Proceedings of the 2006 American Control Conference
Minneapolis, Minnesota, USA, June 14-16, 2006

On the Development of an Intelligent Computer Player for CLUE®: a Case Study on Preposterior Decision Analysis

Chenghui Cai and Silvia Ferrari

Abstract— **The detective boardgame of CLUE® can be viewed as an example of preposterior decision problem, where decisions on how to navigate the board are based on the expected utility of the observations, and the observations are aimed at improving an inference process. The same principles arise in modern surveillance systems, such as demining sensor networks, where the sensor platforms (e.g., autonomous ground or vehicles) move about the environment in order to collect measurements or *evidence* from unknown targets and improve inference of unknown features. The boardgame of CLUE® serves as a well-known and intuitive example problem, that displays the same couplings between motion planning and inference, as modern surveillance systems. In this paper, a Bayesian network (BN) approach is used to develop an automated computer player for CLUE®, that is tested through an interactive simulation of the game. The results show that the intelligent player plans its motions according to the evidence that needs to be collected, and is capable of winning the game of CLUE® against experienced human players.**

I. Introduction

THE boardgame of CLUE® is a popular detective game with the purpose of determining the guilty suspect, the weapon, and the room of an imaginary murder. Players ultimately find the answers by entering rooms, making suggestions, and obtaining other players' responses supporting or refuting their suggestions. In this paper, a network modeling approach using Bayesian networks (BNs) is developed to model the process of inferring the answer to the murder posed in CLUE®. Furthermore, intelligent movements and suggestion-making techniques of the individual players can be explored based on the proposed Bayesian network model. The intelligent computer player designed in this paper is capable of automatically developing a strategy for navigating the mansion illustrated on the game board, based on the evidence collected from the responses of the other players. The tools developed in this paper, which include an interactive simulation of the game and an intelligent computer player implementing BNs, also allow to investigate and illustrate the underlying principles of the decision problems that arise in modern surveillance systems.

Recently, there has been much interest in modern surveillance systems with increased flexibility and functionality, where the sensors and their platforms are characterized by a high degree of autonomy, reconfigurability, and redundancy [1]-[3]. The initial motivation of the CLUE® research is the similarity between these surveillance systems objectives and those of the game of CLUE®. Recent research in the AI community seeks to develop computer games implementing advanced AI [4,11], supporting the idea that computer game simulations are useful for investigating and demonstrating intelligent systems and algorithms.

II. Background

A. The Game of CLUE®

There are nine rooms in the mansion shown in Fig. 1, i.e., dining room, library, billiard room, hall, kitchen, lounge, ballroom and conservatory, and six suspects, i.e., Col. Mustard, Miss Scarlet, Prof. Plum, Mr. Green, Mrs. White and Mrs. Peacock. All of these pawns are potential suspects. As for the weapon, there are six possibilities: knife, rope, candlestick, lead pipe, revolver, and wrench. An interactive simulation of the game of CLUE® is developed, as described in Section IV, in order to allow a human to play against the intelligent computer player.

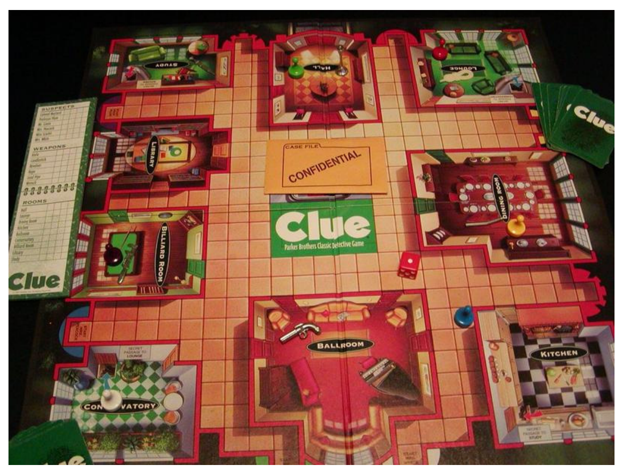

Fig. 1. CLUE® game board. CLUE® & ©2006 Hasbro, Inc. Used with permission.

Each item in the three categories of suspect, weapon, and room is represented by an illustrated card and there are a total of twenty-one cards in the deck. One card from each category is randomly selected and removed from the deck. The selected cards represent the true killer, weapon, and room of the crime and are hidden in an envelope. The remaining cards are then dealt to the players. Players move their pawns from room to room and upon entering one of the nine rooms make a "suggestion" about their belief of the guilty suspect, the weapon and the room where their pawns have moved into. By making suggestions, players collect information about other players' cards, and can infer the

Manuscript received September 15, 2005. This work was supported by the Office of Naval Research Young Investigator Program (Code 321).

C. Cai is a graduate student of Mechanical Engineering at Duke University, Durham, NC, 27708, USA cc88@duke.edu

S. Ferrari is with Faculty of Mechanical Engineering at Duke University, Durham, NC, 27708, USA sferrari@duke.edu

1-4244-0210-7/06/$20.00 ©2006 IEEE

hidden cards in the envelope.

There are three ways for a player to enter a room: (i) entering through one of the doors, (ii) via the Secret Passages, corner to corner (see Fig. 1), or (iii) being taken to the room by another player who is suggesting the player's pawn as a suspect. Both (i) and (ii) are moves made only during his/her turn of play, whereas (iii) is done during another player's turn. When a suggestion has been made by a particular player (Player 1) upon entering a room by method (i) or (ii), the next player (Player 2) proves or disproves the suggestion either by saying, "I have NONE of the suggested cards" or by showing one of the cards to the suggestion maker. If Player 2 has none of the cards, it is the turn of the second next player (Player 3) to prove or disprove the suggestion made by Player 1. This is continued until the suggestion has been disproved or all of the players have been asked. By making suggestions and observing the outcomes of the other players' responses, information is obtained about their cards. When players are confident in their knowledge of the hidden cards, they can make an accusation. If the accusation is correct, the player wins the game; otherwise he/she will lose and exit the game.

B. Bayesian Network Inference

In CLUE®, players infer the three hidden cards based on information gathered from the responses to previous suggestions. By viewing the information as evidence, inference can be carried out by a BN [6] that represents the relationships between the hidden cards and those dealt to the players.

In this paper, capital letters denote variables and lowercase letters denote the *states* or instantiations of the variables (i.e., X_i is said to be in its j^{th} instantiations when $X_i = x_{i,j}$). A BN is a directed acyclic graph (DAG) with *conditional probability tables* (CPTs) attached to each node. If there is a link from node A to node B, it is said that B is a child of A and A is a parent of B [5]. A CPT lists in tabular form the conditional probabilities of each node or state variable, X_i, given its parents, or $p(X_i|\pi(X_i))$, where $\pi(X_i)$ denotes the parents of X_i. Let μ_i denote the instantiations of the children of variable X_i. By utilizing Bayes' rules of inference, given evidence about the observed variables in the network, the posterior probability distribution of variable X_i can be computed as follows:

$$p(X_i \mid \mu_i) = \frac{p(\mu_i \mid X_i)\, p(X_i)}{p(\mu_i)}. \tag{1}$$

The prior probability of X_i, $p(X_i)$, is the known probability distribution over the states of X_i, $(x_{i,1}, \ldots, x_{i,r_i})$. The likelihood function, $p(\mu_i|X_i)$, contains the conditional probabilities of the instantiated children variables connected to X_i. This probability is the product of the likelihood probabilities of the instantiated variables $p(\mu_i \mid X_i) = \prod_j p(\mu_{i(j)} \mid X_i)$, where $\mu_{i(j)}$ represents the instantiation of the j^{th} child of X_i. The marginalization over the observed variables, $p(\mu_i)$, accounts for the relationship between the instantiated variables and all of the possible states of X_i,

$$p(\mu_i) = \sum_{k=1}^{r_i} p(X_i = x_{i,k}) \prod_j p(\mu_{i(j)} \mid X_i). \tag{2}$$

The posterior probability $p(X_i = x_{i,k}|\mu_i)$, also referred to as the marginal probability of $x_{i,k}$, represented the likelihood or confidence level in $x_{i,k}$ given the μ_i. Therefore, using (1), X_i can be inferred from the available hard evidence.

Hard evidence refers to perfect knowledge of a node's instantiation, and is accounted for in (1). If the probability distribution of a node over its possible values is known, it is referred to as *soft evidence*, and denoted by *se*. In this case, Jeffrey's rule can be used as a mechanism to update soft evidence, represented as a distribution $Q(\mu_i)$ [7, 8]; the rule can be written in (3)-(4),

$$p(\mu_i \mid se) = Q(\mu_i) \tag{3}$$

$$\begin{aligned} p(X_i \mid se) &= \sum_{\mu_i} p(X_i \mid \mu_i)\, p(\mu_i \mid se) \\ &= \sum_{\mu_i} p(X_i \mid \mu_i)\, Q(\mu_i). \end{aligned} \tag{4}$$

In (3), the soft evidence distribution $Q(\mu_i)$ can be explained as the conditional probability of μ_i being instantiated given soft evidence *se*. Therefore, $Q(\mu_i)$ can be viewed as the weights of $p(X_i |\mu_i)$ which can be calculated by (1). By comparing (1) and (4), hard evidence can be viewed as a special type of soft evidence whose conditional possibility given soft evidence is 1 for the instantiated or observed value, and 0 on other values.

III. Methodology

A. BN Inference in the Game of CLUE®

An intelligent computer player (ICP) for CLUE® is developed by using a BN to infer the hidden cards from the collected evidence. The evidence pertains to the cards belonging to the other players, which at the onset of the game are unknown to the ICP. In order to replicate the inference process, a so-called CLUE® BN is developed where each card is represented by a BN node. Both the relationships between the nodes (BN arcs) and the CPTs are determined by inspecting the game rules.

Assume there are three players in the game. After three guilty cards comprised of one suspect card (K), one weapon card (W), and one room card (R) are hidden, the remaining eighteen cards are thoroughly mixed together and shuffled, and then dealt one by one such that six cards go to each player. Every time a card is dealt, its value influences the cards that are dealt later. Let C_{ji} denote the j^{th} card of i^{th} player. Then, the general CLUE® BN model represents the relationships between all of the CLUE® cards, shown in Fig. 2. Each card C_{ji} will influence the cards of the same type (suspect, weapon, or room) that are dealt later in the game, i.e., $C_{ji} \rightarrow C_{mn}$, when $n > i$ or when $i = n$, $m > j$. Although cards of a different type do not influence each other (i.e., are conditionally independent), the typology of the dealt cards is unknown. Therefore, in this BN none of the arcs can be eliminated *a priori*.

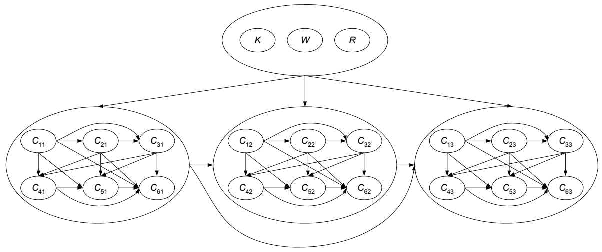

Fig. 2. General CLUE® BN model

Let the information gathered by the ICP be the evidence e. If its CPTs are known this BN can be used to compute the posterior probability distributions $p(K|e)$, $p(W|e)$, and $p(R|e)$, i.e., to infer the hidden cards K, W, and R. The size of the CPT attached to a node X_i, referred to as $|\mathrm{CPT}(X_i)|$, is equal to the product of the number of instantiations of X_i times the number of instantiations of each node in $\pi(X_i)$. In the general CLUE® BN (Fig. 2), the suspect cards have six possible instantiations, i.e., the *suspect domain* is $dom(K)$ = {Col. Mustard, Miss Scarlet, Prof. Plum, Mr. Green, Mrs. White, Mrs. Peacock}; similarly, $dom(W)$ = {knife, rope, candle-stick, lead pipe, revolver, wrench} and $dom(R)$ = {dining room, library, billiard room, hall, kitchen, lounge, ballroom, conservatory}. It follows that C_{ji}, whose typology is unknown *a priori*, has twenty-one possible instantiations and its domain is $dom(C_{ji}) = dom(K) \cup dom(W) \cup dom(R)$. Also, all previous hidden or dealt cards are parents of C_{ji},

$$|\pi(C_{ji})| = 3 + 6(i-1) + (j-1) = 6i + j - 4, \qquad (5)$$

where, $|\cdot|$ denotes the member number of a set. For example, $|\pi(C_{ji})| = 20$, when $j = 6$ and $i = 3$. The size of the CPT attached to C_{ji}, denoted by $|\mathrm{CPT}(C_{ji})|$, is given by:

$$|\mathrm{CPT}(C_{ji})| = \prod_{X \in \pi(C_{ji}) \cup \{C_{ji}\}} |dom(X)|. \qquad (6)$$

For example, $|\mathrm{CPT}(C_{63})| = 6 \times 6 \times 9 \times 21^{17} = 9.7336 \times 10^{24}$.

The children of K, W and R, denoted by μ_K, μ_W and μ_R respectively are all C_{ji}, j = 1, 2, …, 6 and i = 1, 2, 3. The object of the BN model is to calculate $p(K|se)$, $p(W|se)$ and $p(R|se)$. From (1) and (4), the probability of the hidden card, T, is,

$$p(T \mid se) = \sum_{\mu_T} p(T \mid \mu_T) Q(\mu_T), \qquad (7)$$

where,

$$p(T \mid \mu_T) = \frac{p(\mu_T | T)\, p(T)}{p(\mu_T)}, \qquad (8)$$

and $T = K$, W or R, represents the card type. The general CLUE® BN model shown in Fig. 2 is exact. However, the computational complexity of the inference task using BN is NP-hard [9]. Therefore, inference is computationally infeasible for the general CLUE® BN.

The computational complexity of the inference task can be reduced by exploiting the fact that cards of different typology are conditionally independent through the following assumptions:

1) Player P_1 (the ICP), always has two suspect cards, two weapon cards, and two room cards.
2) Player P_2 (the player next to P_1) always has 2 suspect cards, 1 weapon card and 3 room cards.
3) Player P_3 (the player next to P_2) always has 1 suspect card, 2 weapon cards and 3 room cards.

Subsequently, the cards can be labeled not only by the order in which they are dealt and by the player's number, but also by their typology. For example, $C_{ji}^{(k)}$ denotes the suspect card that is number j in the deck of the i^{th} player; similarly, $C_{ji}^{(w)}$ and $C_{ji}^{(r)}$ respectively denote the weapon and room cards. Although each player can shuffle its own deck, since only in the cards' instantiations are of interest, the cards can be labeled in this order without loss of generality. In order to label the cards according to their typology, the cards are shuffled and dealt using three separate decks containing the suspect, weapon, and room cards, respectively.

Since the cards dealt to the ICP player cannot appear in the hidden deck nor in the other players' hands, they can be removed from the BN model, thereby reducing the domains of the three card types. The simplified BN model is shown in Fig. 3. Here, the domain of node K, $dom(K) = \{K_j: j = 1, 2, \ldots, 4\}$, excludes the two suspect cards dealt to the ICP. Similarly, the ICP's weapon and room cards can be excluded, such that, $dom(W) = \{W_j: j = 1, 2, \ldots, 4\}$ and $dom(R) = \{R_j: j = 1, 2, \ldots, 7\}$. Equations (6)-(8) still hold for the simplified CLUE® BN model and, therefore, are used to infer the posterior probabilities of K, W, and R. Now that a feasible CLUE® BN structure has been identified, the corresponding CPTs are determined by inspection.

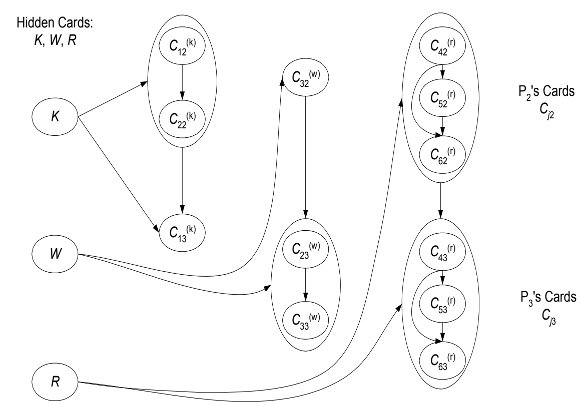

Fig. 3. Simplified Clue® BN model for intelligent computer player

A variable *unique*(.) is introduced such that $unique(C_{ji})$ = 1 if the instantiated cards of all members of $\pi(C_{ji}) \cup \{C_{ji}\}$ are different from each other, otherwise, $unique(C_{ji}) = 0$. The CPTs for this simplified BN model are:

$$p(T = T_j) = 1/|dom(T)|, \forall T_j \in dom(T), \qquad (9)$$

where $T = K$, W or R.

$$p(C_{ji}^{(k)} \mid \pi(C_{ji}^{(k)})) = \begin{cases} 0, \text{ if } unique(C_{ji}^{(k)})=0 \\ 1/(|dom(K)|-|\pi(C_{ji}^{(k)})|), \text{ else} \end{cases} \qquad (10)$$

$$p(C_{ji}^{(\mathrm{w})} \mid \pi(C_{ji}^{(\mathrm{w})})) = \begin{cases} 0, \text{ if } unique(C_{ji}^{(\mathrm{w})})=0 \\ 1/(|dom(W)|-|\pi(C_{ji}^{(\mathrm{w})})|), \text{ else} \end{cases} \quad (11)$$

$$p(C_{ji}^{(\mathrm{r})} \mid \pi(C_{ji}^{(\mathrm{r})})) = \begin{cases} 0, \text{ if } unique(C_{ji}^{(\mathrm{r})})=0 \\ 1/(|dom(R)|-|\pi(C_{ji}^{(\mathrm{r})})|), \text{ else} \end{cases} . \quad (12)$$

B. Evidence Table (ET) Construction

The approach for formalizing and updating the evidence collected during the game consists of using matrices of probabilities, referred to as *Evidence Table*s (ETs), to store hard and soft evidence about the cards of the two ICP adversaries. Both hard and soft evidence about suspect, weapon, and room cards is presented in the form of probability distributions $Q(\mu_K)$, $Q(\mu_W)$, and $Q(\mu_R)$, respectively, and used in (7) and (8) to infer the three hidden cards. Let $\mu_K(m)$, $\mu_W(m)$ and $\mu_R(m)$ represent the instantiation of the m^{th} child of K, W and R, respectively. Then, the evidence table for the cards of type T, denoted by E_T, is defined in terms of the probability distribution $Q(\mu_T)$,

$$E_T(m,n) = Q(\mu_T(m) = T_n), \quad (13)$$

where T represents K, W or R, $m = (1, \ldots, |\mu_T|)$, and $n = (1, \ldots, |dom(T)|)$. In Fig. 3, $\mu_K = \{C_{12}^{(\mathrm{k})}, C_{22}^{(\mathrm{k})}, C_{13}^{(\mathrm{k})}\}$, $\mu_W = \{C_{32}^{(\mathrm{w})}, C_{23}^{(\mathrm{w})}, C_{33}^{(\mathrm{w})}\}$ and $\mu_R = \{C_{42}^{(\mathrm{r})}, C_{52}^{(\mathrm{r})}, C_{62}^{(\mathrm{r})}, C_{43}^{(\mathrm{r})}, C_{53}^{(\mathrm{r})}, C_{63}^{(\mathrm{r})}\}$. Thus, the dimensions of E_K, E_W and E_R are $|\mu_K|\times|dom(K)|$, $|\mu_W|\times|dom(W)|$ and $|\mu_R|\times|dom(R)|$, respectively.

Hard and soft evidence are stored in the ETs in the same fashion, but differ in their updating. Hard evidence represents perfect knowledge of an adversary's card. However, cards are shown only to the player making the suggestions, and the other players can gather information about their adversaries' cards from the suggestions themselves. This type of observation is referred to as soft evidence.

There are two ways for the ICP to obtain soft evidence. The first situation arises when the ICP makes a suggestion $\lambda = \{K_{s1}, W_{s2}, R_{s3}\}$ that cannot be disproved by P_d (P_2 or P_3), or both. This is referred to as Case (*a*), and indicates that the probabilities for the instantiations of λ are equal to 0. Case (*b*) refers to the situation where P_d shows another player his/her card, C_{show}, in order to disprove one element of the suggestion λ.

Before any evidence is entered, the ETs reflect uniform distributions for all cards, and are initialized as follows,

$$E_T(m,n) = 1/|dom(T)|, \quad (14)$$

where $T = K$, W or R, $m = (1, \ldots, |\mu_T|)$, and $n = (1, \ldots, |dom(T)|)$. The ETs are updated after each player's turn based on the evidence collected from his/her suggestion and the corresponding replies, according to the following rules:

1) The probabilities of mutually-exclusive instantiations for a variable must sum to one, i.e.,

$$\sum_n E_T(m,n) = 1, \quad (15)$$

where $T = K$, W or R, $m = (1, \ldots, |\mu_T|)$, and $n = (1, \ldots, |dom(T)|)$.

2) Hard evidence supporting the negation of an instatiantion holds for the remainder of the game, i.e.

$$\text{if } \mu_T(m) \neq T_n, \; E_T(m,n) = 0, \; \forall t > t_i, \quad (16)$$

where $T = K$, W or R, and t_i is the present turn.

3) Hard evidence supporting an instantiation holds for the remainder of the game, i.e.

$$\text{if } \mu_T(o) = T_q, \text{ then } \begin{cases} E_T(o,n)=0 \\ E_T(m,q)=0, \; \forall t > t_i \\ E_T(o,q)=1 \end{cases}, \quad (17)$$

where $T = K$, W or R, $m = (1, \ldots, |\mu_T|)$ and $m \neq o$, $n = (1, \ldots, |dom(T)|)$ with $n \neq q$, and t_i is the present turn.

The ETs are updated by first selecting the elements to be updated upon a player's move (step 1), and then by normalizing the corresponding rows and columns (step 2).

C. Evidence Tables Updating

Two normalization procedures called row normalization and column normalization are developed to guarantee that the laws of probability apply to the updated ETs.

1) Step 1 of ETs Update

In Fig. 3, the child nodes of type T belonging to P_d are denoted as $\mu_T(P_d)$, where $P_d = P_2$ or P_3. For example, $\mu_K(P_2) = \{C_{12}^{(\mathrm{k})}, C_{22}^{(\mathrm{k})}\}$. During the game of CLUE® hard evidence is obtained only when the ICP suggestion is disproved by a player, say P_d. When P_d shows a card of type T to the ICP, the row with the highest uncertainty among $\mu_T(P_d)$, denoted as the s^{th} row in E_T, is selected for updating and set equal to one (Rule 3). The Shannon's entropy [10] is used to quantify the uncertainty of the discrete probability distributions contained in the ETs. For an ET of type T this entropy is,

$$H_T(m) = -\sum_n E_T(m,n)\log_2(E_T(m,n)), \quad (18)$$

where, $T = K$, W or R, $m = (1, \ldots, |\mu_T|)$, and $n = (1, \ldots, |dom(T)|)$.

In Case (*a*), presented in Section III. B., player P_d has none of member cards in the suggestion of λ. Only those elements corresponding to the instantiation values in λ but not held by ICP, i.e., the intersections of λ and $dom(K) \cup dom(W) \cup dom(R)$, are updated. Denote these intersections as $\alpha = \{T_{h(v)}: 1 \leq v \leq |\alpha|\}$. $T_{h(v)}$ means the instantiated value is the $h(v)^{\text{th}}$ member in $dom(T)$. For any $T_{h(v)}$ in α, determine the row in E_T corresponding to every member in $\mu_T(P_d)$ to be $s(v)^{\text{th}}$ row, and then set the appropriate element $E_T(s(v),h(v)) = 0$, showing that any T card held by player P_d is not $T_{h(v)}$.

In Case (*b*), player P_d shows another one of his/her cards, C_{show}, to disprove one element of the suggestion $\lambda = \{K_{s1}, W_{s2}, R_{s3}\}$. The possible instantiations of C_{show} are the *free* instantiations that are also the intersections of λ and $dom(K) \cup dom(W) \cup dom(R)$, denoted by $\alpha = \{T_{h(v)}: 1 \leq v \leq |\alpha|\}$:

$$p(C_{\text{show}} = T_{h(v)}) = 1/|\alpha|. \quad (19)$$

For any $T_{h(v)}$ in α, determine the highest uncertainty among $\mu_T(P_d)$, say the s^{th} row in E_T, and then select the appropriate element to be $E_T(s,h(v))$, such that,

$$E_T(s,h(v)) = Q(\mu_T(s) = T_{h(v)})$$
$$= p(\mu_T(s) = T_{h(v)} \mid C_{\text{show}} = T_{h(v)})p(C_{\text{show}} = T_{h(v)}) \quad (20)$$
$$+ p(\mu_T(s) = T_{h(v)} \mid C_{\text{show}} \neq T_{h(v)})p(C_{\text{show}} \neq T_{h(v)})$$

$$p(\mu_T(s) = T_{h(v)} \mid C_{\text{show}} = T_{h(v)}) = 1 \quad (21)$$

$$p(C_{\text{show}} \neq T_{h(v)}) = 1 - p(C_{\text{show}} = T_{h(v)}) \quad (22)$$

$$p(\mu_T(s) = T_{h(v)} \mid C_{\text{show}} \neq T_{h(v)})$$
$$\approx \text{previous} \quad p(\mu_T(s) = T_{h(v)}) \quad (23)$$
$$= \text{previous} \quad E_T(s,h(v))$$

where $1 \le v \le |\alpha|$. Under the condition that C_{show} is $T_{h(v)}$, the statement that $\mu_T(s)$ equals the value of $T_{h(v)}$ is always true, as is shown in (21).

2) Step 2 of ETs Update

The procedure of row normalization has two inputs: a row in ET, say $E_T(s,:)$ and the last set element to be updated, say $E_T(s,h)$. Row normalization procedure first finds "1" and "0" elements in $E_T(s,:)$. Then, if any element is equal to 1, the other elements are set equal to 0; if none of the element equals 1, calculate the variation $\Delta = 1 - \sum_n E_T(s,n)$ and divide Δ by those elements that are not equal to 0 nor $E_T(s,h)$. The negative element(s) in $E_T(s,:)$ is set equal to 0. The above procedure is used again until the sum of $E_T(s,:)$ equals 1 and all the elements in $E_T(s,:)$ are non-negative. It can be seen that row normalization guarantees that ET-updating Rules (1), (2) and the first part of Rule (3) are always satisfied.

The procedure of column normalization first finds "1" elements in ETs. If there is any element of value 1, say $E_T(s,q)$, in E_T, then the other elements along the same column in the same matrix, (i.e, $E_T(m,q)$ for all $m \neq s$) are set to 0. This procedure guarantees that the second part of rule (3) is always satisfied.

IV. CLUE® Simulation

The simulation of CLUE® is developed through MATLAB Graphical User Interface (GUI) toolbox and includes three steps, choose players, deal cards, and start game.

A. Choose Players and Deal Cards

In the simulation of CLUE®, the ICP plays against two other players who are either human or random computer players. The human player(s) can choose three pawns to be computerized or human, among which is at least one computer player (pawn). In the sequence Col. Mustard, Miss Scarlet, Prof. Plum, Mr. Green, Mrs. White, Mrs. Peacock, the first computer player (pawn) is always the ICP. The other computer player, if any, is random and performs random moves and suggestions. Otherwise, the ICP plays against two human players.

The CLUE® cards are dealt through the rules described in Section III. A. Subsequently, the CLUE® BN model for the ICP is created based on the remaining cards or available instantiations. The BN CPTs are computed as described in Section III. A., and the ETs are initialized according to (14).

B. Start and Play Game

The necessary steps for a player of CLUE® include: 1) roll the die and get a random number from 2 to 12 to determine how far the pawn can move, 2) move to a position on the board with the number of passing bins no greater than die number, 3) transfer the turn to next player. If a player enters a room during his/her move, the player makes a suggestion and waits for other players' responses. If the player is taken into a room by another player, he/she must wait for his/her turn to make a suggestion. An accusation can be made at the beginning or at the end of players' turn. Secret passages to other rooms are located in the corner rooms, as in the real boardgame. The human players can see their CLUE® cards, and their actions are observed by the computer and the human players during the game via the interfaces developed in MATLAB.

At the beginning of every turn, the ICP infers posterior probabilities $p(K|se)$, $p(W|se)$ and $p(R|se)$, based on (7) and (8). If the maximum values $max(p(K|se))$, $max(p(W|se))$ and $max(p(R|se))$ are all $\geq$ 90%, then the ICP makes an accusation $\{K_{\text{estimated}}, W_{\text{estimated}}, R_{\text{estimated}}\}$. Otherwise, the ICP rolls the die and moves its pawn. The movement rule is that if $max(p(R|se)) \geq 90\%$, then ICP enters the room $R_{\text{estimated}}$ of $p(R|se) \geq 90\%$ as often as possible and makes suggestions on suspect and weapon cards; if $max(p(R|se)) < 90\%$, then ICP enters the nearest room where $0 < p(R|se) < 90\%$ as often as possible and makes suggestions as often as possible. Whenever the ICP enters a room, he makes a suggestion based on the following rule: if there is one suspect card K_s of $0 < p(K = K_s|se) < 90\%$, the ICP chooses K_s as the suggested suspect card, otherwise, the ICP chooses a suspect weapon card in hand as the suggested card; if there is one weapon card W_s of $0 < p(W = W_s|se) < 90\%$, the ICP chooses W_s as the suggested weapon card, otherwise, the ICP chooses a weapon card in hand as the suggested card. After making a suggestion, the ICP observes the response(s) of other player(s) and updates his ETs accordingly. At the end of his turn, the ICP makes an accusation if $max(p(K|se))$, $max(p(W|se))$, $max(p(R|se))$ are all $\geq$ 90%, otherwise play continues.

The user-interface allows the human players to manually select the suggestion. He/she also can make an accusation. If it is incorrect, he/she loses the game and will not move further but still can prove or disprove others' suggestions. To analyze CLUE® playing results, history record files are written to store useful information in CLUE® playing, such as information of players, ETs, posterior probabilities.

V. Results and Discussion

The CLUE® simulation begins with the simplified CLUE® BN model and the standard CLUE® rules. Then, the ETs are continually updated, allowing the ICP to follow the rules for movement and suggestion-making and to intelligently move about the board and make suggestions. Through many simulations involving the ICP playing against a random computer player and a human, ICP's intelligent suggestion-making is rationalized via game playing statistics.

Two testers, tester 1 and tester 2, respectively, were invited to test the CLUE® simulation. Tester 1 is a 15 year old 10th grader at Chapel Hill High School, Chapel Hill, North Carolina, USA. Tester 2 is a 12 year old 7th grader at Grey Culbreth Middle School, Chapel Hill, North Carolina, USA. Both testers have more than 2 years experience playing the game of CLUE®. A simulation begins with one tester choosing three of the six colored pawns, which represent a human player (the respective tester), an ICP, and a random computer player. The game is then played according to the rules outlined in Section II. A. and repeated until ICP or the tester wins the game. The test results are shown in Table I.

TABLE I
TEST RESULTS OF CLUE® SIMULATION

	Tester 1	Tester 2
Total Times of Playing	14	12
Times human player wins	9	4
Times ICP wins	5	8
Rate ICP wins	35.7%	66.7%

It can be seen from Table I that the ICP winning rate is lower for tester 1 (35.7%) than tester 2 (66.7%). The reason for the change in the rate ICP wins is related to the additional experience in playing the game of CLUE® of tester 1 over tester 2. The initial results about ICP winning rates are inspiring and show that ICP can defeat human players.

VI. CONCLUSION AND FUTURE WORK

A BN approach is developed to derive an intelligent computer player for the board game of CLUE®, whose objective is to correctly infer the three hidden cards, representing the suspect, weapon, and room of the murder. Evidence tables are constructed to update hard and soft evidence available during the game playing and incorporated by the proposed Bayesian network model. The game of CLUE® is used as a research benchmark to develop intelligent technologies for surveillance systems. Future work will implement ICP's optimal movement using finite state machine (FSM) approach [11] and preposterior analysis technique [12].

ACKNOWLEDGMENT

The authors would like to thank Ian and Nick Albertson for their contributions to testing the ICP, and Kelli Crews Baumgartner for her helpful comments. CLUE® & ©2006 Hasbro, Inc. Used with permission.

REFERENCES

[1] J. MacDonald, et al., *Alternatives for Landmine Detection*. RAND, Science and Technology Policy Institute, 2003.
[2] R. Popoli, "The Sensor Management Imperative," Chapter 10 in *Multitarget-Multisensor Tracking: Advanced Applications*, Vol. II, Y. Bar-Shalom, Ed., Artech House, 1992.
[3] S. Ferrari and A. Vaghi, "Sensor modeling and feature-level fusion by Bayesian networks," *Journal of Smart Structures and Systems*, vol. 1, no. 1, pp1-9, Nov. 2004
[4] J.E. Laird, "Using a Computer Game to Develop Advanced AI," *Computer*, vol.. 34, no. 7, pp. 70-75, July 2001.
[5] F. V. Jensen, "Bayesian Networks and Decision Graphs". Springer. 2001, ch. 1.
[6] R. Cowell, "Introduction to inference for Bayesian networks," in *Learning in Graphical Models*, M. Jordan, Ed., 1998, pp. 9-26
[7] J. Pearl, "Jeffery's rule, passage of experience, and neo-Bayesianism," In H.E. et al. Kyburg, Jr., editor, *Knowledge Representation and Defeasible Reasoning*, Kluwer Academic Publishers, 1990, pp. 245-265.
[8] M. Valtorta, Y. Kim and J. Vomlel, "Soft evidential update for probabilistic multiagent systems", *International Journal Approximate Reasoning*, 29(1) 71-106, 2002.
[9] G.F. Cooper, "The computational complexity of probabilistic inference using Bayesian belief networks," *Artificial Intelligence*, vol. 42, issues 2-3, pp. 393-405, March 1990,
[10] C. E. Shannon, "A mathematical theory of communication," *Bell System Tech. J.* 27, 379-423, 623-656, 1948.
[11] D. Johnson and J. Wiles, "Computer Games With Intelligence," 10th IEEE International Fuzzy Systems Conference, The University of Melbourne, Australia, pp. 1356-1358, December, 2001.
[12] R. L. Winkler, *An Introduction to Bayesian Inference and Decision*, 2nd ed. Gainesville, FL: Probabilistic Publishing, 2003, ch. 6.
[13] K. Murphy, *How to Use Bayes Net Toolbox*. [Online]. Available: http://www.cs.ubc.ca/~murphyk/Software/BNT/bnt.html, 2004.