some papers use Gaussian mixtures to approximate arbitrary continuous probability (Poland 1994), or use mixtures of truncated exponential potentials to approximate probability density functions and utility functions (Cobb and Shenoy 2008). Some papers use discretization to provide the approximate solution, such as the moment methods for decision analysis (Smith 1993). The simulation method has also been studied. (Jenzarli 1995) investigated the use of Gibbs sampling for solving IDs; (Bielza, Muller, and Insua 1999) explored the approach of using Markov Chain Monte Carlo (MCMC) method to solve a single-stage decision problem by defining an artificial distribution of the product

space of alternatives and chances; (Charnes and Shenoy 2004) proposed a multi-stage method to decompose the problem of multiple decisions into many subproblems involving single decisions; (Cano, Gomez, and Moral 2006) proposed a forward-backward Monte Carlo method to solve the decision with very large models. However, most of the papers considered the problems that only involve the continuous chance variables, and few of them studied the problems with both continuous chance and continuous decision variables. In addition, in these simulation methods, they need the "roll-back" procedure and the Bayesian inference to generate the samples.

In this paper, a method of solving IDs with a mixture of discrete and continuous chance and decision variables by Monte Carlo simulation is proposed. We use simulation to evaluate the conditional expected utility, and also solve the optimization problem by simulation-based method to get the approximate optimal decision on a continuous support. The proposed method is suitable for the problem with discrete information sets, which means that the information known before making each of the decisions belongs to a discrete set. In the proposed method, the continuous decision nodes are treated as random variables, which are independent to all of the chance variables and other decisions at first, and the ID becomes a Bayesian network with a mixture of discrete and continuous nodes. By doing so, the simulation method to generate random variables in Bayesian network could be used to generate samples of the decision problem. We then classify these samples based on different associated information, and calculate the estimated expected utility of each category. To solve for the optimal decision strategies, we use the simulation-based optimization method. In particular, we use the cross-entropy method (Boer et al. 2005, Rubinstein 1999). In the proposed method, we only need the samples of the chance and decisions variables with specified joint probability distribution. Thus, we do not need to calculate the Bayesian inference, which reduces the computation, and is useful especially for problems where the Bayesian inference is intractable. Using cross-entropy method to solve the optimization allows us to deal with problems with complex expected utility functions on continuous domain.

The organization of the paper is as follows. In section 2, the ID representation of decision making problem is presented, and an example of an asymmetric decision problem involving Bayesian inference, a bidding problem, is described. In section 3, the simulation-based method for solving IDs is proposed. We first present the forward Monte-Carlo method for calculating the expected utility, and then present the cross-entropy method for solving the optimal strategies. In section 4, the simulation results of the bidding example are provided, and we discuss the comparison of the results by different smoothing update approaches in cross-entropy method. Finally, we provide concluding remarks and indicate some directions for future research in section 5.

2 INFLUENCE DIAGRAM REPRESENTATION OF BIDDING PROBLEM

To illustrate the proposed method, we consider a bidding problem as an example. This problem contains sequential decisions, and both discrete and continuous chance variables and decision variables. The ID of this decision making problem is a hybrid ID, which is usually hard to solve using the exact "roll-back" method. Moreover, this problem involves Bayesian inference of probabilities when solving for the best strategies. With the hybrid structure, the Bayesian inference is usually intractable.

2.1 Problem Statement

A company needs to decide how much to bid on the ownership of the patent of a new product. The profit of producing the new product is uncertain. The bid in this example is blind bid, in which the bidders do not know other bidders' bidding amount, and the highest bidder gets the patent immediately. That is, the company will get the patent if the value of its bidding amount is larger than or equal to the maximal competitive bid (MCB). To get a better bidding strategy, the company also has the option to pay some money to get the results of a survey, which provides the range of the MCB. But the range of MCB provided by the survey is imperfect.

Let P denote the profit of producing the new product in thousand dollar. P is a continuous random variable with truncated normal distribution $N(80, 15^2)$ in the support $[-20, 100]$. The instantiation of the profit is denoted by the lowercase p. The maximal competitive bid (MCB) in thousand dollar, denoted as B, given the profit of producing the product p, is a beta distribution in the support $B \in [-20, p]$. Given the profit p, the random variable $\frac{B+20}{p+20}$, which is normalized between 0 and 1, is according to beta distribution $Beta(10, 3)$. The company also has an option to pay $c = 2.2$ thousand dollar to get the result of the survey. The decision of whether to do the survey is denoted as S. Let $S = 1$ denote doing the survey, and $S = 0$ denote not doing the survey. The result of the survey is denoted as SR, and it provides the imperfect information about the range of the MCB. If $SR = low$ or 1, it implies that the maximal competitive bid lies in $[-20, 40]$; if $SR = mediun$ or 2, it implies that the maximal competitive bid is in $[40, 70]$; if $SR = high$ or 3, it implies that the maximal

competitive bid lies in $[70, 100]$. However, the result of the survey is imperfect.

$$SR = \quad 1 \quad w.p. \quad 1 \quad, \quad \text{if} \quad B \in [-20, 30]$$

$$SR = \begin{cases} 1 & w.p. \quad 0.5 \\ 2 & w.p. \quad 0.5 \end{cases}, \quad \text{if} \quad B \in (30, 50]$$

$$SR = \begin{cases} 2 & w.p. \quad 0.5 \\ 3 & w.p. \quad 0.5 \end{cases}, \quad \text{if} \quad B \in (50, 80]$$

$$SR = \quad 3 \quad w.p. \quad 1 \quad, \quad \text{if} \quad B \in [80, 100]$$

The decision of how much to bid is denoted as X, and the lowercase x denotes an instantiation of the bidding amount. The goal is to solve for the optimal decision whether to do the survey or not, and the optimal bidding amount, which is denoted as x^*, based on different information or survey results.

2.2 Influence Diagram Representation of Decision Problems

An influence diagram (ID) (Howard and Matheson 2005) is an acyclic graphical representation of the decision problems. An ID consists of directed arrows and three types of nodes: decision nodes, chance nodes, and value nodes. Decision node is denoted as a square, and represents the decision variables and set of alternatives available to the decision maker; chance node is denoted as a circle, and represents the uncertain events in the decision situation; value node is denoted as a diamond, and represents the value function or the utility function associated with the decisions and instantiations of the chance variables. There is zero or more decision and chance nodes, and at most one value node, which represents the objective or utility to be maximized in expectation. The arrows from chance nodes to chance nodes represent the probability dependence of the random variables; the arrows from chance or decision nodes to decision nodes represent the information known before making the decision; the arrows from decision nodes to chance nodes represent the decisions impact on chance; the arrows from chance and decision nodes to value node represent the arguments of the value or utility function.

The ID of the above bidding problem is shown in figure 1. As we may see from the diagram, there are
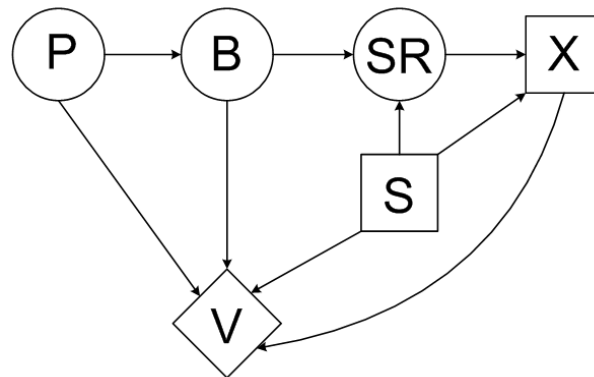

Figure 1: Influence diagram of the bidding problem

two decision nodes: S, whether to do the survey, and X, how much to bid; three chance nodes: P, the profit of producing the new product, B, the maximal competitive bid (MCB), and SR, the survey result. The MCB, B, depends on the profit P, the survey result SR depends on B. The decision whether to do the survey is made at the beginning, and the decision on how much to bid is made based on the observation of the survey result and whether the decision maker did the survey. The value function in this example is the payoff, which is a function of the profit P, the bidding amount X, the maximal competitive bid B, the decision of whether to do the survey S, and the cost of doing the survey c. Let $\mathbb{I}_{\{\cdot\}}$ denote the indicator function: $\mathbb{I}_{\{\cdot\}} = 1$ if the statement $\{\cdot\}$ holds, and $\mathbb{I}_{\{\cdot\}} = 0$ otherwise. Assume the decision maker is risk neutral, and the utility is equal

to the payoff of this decision situation. The utility function of this bidding problem is

$$U(p,x,b,S) = (p-x)\mathbb{I}_{\{x \geq b\}} - c\mathbb{I}_{\{S=1\}} \tag{1}$$

The decision maker will get the patent if and only if his or her bidding amount is larger than or equal to the MCB. We introduce an artificial state $SR = 0$, which indicates no survey result and $S = 0$. The objective function to be maximized is the expected payoff given the survey results. The optimal decision strategy is

$$(x^*, S^*) = \arg\max_{x,S} \mathbb{E}_{P,B}\left[U(P,B,X,S)|SR\right] \tag{2}$$

where $U(P,B,X,S)$ is represented in (1).

When solving the ID, the standard method is the "roll-back" method (Shachter 1986, Bielza and Shenoy 1999). That is to solve the ID by removing nodes in the backward sequence of the information flow. The node reduction order of this bidding problem is shown in Fig. 2. The chance node is removed by calculating the
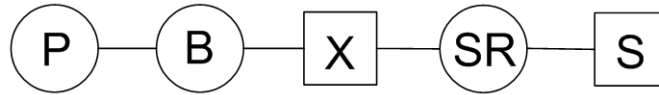

Figure 2: Nodes reduction order of the bidding problem

conditional expected value given its information set. At the first step, the value is the overall multiattribute utility function, and the value after removing a chance node is replaced by the conditional expected utility. The decision node is removed by maximizing the conditional expected value, and the value after removing a decision node is replaced by the maximal value. The procedure of node reduction of the "roll-back" method to solve ID for this bidding problem is listed as below:

1. Remove the chance node P. The value for node P is the multiattribute utility, and after removing node P, the value is replaced by

$$v_B(b,x,sr,s) = \mathbb{E}_P\left[U(P,B,X,S)|B=b,X=x,SR=sr,S=s\right]$$

2. Remove the chance node B. The value is replaced by

$$v_X(x,sr,s) = \mathbb{E}_B\left[v_B(B,X,SR,S)|X=x,SR=sr,S=s\right]$$

3. Remove the decision node X. The value is replaced by

$$v_{SR}(sr,s) = \max_x v_X(x,SR=sr,S=s)$$

The optimal bidding amount is $x^* = \arg\max_x v_X(x,SR=sr,S=s)$.
4. Remove the chance node SR. The value is replaced by

$$v_S(s) = \mathbb{E}_{SR}\left[v_{SR}(SR,S)|S=s\right]$$

5. At last, remove the decision node S. The value is replaced by

$$v = \max_s v_S(s)$$

The optimal decision of whether to do the survey is $s^* = \arg\max_s v_S(s)$.

By the standard "roll-back" method to solve IDs, the Bayesian inference to calculate the reversed conditional probability according to the information flow is required, since we only know the conditional or marginal probability given in the ID. Specifically, for solving the above bidding problem, the conditional probability $p(p|b,sr)$ and $p(b|sr)$ are needed in the "roll-back" procedure, but the ID only provides the conditional

probability $p(b|p)$ and $p(sr|b,p)$. With the probability densities given in this problem, the reversed conditional probability is intractable to solve. Using Bayesian inference to solve the reversed conditional probability is required in both the exact methods and the simulation methods in literature. In the problems with a mixed joint distribution of continuous and discrete chance variables and with non-Gaussian distributed chance variables, the exact method to solve the Bayesian inference is usually intractable. The optimal decision strategies could also be hard to solve if the expected conditional utility as an objective to maximize is complicated, especially for the problems with multiple local maxima.

3 SIMULATION METHOD IN SOLVING INFLUENCE DIAGRAMS

Because of the limitation of the exact method and the previously proposed simulation method to solve IDs, we present a new simulation method to solve IDs in this section, and simulation is used in both calculating the expectation and solving the optimization.

3.1 Simulation Method in Solving Expectation

Monte-Carlo simulation method is an important method to approximate the expectation. In this section, we consider the simulation method in calculating the conditional expected utility when solving IDs.

The probability model in influence diagram provides the order of factorization of the joint probability in terms of the marginal and conditional probability according to the Bayes's rule. If we change the decision node to a probability node, and delete the value node, the ID in Fig. 1 becomes the Bayesian network shown in Fig. 3. Thus, we may use the simulation method in Bayesian network to solve IDs. In the previous
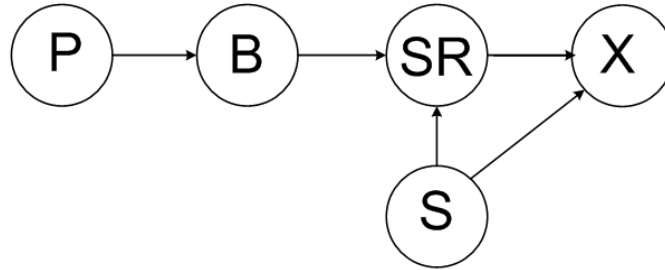

Figure 3: Probability model of bidding problem

"roll-back" method, the conditional probabilities given the observation are required, and we need to calculate the Bayesian inference in most of the decision making problems. However, the Bayesian inference may be intractable using the exact method. In order to avoid calculating the Bayesian inference, we use forward Monte-Carlo simulation method to generate the samples of the chance variables with given joint probability function. In this method, we do not solve the ID by reducing decision and chance nodes in the reverse order of the sequence determined by the information flow. However, we consider the joint probability and the whole domain of the variables of the decision problem, and only calculate the expected utility given the observation before each of the decisions.

For the bidding problem, we generate N independent sample sets (P_k, B_k, SR_k) $(k = 1, \cdots, N)$ in the same order as in the ID, $P \to B \to SR$. Thus, we only need the conditional probability density function $f(p)$, $f(b|p)$, and $f(sr|b)$, which are given in the ID. Then the samples (P, B, SR) are empirically distributed with the joint probability defined in the ID. Having the samples with desired joint distribution, we classify these samples based on the different scenarios. In the bidding problem, there are four different scenarios: bid without survey, bid given the survey result indicates low MCB, medium MCB, and high MCB. The expected conditional utility given different observations can be approximate by the summation. That is

$$\mathbb{E}^a[U(P,B,X,S)|SR = i] = \frac{1}{N_i} \sum_{k=1}^{N} \left[(p_k - x)\mathbb{I}_{\{x > b_k\}} - c\mathbb{I}_{\{S=1\}} \right] \mathbb{I}_{\{SR_k = i\}} \tag{3}$$

where $i = 0, 1, 2, 3$ and $N_i = \sum_k \mathbb{I}_{\{SR_k = i\}}$.

In general, the steps of the forward Monte-Carlo method are listed as below:

1. Consider the decision variables as chance variables with specified distribution.
2. Generate independent sample sets with the joint probability distribution specified in the ID in the same order defined in the ID according to forward Monte-Carlo method.
3. Classify the samples based on different information or observation of each of the decision scenarios.
4. Calculate the approximate conditional expected utility given different scenarios.

By the forward Monte-Carlo method, we generate samples over the overall joint space of the utility function, and do not calculate the expected utility locally as in the "roll-back" method. Our proposed method reduce the calculation especially in the problem that needs Bayesian inference when using standard "roll-back" method. Note that this method can only be used in the problems with discrete observations and with small number of instantiations of observation, since we need to classify the samples based on different scenarios and treat the problem globally.

3.2 Simulation Method in Solving Optimization

In this section, we present the simulation-based optimization method in solving decision making problems, where the conditional utility as objective function is complicated, and the optimal solution to maximize the expected utility is hard to solve by the exact method. We choose cross-entropy (CE) method (Boer, Kroese, Mannor, and Rubinstein 2005, Rubinstein 1999, Rubinstein and Kroese 2004), which is a Monte Carlo approach to optimization, and returns a near optimal solution.

The cross-entropy method is originated from rare event simulation and importance sampling. In the decision making problems, the decision variables are considered as random variables with specified parameterized family probability distribution. Then, we use the rate event simulation technique to update the probability density for simulation, such that there will be more samples generated around the maxima of the objective function. The probability density function updates iteratively, and for each iteration, there involves two stages as shown in Fig. 4.
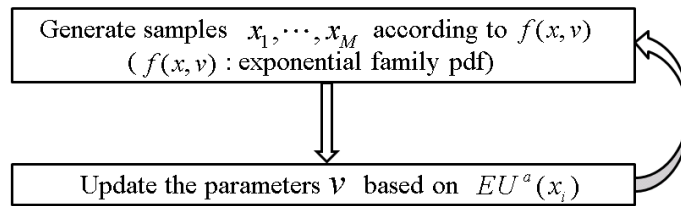

Figure 4: Two stages of cross-entropy method

Let the objective function we want to maximize be

$$x^* = \arg\max_x EU(x)$$

We first generate random samples $x_1, x_2, \cdots, x_M$ according to the specified probability density function $f(x, v)$, where v is the parameter we need to update. The random variable x is usually chosen to be an exponential family for the continuous optimization, such that the solution of the minimal cross-entropy has an analytical solution. It could also be used in discrete problems with updating the specified probability mass function (pmf). Then, we update the parameter v of the probability density function based on the value of $EU(x)$, such that we will generate better samples in the next iteration.

For the bidding problem in section 2, the decision variable of how much to bid, X, is a continuous decision variable, and is treated as a random variable according to the Gaussian distribution $N(\mu, \sigma^2)$, when using cross-entropy method to solve optimization. The variance σ^2 is large at the first step, and the variable X is approximately uniformly distributed, since we have no information about which decision value is better at the beginning. Based on the value of expected utility, we update the parameter μ and σ of the Gaussian distribution, and the samples become more and more concentrated on the maxima of the objective function, as the number of iteration increases. We stop until the variance become small enough, and the probability becomes a near Dirac delta function.

In some of the cases, the standard cross-entropy method presented above has the "freezing" problem, which means that the probability density will converge to a Dirac delta function very fast, and the iteration may stop at a local optimum instead of a global optimum. In order to improve the performance of the CE method, some smooth updating procedures have been proposed. One method is to use smoothing parameter (Boer, Kroese, Mannor, and Rubinstein 2005). Instead of updating the parameter $\hat{v}_{t-1}$ to $\hat{v}_t$ directly, we may introduce a smooth parameter θ, and the updated parameter of the probability density function is

$$\hat{v}_t = \theta \tilde{v}_t + (1-\theta)\hat{v}_{t-1}$$

where $\theta \in [0,1]$. Thus, the parameter $\hat{v}_t$ will not converge to fast. However, by introducing the smooth parameter, there will be a tradeoff between the accuracy and the computing speed. Sometimes, it still does not solve the "freezing" problem well.

To provide a better solution as well as reduce computation time, we could also add an artificial noise to the samples and disturb the samples, such that it will not converge too fast (Zhou, Fu, and Marcus 2010). The artificial noise is large at the beginning, and gradually becomes smaller as it converges to the true optima. Thus, instead of using the samples x_k generated from the probability density function $f(x, \hat{v}_t)$, we disturb the sample by a noise, and use

$$x_k^a = x_k + \omega_k, \qquad k = 1, 2, \cdots, M$$

where

$$\omega_k = \alpha \beta^{k-1} \omega, \qquad \beta \in (0, 1)$$

ω is the noise, and is usually chosen to be the standard Gaussian noise $N(0, I)$, α is the initial standard deviation of the noise, and β is the coefficient of the noise reduction rate. This method is called noisy CE method, and is denoted as NCE. By introducing the artificial noise, the number of iterations needed may reduce, and it may return better solution especially in the problem with multiple local optima, since by adding the noise, the solution is not easily trapped in local optima.

We may also use an adaptive smoothing parameter to provide better solution with less iterations, and we denote the adaptive smoothing CE method by ACE. It is a variation of the adaptive smoothing in (Rubinstein and Kroese 2004) with different adaptive smoothing parameter. Instead of using the constant smoothing parameter θ, we change the parameter every iteration, and the adaptive parameter is denoted by θ_t. At first, θ_t is chosen to be near 1, and the probability parameter $\hat{v}_t$ changes fast; as iteration increases, the smoothing parameter θ_t becomes smaller, such that it may return more accurate solution. The parameter $\hat{v}_t$ is updated by

$$\hat{v}_t = \theta_t \tilde{v}_t + (1-\theta_t)\hat{v}_{t-1}$$

where

$$\theta_t = \theta + (t+1)^{-a}$$

θ is chosen to be a small number such as 0.1 or 0.2, and $a \in [0, 1]$. As the iteration number t goes to infinite, the smoothing parameter θ_t converges to θ; at the beginning, the smoothing parameter θ_t is near 1. The rate of the change of θ_t may be adjusted by the value of the parameter a.

4 SIMULATION RESULTS

To illustrate the effectiveness of the proposed simulation method, we provide the simulation results of the bidding problem in section 2. The decision maker needs to make the decision how much to bid and whether to do the survey to get information about the MCB before bidding. But the survey gives imperfect information about the MCB. The uncertainty involved in the problem is the profit of the bidding item P, the MCB B, and the survey results SR.

The ID of this bidding problem is given in Fig. 1. The conditional probabilities associated with the ID are given in section 2. The profit P is with truncated normal distribution $N(80, 15^2)$ in the support $[-20, 100]$; the maximal competitive bid B given P is beta distributed with $\alpha = 10$ and $\beta = 3$, that is $\frac{B+20}{P+20} \sim Beta(10, 3)$. The cost of the survey is $c = 2.2$, and the survey gives imperfect information. If the real maximal competitive bid B is in $[-20, 30]$, SR is equal to 1 with probability 1; if B is in $(30, 50)$, SR is equal to 1 or 2 with probability 0.5; if B lies in $[50, 80]$, SR is equally likely to be equal to 2 or 3; if B lies in $(80, 100]$, SR is equal to 3 for sure.

We use the proposed simulation method to solve this decision problem, and choose the sample size N for approximating the expected utility to be 10^4, and the sample size M for calculating the optimal decision strategies to be 300. We use the CE method with fixed smooth parameter θ, the noisy CE method (NCE), and the adaptive smoothing CE method (ACE) to calculate the optimal bidding amount, and compares the results in table 1. In the CE method the smoothing parameter is chosen to be $\theta = 0.8$ and $\theta = 0.3$. In

Table 1: Optimal decision strategies of bidding Problems

	CE ($\theta = 0.8$)	CE ($\theta = 0.3$)	NCE	ACE
Iteration	11	32	25	13
$x^*|SR = 0$	57016	57012	57013	57013
$x^*|SR = 1$	49999	49969	49968	49969
$x^*|S = 2$	57041	57044	57043	57045
$x^*|S = 3$	64765	64763	64764	64764
s^*	1	1	1	1
EU^*	7768	7769.8	7769.8	7769.7

NCE, the artificial noise is chosen to be the standard Gaussian noise $N(0, I_{n \times n})$, and the parameter $\alpha = 2$ and $\beta = 0.98$. In ACE, θ is chosen to be 0.1 and a is 0.4. From the table, we may see that the CE method with smoothing parameter $\theta = 0.3$, the NCE and ACE gives better optimal solution compared with the CE method with $\theta = 0.8$. However, the CE method with $\theta = 0.3$ needs more iterations, and NCE and ACE method need less iterations. The adaptive smoothing parameter CE method needs almost the same number of iterations as the CE method with $\theta = 0.8$, but the ACE method has better solution. We also did the simulation by changing the number of sample size M in calculating the optimization, and observe the average performance of the standard CE, NCE and ACE. The results are shown in Fig. 5. The black curve is the average maximal
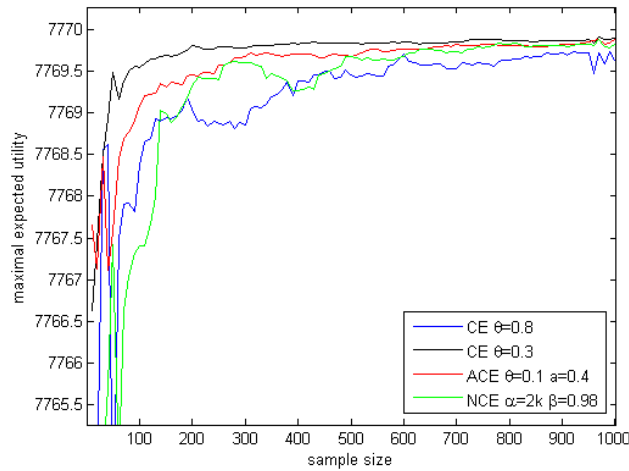

Figure 5: Comparison of the smooth updating approaches in cross-entropy method

expected utility by CE method with smoothing parameter $\theta = 0.3$, the blue curve is for CE with $\theta = 0.8$, the red curve is for adaptive smoothing CE (ACE), and the green curve is for noisy CE (NCE). In this particular bidding problem, the performance of ACE and NCE is better than the CE with smoothing parameter $\theta = 0.8$, but they needs less iterations compared with the CE with $\theta = 0.3$.

Note that the presented bidding problem is only a simple example with commonly used distribution. In the problems that the objective function to maximize has multiple local optima, the NCE and ACE method may have much better performance than the standard CE.

5 CONCLUSION

In this paper, we proposed a simulation method to solve IDs in complex decision making problems. We consider the decision problems with a hybrid of continuous and discrete chance and decision variables. The proposed method uses simulation in both calculating the expectation and optimization involved in solving IDs. We use forward Monte-Carlo method to approximate the conditional expectation, and avoid to solve Bayesian inference, which is intractable in most hybrid IDs. For solving optimization, we use cross-entropy method, which returns near optimal global solutions. The proposed method applies in most hybrid decision making problems with discrete information sets. As an illustration, we incorporate the proposed method in the bidding problem, and discuss the comparison of smooth updating approaches in cross-entropy method.

Our current and future research directions of interest include solving more sophisticated decision making problems such as bidding problems with sequential bidding decisions and the updated information of the maximal competitive bid, and solving the problem with different simulation-based optimization.

REFERENCES

Bielza, C., P. Muller, and D. R. Insua. 1999. Decision analysis by augmented probability simulation. *Management Science* 45 (7): 995–1007.

Bielza, C., and P. P. Shenoy. 1999. A comparison of graphical techniques for asymmetric decision problems. *Management Science* 45 (11): 1552–1569.

Boer, P., D. Kroese, S. Mannor, and R. Rubinstein. 2005. A tutorial on the cross-entropy method. *Annals of Operations Research* 134 (1): 19–67.

Cano, A., M. Gomez, and S. Moral. 2006. A forward-backward monte carlo method for solving influence diagrams. *International Journal of Approximate Reasoning* 42:119–135.

Charnes, J. M., and P. P. Shenoy. 2004. Multistage monte carlo method for solving influence diagrams using local computation. *Management Science* 50 (3): 405–418.

Cobb, B. R., and P. P. Shenoy. 2008. Decision making with hybrid influence diagrams using mixtures of truncated exponentials. *European Journal of Operational Research* 186 (1): 261–275.

Howard, R. A., and J. E. Matheson. 2005. Influence diagrams. *Decision Analysis* 2 (3): 127–143.

Jensen, F., F. V. Jensen, and D. Dittmer. 1994. From influence diagrams to junction trees. In *Proceedings of the 10th Conference on Uncertainty in Artificial Intelligence*, 367–373.

Jenzarli, A. 1995. *Modeling dependence in project management*. PhD Thesis, University of Kansas Business School, Lawrence, KS.

Madsen, A. L., and F. Jensen. 2005. Solving linear-quadratic conditional gaussian influence diagrams. *International Journal of Approximate Reasoning* 38:263–282.

Poland, W. B. 1994. *Decision analysis with continuous and discrete variables: A mixture distribution approach*. PhD Thesis, Stanford University, Department of Engineering-Economic Systems, Stanford, CA.

Rubinstein, R. 1999. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability* 1:127–190.

Rubinstein, R. Y., and D. P. Kroese. 2004. *The cross-entropy method: A unified approach to combinatorial optimization, monte-carlo simulation, and machine learning*. New York, NY: Springer.

Shachter, R. D. 1986. Evaluating influence diagrams. *Operations Research* 43 (6): 871–882.

Shachter, R. D., and C. R. Kenley. 1989. Gaussian influence diagrams. *Management Science* 35 (5): 527–550.

Smith, J. E. 1993. Moment methods for decision analysis. *Management Science* 41 (2): 340–358.

Tatman, J. A., and R. D. Shachter. 1990. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics* 20 (2): 365–379.

Zhang, N. L. 1998. Probabilistic inference in influence diagrams. In *Proceedings of the 14th Annual Conference on Uncertainty in Artificial Intelligence*, 514–522.

Zhou, E., M. Fu, and S. Marcus. 2010. A particle filtering framework for randomized optimization algorithms: EDAs, CE, MRAS, and more. under revision.

AUTHOR BIOGRAPHIES

XI CHEN is a Ph.D student in the Department of Industrial & Enterprise Systems Engineering, University of Illinois at Urbana-Champaign. Her research interests include simulation optimization and decision analysis. Her email address is <xchen37@illinois.edu>.

ENLU ZHOU is an Assistant Professor in the Department of Industrial and Enterprise Systems Engineering, at the University of Illinois at Urbana-Champaign. She received the B.S. degree with highest honors in electrical

engineering from Zhejiang University, China, in 2004, and received the Ph.D. degree in electrical engineering from the University of Maryland, College Park, in 2009. Her research interests include stochastic control, nonlinear filtering, and simulation optimization. Her email address is <enluzhou@illinois.edu>.