A Multi-State Bayesian Network for Shill Verification in Online Auctions*

Ankit Goel[1], Haiping Xu[1] and Sol M. Shatz[2]
[1]Computer and Information Science Department
University of Massachusetts Dartmouth, North Dartmouth, MA 02747, USA
{agoel, hxu}@umassd.edu
[2]Computer Science Department
University of Illinois at Chicago, Chicago, IL 60607, USA
shatz@uic.edu

Abstract. *Online auction systems have made remarkable progress in recent years. However, one of the most severe and persistent problems in such systems is shilling behavior, which is a type of auction fraud where a bidder artificially drives up the bidding price so that the winner of the auction has to pay more than he otherwise would pay. Verification of shill bidders in an online auction is difficult due to incomplete knowledge about suspicious bidders. In this paper, we introduce a novel approach for verifying shill bidders using a multi-state Bayesian network, which supports reasoning under uncertainty. We describe how to construct the multi-state Bayesian network and present formulas for calculating the probabilities of a bidder being a shill and being a normal bidder. To illustrate the effectiveness of our approach, we provide a case study for shill verification, and demonstrate that a multi-state Bayesian network performs better than a bi-state Bayesian network.*

1. Introduction

Online auctions have become an integral part of e-commerce. An online auction system provides a platform for people from different walks of life and geographic locations to come together for the purpose of exchange of services, goods or money. In the spirit of the auction, the highest bidder becomes the winner of the auctioned item. As the number of people participating in online auctions increases, online auction systems are experiencing an immense volume of auction-based trading, as well as the problems associated with this volume of traffic.

Despite the popularity of online auctions in recent years, one of the biggest problems in online auctions, namely *shill bidding* [1-3], remains unchecked. Shill bidding is a bidding activity that artificially increases an auctioned item's price or apparent desirability. To engage in shill bidding, a seller might recruit a fake bidder, or create a pseudo-bidder account, to place bids that are solely intended to raise the price of an auctioned item. In this case, an honest bidder may become a victim of shill bidding and not even be aware that such activity has occurred. As a concrete example, consider auction *A* held by seller *S*. If the current price of the auctioned item is below the expectation of the seller, *S* can employ bidder *B* to bid and raise the price. Once the auction price has reached a certain satisfactory value, *B* stops bidding to avoid accidentally winning the auction.

Shill bidding is not a new phenomenon in the domain of online auctions; however, the development of effective ways to detect and verify shill bidding in online auctions is still an open problem [3]. Since both a seller and a buyer can be anonymous and managed by a single person with multiple accounts, it is very hard to discover the actual relationship between sellers and bidders. Furthermore, the actions of a shill bidder are seemingly close to normal bidding, which make it difficult to differentiate between such a shill bidder and a normal bidder. To safeguard the interests of legitimate users, researchers have attempted to identify various shill bidding patterns [4-6]. Meanwhile, many popular auction systems have implemented techniques aimed at curbing shilling activity. For example, eBay has a reputation point system where each buyer or seller has a reputation score called a *feedback score*. This score reflects a user's trustworthiness as evaluated by other users with whom this user has traded. However, most of the methods deployed so far cannot establish a shill bidder due to a lack of complete information about the bidder, thus allowing guilty parties to go unpunished.

In this paper, we propose a Bayesian network based system that can handle incomplete knowledge and help verify whether a shill suspect is an actual shill. Figure 1 shows a framework for shill detection and verification. Based on online auction data and shill bidding patterns, we can detect shill suspects using existing approaches, such as the data mining method [7] and real-time model

* This material is based upon work supported by the U.S. National Science Foundation under grant numbers CNS-0715648 and CNS-0715657.

checking mechanism [8]. However, there is no guarantee that a shill suspect is an actual shill because it is possible for a normal bidder to coincidentally demonstrate some shilling behaviors, although the bidder has no intention at all to be a shill. In order to verify if a shill suspect is an actual shill, we must use additional evidence to reason about the shill suspect's possibility for being a shill or a normal bidder. Our approach employs a multi-state Bayesian network as a verification engine. Once an actual shill is confirmed, the involved auction would likely be cancelled in order to protect other users' interests. Note that although shill detection is an important component of our proposed framework, details about detection mechanisms are beyond the scope of this paper.

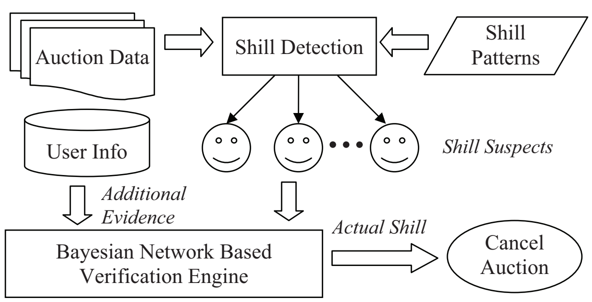

Figure 1. Shill detection and verification framework

2. Related Work

Most of the previous work related to auction fraud focuses on detection of in-auction fraud, which happens while transactions are in progress. Such auction fraud occurs disguisedly, leaving behind no obvious evidence; thus the victims typically do not even realize that the fraud has taken place. In order to effectively detect shill bidding, which is a major type of in-auction fraud, researchers have summarized various shill and normal bidding patterns [4-6]. Such bidding patterns are important knowledge for detection of shill suspects; however, they are not sufficient for justifying the presence of shills because some shill bidding patterns can be considered as normal behaviors under certain circumstance. For example, when an auctioned item is a rare and very hard to find item, a bidder may place very high bids in order to win the auction since the opportunity to bid on such an item might not present itself in the future. In this case, the bidder's bidding behavior may match with some shill bidding patterns; however, the bidder has no intention at all in being a shill.

Other researchers have utilized statistical and data mining techniques to detect abnormal bidding behaviors. Kauffman and Wood used a statistical approach to examining reserve price shilling behaviors and presented the factors that lead to this behavior [9]. They also showed how to use an empirical model to test for questionable and opportunistic bidding behaviors. Chau and Faloutsos proposed a data mining method to detect auction fraud by extracting characteristic features from exposed fraudsters [10]. They determined the significant features related to auction fraud by analyzing the fraudsters' transaction histories, which exist as graphs. More recently, Xu, et. al proposed a real-time model checking approach to detect shill suspects in auctions that are in progress, or "live" [8]. This approach introduced a dynamic auction model (DAM) that can be used to detect shilling behaviors formally specified using linear temporal logic (LTL). Although the above approaches can be effective in detecting shill suspects in online auctions, they are not sufficient for determining whether a shill suspect is an actual shill. In contrast, the focus of this paper is to introduce a verification engine that supports reasoning under uncertainty for shill verification. As such, this work is complementary to other research efforts for detection of shill suspects.

Previous work on shill verification is rare. Dong, et. al used Dempster-Shafer (D-S) theory to verify whether a user is a shill using additional evidence from both auction data and user information [11]. This is different from our approach because in our approach, the evidence set is based solely on user information. Note that using our shill detection and verification framework, the auction data has already been used in the first stage – for detection of shill suspects, we do not use the same information again in the second stage – for shill verification. Therefore, our approach requires less computation, and thus it is more efficient for shill verification than the D-S based method.

3. Bayesian Network Based Verification Engine

3.1 Bayesian Network for Shill Verification

Bayesian network (BN) or belief network is a probabilistic graph model that can be used to capture uncertain knowledge in a natural and efficient way [12]. BN models the dependencies among variables and gives a concise specification of full joint probability distribution. BN is a directed acyclic graph (DAG) where each variable is denoted by a node, and the probability of a node is conditionally dependent on its parent node(s). The nodes are selected from the same domain, such that they help in the decision making process. It is vital that these nodes represent major features in the domain, which contribute factors in decision making or influencing the variables that affect the major calculation.

The basic task for a probabilistic inference system using BN is to calculate the posterior probability for a query variable X, given a set of observed evidence $\mathbf{e}$ for a set of evidence variables $\mathbf{E}$. This is done by summing terms of the full joint distribution as follows [12]:

$$\mathbf{P}(X|\mathbf{e}) = \alpha\,\mathbf{P}(X,\mathbf{e}) = \alpha \sum_{\mathbf{y}} \mathbf{P}(X,\mathbf{e},\mathbf{y}) \tag{1}$$

where X is the query variable, $\mathbf{e}$ is the observed values for a set of evidence variables **E**, **y** is the values for a set of unobserved (or hidden) variables **Y**, and α is the normalization factor. Note that **P**(X, **e**, **y**) is the full joint probability of X, **e** and **y**.

Figure 2 shows a BN with three layers that we designed for shill verification. The first layer defines three groups of evidence nodes, namely evidence nodes for determining the strength of bidder's evidence (N1-N4), evidence nodes for determining the strength of seller's evidence (N7-N9), and evidence nodes for determining the interaction strength of the involved bidder and seller (N5, N6). Nodes N1-N9 represent evidence sources that are observable in the context of eBay. To simplify the network, we define three unobservable nodes (N10-N12) in the second layer, whose parents are the aforementioned three groups of nodes, respectively. The values of the three layer-two nodes (N10-N12) can be used to determine the conditional probabilities of being a shill or a normal bidder, which are defined as nodes *Shill* (N13) and *Normal* (N14) in the third layer.

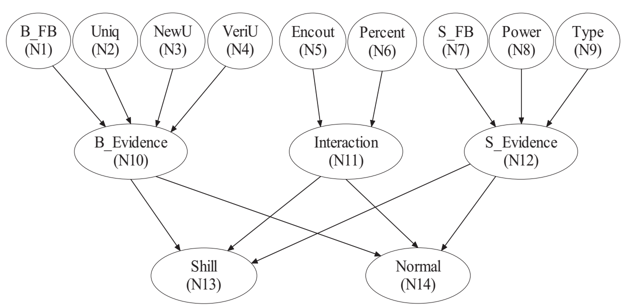

Figure 2. Bayesian network for shill verification

We now explain each of the nodes (N1-N14) defined in Figure 2.

Bidder Feedback (N1) represents a range of feedback score values for a particular bidder. In eBay, feedback scores for bidders are represented by different colored stars. For example, a *yellow* star represents that the bidder has a feedback score in the range of 10 to 49. A reputation mechanism, such as feedback stars, serves as a deterrent for bad user service and illegal actions; it can also be used to gauge the trustworthiness of a user. Obviously, a bidder with a high score has been in the system longer and is trusted by many people. Since it can take a long time and serious effort for a user to obtain a high feedback score, a legitimate bidder is not likely to risk sacrificing this score lightly by engaging in shill bidding.

Number of Unique Sellers (N2) is the number of different sellers whose auctions the bidder has participated in over the past 30 days. Most shill bidders frequently trade with the same seller or a set of sellers. Knowing the number of different sellers with whom the bidder has interacted, helps us understand the bidder's general activity with the auction community.

New User (N3) is a person who registered recently, possibly within the last 30 days. EBay assigns a special icon to such users. A new user has little or no associated history and has the least to lose if banned for fraud. Thus, new users are more likely to appear on a suspect list than experienced users. Note that since a user can have multiple accounts, it is easy for a user to create a new account for shill bidding.

Verified User (N4) is a registered user who has provided eBay with further proof of identity. To become a verified user with eBay, the user must provide personal information (e.g., home address and birth date), which is processed through a third party company such as Equifax. Since a verified user's personal information is recorded, such a person typically would not risk shill bidding and accidentally being caught.

Encounters (N5) is the number of auctions held by seller *S*, in which the bidder placed a bid in the past 30 days. Although a high number of encounters may imply that the bidder simply prefers seller *S*, it is also possible that the bidder is serving as a shill bidder for *S*.

Bid Percentage (N6) is the percentage of bids placed by the bidder in auctions held by seller *S* over the total number of bids placed by the bidder in the past 30 days. Although a high bid percentage may be due to the competitiveness of the auctions held by seller *S*, it is also an indicator of shill bidding because it implies a very close relationship between the bidder and seller *S*.

Seller Feedback (N7) is the feedback assigned to a seller. Again, eBay uses different colored stars for different feedback score ranges. Similar to the bidder feedback star, a seller with a good feedback score typically does not want to risk loosing that score by being involved in shill bidding. Note that a high feedback score is critical to a seller because most bidders prefer to buy items from a trustworthy seller, even at the cost of paying more money for the auctioned item, as it guarantees them proper and timely delivery.

Power Seller (N8) is a status assigned by eBay to some sellers who consistently sell a significant volume of items, maintain a 98% positive feedback rating, and provide a high level of service to their buyers. There are multiple levels of power seller like *Gold*, *Silver* and *Bronze*. Obviously, being a power seller is a positive indicator that that seller is not involved in shill bidding.

Seller Type (N9) of a seller can either be *private* or *store* at eBay. To be an eBay store, the seller must have an eBay seller account with credit card information on file. The seller must also be a verified user or have a feedback score of at least 20. So a "store seller" tends to be trusted by bidders, and is typically not likely to be involved in shill bidding.

Bidder Evidence Level (N10) is evaluated based on the bidder's evidence information from its parent nodes. A high bidder-evidence (strength) level implies the bidder is more likely an actual shill.

Interaction Level (N11) represents the strength of interaction between the bidder and seller in question. A high level of interaction implies a preferred seller or an unhealthy relationship between the two parties.

Seller Evidence Level (N12) considers the probability that the seller is involved in shill bidding. Note that if we implicate a bidder in an auction, the seller of the same auction must also be implicated. Similar to bidder evidence level, a seller is judged based on the seller's evidence information from its parent nodes.

Shill Bidder (N13) indicates whether a bidder suspect is an actual shill given additional evidence. This node represents the prior probability that the bidder is an actual shill based on the states and probabilities of its parent nodes, namely N10-N12.

Normal Bidder (N14) indicates whether a bidder suspect is actually a normal bidder given additional evidence. This node represents the prior probability that the bidder is actually a normal bidder based on the states and probabilities of its parent nodes, namely N10-N12.

Note that we consider the two nodes N13 and N14 as conditionally independent. Thus given a set of evidence **e** (i.e., values of evidence variables N1-N9), the summation of probabilities of being a shill and a normal bidder, P(*shill*|**e**) + P(*normal*|**e**), does not necessarily equal to 1.

3.2 Multi-State Bayesian Network

In a BN, each node typically takes only one of two values, such as *true* and *false*. This type of BN is called a *bi-state* BN. The second column in Table 1 shows the two states for each of the nodes we are considering. For example, node N1 in a bi-state BN can be in a state of *low* or *high*. N1 has the value of *low* when the bidder feedback star is *Turquoise* or below; otherwise, N1 has the value of *high*. Note that the values corresponding to each state are determined based on the actual data distribution of the original information from eBay. In other words, the states are assigned based on how shills and normal bidders are clustered, but only through approximation and observation.

In order to make the BN reflect the dependencies among different nodes more precisely, we consider a *multi-state* BN, where the nodes are not limited to two states. The third column of Table 1 shows the use of multiple states for some nodes. For example, node N1 in a multi-state BN can be in one of the following states *None*, *Yellow*, *Blue*, *Turquoise*, or *Other*, which correspond to different levels of a bidder feedback star. In our case study, we demonstrate that a multi-state BN performs better than a bi-state BN for shill verification.

Table 1. State values for nodes N1-N9

Node	Bi State	Multi State
N1	Low/High	None/Yellow/Blue/Turquoise/Other
N2	Low/High	1 / 2-5 / 6-15 /Other
N3	True/False	True/False
N4	True/False	True/False
N5	Low/High	Low/High
N6	Low/High	No more than 30 / 31-80 / Other
N7	Low/High	None/Yellow/Blue/Turquoise/Other
N8	Low/High	None/Bronze/Other
N9	Store/Private	Store/Private
N10	Low/High	Low/High
N11	Weak/Strong	Weak/Strong
N12	Low/High	Low/High
N13	Yes/No	Yes/No
N14	Yes/No	Yes/No

3.3 Conditional Probability Table

The prior knowledge of the BN can be derived from eBay auction data as well as user information. The auction data and user information were collected from eBay using the Trading APIs available for developers [13], which allows a developer to retrieve various types of data via different web service invocations. It is important to note that information retrieved about a bidder is limited; one can only obtain information that is available for the last 30 days before the time of data collection. This is in accordance with eBay's privacy policies, which prevents other users from acquiring unlimited information regarding a bidder in an auction. Furthermore, all bidders are named anonymously (e.g., *x***y*) – only a seller of an auction may see the real user identifications of the bidders participating in the auction.

We represent the prior knowledge of the BN as conditional probability tables (CPT) associated with each node in the network. A CPT is similar to a truth table and shows the conditional probability of the node with respect to every state of its parent nodes. In order to complete the CPT for each node, we retrieved auction data for the query "laptop." Auction data for a total of 109 auctions with 1127 bidders was accumulated. With the help of an existing shill detection tool [8], we created a training data set by identifying shill suspects (bidders or sellers) and verifying them manually through investigation of their profiles as well as their shill patterns. In the following, we provide a few examples of shill bidding patterns that were useful for creating the training data set.

High and Irregular Incremental Bids: Shill bids have a tendency to be relatively high as compared to other bids placed by normal bidders. A normal bid usually has an increment of \$1 to \$5, or \$10 to \$50 for high-end auctions. If a bidder's bid increment is very high and irregular, the bidder is likely a shill.

Early and Middle Stage Bidding: Most shills only bid in the initial and middle stages of an auction, but stop bidding in the final stage to avoid winning the auction. On the other hand, an active bidder in the final stage of an auction is not likely a shill.

Successive Bidding: A normal bidder typically does not outbid himself. Therefore, any successive bidding activities that outbid oneself can be considered as a good indicator of shilling behavior.

Total Increase in Price: An interesting observation is the total increase in the price of an item due to a single bidder or multiple bidders who collude to raise the auction price. Thus, it is a good clue to shilling.

Once the training data set is ready, we can create a CPT for each node in both a bi-state BN and a multi-state BN. Note that to create CPTs for node N10 and N12, we manually identified bidders and sellers with high evidence level based on their suspiciousness. We also manually identified bidders and sellers with strong interaction strength for creating CPT for node N11. Table 2 shows part of the CPT for node N10 in the bi-state BN. For example, the second row of Table 2 tells that when the bidder's feedback score is *low* (i.e., *Turquoise* or below, for N1), and the number of unique sellers is also *low* (i.e., less than 10, for N2), and the bidder is a new user (N3), and the bidder is not a verified user (N4), then the conditional probability for N10 being *high* (i.e., high bidder evidence level) is 0.83. This is reasonable because an unverified new user with a low number of unique sellers and low feedback score is likely a shill (if the bidder has already been detected as a shill suspect). As another example, since there are no training examples in the category "N1 = H, N2 = L, N3 = T, N4 = F" (i.e., the third row in Table 2), the conditional probability for N10 being *high* is set as 0.00. This setting will have no impacts on shill verification results because a new user is not likely to have a high feedback score.

Table 2. Part of the CPT for node N10 in bi-state BN

VeriU (N4)	NewU (N3)	Uniq (N2)	B_FB (N1)	P(N10 = H)
F	T	L	L	0.83
F	T	L	H	0.00
F	T	H	L	0.15
F	T	H	H	0.00
F	F	L	L	0.86
F	F	L	H	0.15
F	F	H	L	0.30
F	F	H	H	0.10

Similarly, Table 3 shows part of the CPT for the node N10 in the multi-state BN. The CPT for node N10 in the multi-state BN provides a conditional probability for more specific state values of its parent nodes. For example, the second row in Table 3 indicates that when the bidder's feedback score is *None* (i.e., no feedback star for N1), and the number of unique sellers is 1 (N2), and the bidder is not a new user (N3), and the bidder is not a verified user (N4), then the conditional probability for N10 being *high* is 0.93. This is also reasonable because an unverified experienced user with very low feedback score, who only traded with the same seller of the auction in the past 30 days, is very likely a shill (again, if the bidder has already been detected as a shill suspect).

Table 3. Part of the CPT for node N10 in multi-state BN

VeriU (N4)	NewU (N3)	Uniq (N2)	B_FB (N1)	P(N10 = H)
F	F	1	None	0.93
F	F	1	Yellow	0.75
F	F	1	Blue & Torq	0.72
F	F	1	Other	0.10
F	F	<=5	None	0.89
F	F	<=5	Yellow	0.87
F	F	<=5	Blue & Torq	0.87
F	F	<=5	Other	0.10
F	F	<=15	None	0.30
F	F	<=15	Yellow	0.25
F	F	<=15	Blue & Torq	0.10
F	F	<=15	Other	0.10
F	F	Other	None	0.00
F	F	Other	Yellow	0.00
F	F	Other	Blue & Torq	0.00
F	F	Other	Other	0.00

4. Case Study

To illustrate the effectiveness of our proposed approach, we developed a Bayesian network toolkit (a snapshot of the toolkit for a multi-state BN is shown in Figure 3). The inputs of the toolkit are the BN for a particular auctioned item, and an auction with the same type of auctioned item, along with the user information for all involved bidders and seller. In this study, the auctioned item is of type "laptop," and the auction under investigation is "HP HDX 16t notebook" held from Oct-09-09 to Oct-16-09, as shown in Table 4.

The conditional probabilities for query variables *Shill* and *Normal*, given evidence **e**, can be calculated using Equations (2) and (3), respectively, which are derived from Equation (1) defined in Section 3.1.

$$\mathbf{P}(Shill \mid \mathbf{e}) = \alpha\, \mathbf{P}(Shill, \mathbf{e}) = \alpha \sum_{\mathbf{y}} \mathbf{P}(Shill, \mathbf{e}, \mathbf{y}) \quad (2)$$

$$\mathbf{P}(Normal \mid \mathbf{e}) = \alpha\, \mathbf{P}(Normal, \mathbf{e}) = \alpha \sum_{\mathbf{y}} \mathbf{P}(Normal, \mathbf{e}, \mathbf{y}) \quad (3)$$

In the above equations, *Shill* is the query variable N13, *Normal* is the query variable N14, **e** is the observed values for the set of evidence variables **E** = {N1, N2, ..., N9}, and **y** is the values for the set of unobserved variables **Y** = {N10, N11, N12, N14} in Equation (2), and **Y** = {N10, N11, N12, N13} in Equation (3), respectively.

Table 4. HP HDX 16t notebook bidding history

Bidder (FB)	Bid Amount	Bid Time	Bidder (FB)	Bid Amount	Bid Time
*3***3 (27)*	US $630.00	Oct-16-09 08:19:12 PDT	*9***a (1)*	US $200.00	Oct-10-09 12:40:10 PDT
*p***t (299)*	US $621.00	Oct-15-09 18:34:38 PDT	*g***e (245)*	US $112.50§	Oct-09-09 15:49:21 PDT
*p***t (299)*	US $611.00§	Oct-15-09 18:34:38 PDT	*9***a (1)*	US $110.00	Oct-10-09 12:39:56 PDT
*g***e (245)*	US $601.00	Oct-14-09 16:59:59 PDT	*g***e (245)*	US $102.50§	Oct-09-09 15:49:21 PDT
*p***t (299)*	US $601.00	Oct-15-09 17:54:10 PDT	*9***a (1)*	US $100.00	Oct-10-09 12:39:43 PDT
*g***e (245)*	US $561.00§	Oct-14-09 16:59:59 PDT	*g***e (245)*	US $61.00§	Oct-09-09 15:49:21 PDT
*p***t (299)*	US $551.00	Oct-12-09 17:09:44 PDT	*9***a (1)*	US $60.00	Oct-10-09 12:39:27 PDT
*8***1 (29)*	US $550.00	Oct-14-09 12:31:54 PDT	*g***e (245)*	US $51.00§	Oct-09-09 15:49:21 PDT
*p***t (299)*	US $510.00§	Oct-12-09 17:09:44 PDT	*9***a (1)*	US $50.00	Oct-10-09 12:39:04 PDT
*g***e (245)*	US $500.00	Oct-09-09 15:49:21 PDT	*g***e (245)*	US $31.00§	Oct-09-09 15:49:21 PDT
*m***s (4)*	US $500.00	Oct-10-09 16:49:22 PDT	*a***a (21)*	US $30.00	Oct-09-09 13:58:29 PDT
*g***e (245)*	US $405.00§	Oct-09-09 15:49:21 PDT	*g***e (245)*	US $26.00§	Oct-09-09 15:49:21 PDT
*m***s (4)*	US $400.00	Oct-10-09 16:49:00 PDT	*a***a (21)*	US $25.00	Oct-09-09 13:58:02 PDT
*g***e (245)*	US $202.50§	Oct-09-09 15:49:21 PDT	*a***a (21)*	US $0.99§	Oct-09-09 13:58:02 PDT
§ Automatic bid using eBay proxy bidding system.			*Starting Price*	US $0.99	Oct-09-09 08:20:59 PDT

In order to determine whether a shill suspect is an actual shill, we first define thresholds based on our experience for probabilities of being a shill and a normal bidder. Although the thresholds are subjectively defined, they are sufficient for our case study and can be improved later as we gain more experience with shill verification. The current thresholds are defined as follows:

$$\mathrm{P}(shill|\mathbf{e}) \geq 0.8 \text{ and } \mathrm{P}(normal|\mathbf{e}) < 0.30 \Rightarrow \text{Shill}$$

Now from the bidding history listed in Table 4, we detected three shill suspects, namely *g***e*, *m***s*, and *9***a*, all of which demonstrated some shill patterns. For example, bidder *m***s* has very high and irregular incremental bids, and only placed bids in the middle stage of the auction. The evidence we collected for the three shill suspects is listed in Table 5.

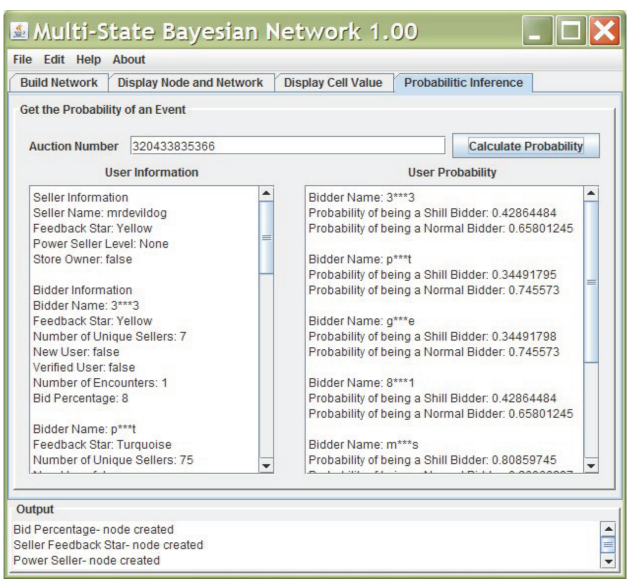

Figure 3. Tool support for multi-state BN

Table 5. Evidence of three shill suspects

Node	Bidder		
	ge***	***m***s***	***9***a***
N1	Turquoise	None	None
N2	14	1	2
N3	False	False	False
N4	False	False	False
N5	2	1	4
N6	6	100	45
N7	Yellow	Yellow	Yellow
N8	None	None	None
N9	Private	Private	Private

Based on the evidence data and the BN we developed for auctioned items of type "laptop," we used our Bayesian network toolkit to calculate P(*shill* | **e**) and P(*normal* | **e**) for the three shill suspects. The results for both the bi-state BN and the multi-state BN are listed in Table 6. From Table 6, we can see that using the bi-state BN, no shill suspects can be confirmed as actual shills, although bidder *m***s* had demonstrated very obvious shilling behaviors. In contrast, when using the multi-state BN, bidder *m***s* can be successfully identified as an actual shill. Thus, the experimental results of the case study conform to our expectation that a multi-state BN performs better than a bi-state BN.

Table 6. Probability of being a shill or a normal bidder

Probability of Being Shill / Normal Bidder		Bidder		
		ge***	***m***s***	***9***a***
Bi-State	P(*shill*|**e**)	0.4008	0.6115	0.6391
	P(*normal*|**e**)	0.5687	0.3052	0.2590
Multi-State	P(*shill*|**e**)	0.3449	<u>0.8086</u>	0.7234
	P(*normal*|**e**)	0.7456	<u>0.2607</u>	0.2454

In our case study, we also calculated P(*shill* | **e**) and P(*normal* | **e**) for other bidders in the auction (as shown in Figure 3). Our experiments show that no other bidders satisfy our defined requirements to be shill, as is expected. For example, bidder *3***3* is the winner of the auction, and is thus not likely to be a shill. This fact is consistent with the results demonstrated in Figure 3, where P(*shill* | **e**) = 0.4286 and P(*normal* | **e**) = 0.6580.

5. Conclusions and Future Work

Shill detection and verification are an imprecise art riddled with uncertainties. However, using a probabilistic inference system such as a BN, we can account for these uncertainties with some degree of belief and reach a decision. The BN gives us results for decision making based on prior knowledge about the domain and thus provides some level of quantification. The biggest dilemma we have when implicating a bidder is whether the bidder was simply "in the wrong place at the wrong time." With additional evidence for suspected bidders, our multi-state BN helps us resolve this uncertainty when implicating a bidder and the seller of the auction.

We also see from the results that a multi-state BN gives more accurate results as compared to a bi-state BN. This follows the fact that using multi-state BN, we can classify shills in fine-grained categories, allowing us to draw more precise conclusions.

For future work, we will implement a clustering algorithm like *k*-means clustering [14], which can help cluster large chunks of data and provide a clearer picture of the knowledge domain. This will help in the discovery of any hidden node states that can give a better result for identifying shills in online auctions. Furthermore, we also plan to incorporate our Bayesian network toolkit with an agent-based trustworthy online auction system we previously developed [15]. We believe once the above features are implemented, we will have a full-fledged, reliable verification engine, which can form the core of a trustworthy agent-based online auction system.

References

[1] B. K. Bhargava, M. Jenamani, and Y. H. Zhong, "Counteracting Shill Bidding in Online English Auction," *International Journal of Cooperative Information Systems*, Vol. 14, No. 2-3, 2005, pp. 245-263.

[2] D. H. Chau, S. Pandit, and C. Faloutsos, "Detecting Fraudulent Personalities in Networks of Online Auctioneers," *Proceedings of the 10th European Conference on Principles and Practice of Knowledge Discovery in Database (PKDD 2006)*, Berlin, Germany, September 2006, pp. 103-114.

[3] F. Dong, S. M. Shatz, and H. Xu, "Combating Online In-Auction Fraud: Clues, Techniques and Challenges," *Computer Science Review (CSR)*, Vol. 3, No. 4, November 2009, pp. 245-258.

[4] H. Xu and Y.-T. Cheng, "Model Checking Bidding Behaviors in Internet Concurrent Auctions," *International Journal of Computer Systems Science & Engineering*, 2007, Vol. 22, No. 4, pp. 179-191.

[5] H. S. Shah, N. R. Joshi, A. Sureka, and P. R. Wurman, "Mining eBay: Bidding Strategies and Shill Detection," *Proceedings of the 4th International Workshop on Mining Web Data for Discovering Usage Patterns and Profiles*, July 2003.

[6] J. Trevathan and W. Read, "Detecting Shill Bidding in Online English Auctions," *Handbook of Research on Social and Organizational Liabilities in Information Security*, M. Gupta and R. Sharman (eds.), Information Science Reference, 2009, pp. 446-470.

[7] S. Pandit, D.H. Chau, S. Wang, and C. Faloutsos, "Netprobe: a Fast and Scalable System for Fraud Detection in Online Auction Networks," *Proceedings of World Wide Web*, 2007, pp. 201-210.

[8] H. Xu, C. K. Bates, and S. M. Shatz, "Real-Time Model Checking for Shill Detection in Live Online Auctions," *Proceedings of the International Conference on Software Engineering Research and Practice (SERP'09)*, July 13-16, 2009, Las Vegas, Nevada, USA, pp. 134-140.

[9] R. J. Kauffman and C. A. Wood, "Running up the Bid: Detecting, Predicting, and Preventing Reserve Price Shilling in Online Auctions," *Proceedings of the 5th International Conference on Electronic Commerce*, 2003, Pittsburgh, PA, pp. 259-265.

[10] D. Chau and C. Faloutsos, "Fraud Detection in Electronic Auction," *Proceedings of European Web Mining Forum (EWMF 2005)*, ECML/PKDD 2005, October 2005, Porto, Portugal.

[11] F. Dong, S. M. Shatz, and H. Xu, "Reasoning Under Uncertainty for Shill Detection in Online Auctions Using Dempster-Shafer Theory," To appear in *International Journal of Software Engineering and Knowledge Engineering (IJSEKE)*, Vol. 20, No. 7, November 2010.

[12] S. Russel and P. Norvig, *Artificial Intelligence: a Modern Approach*, Second Edition, Prentice Hall, 2005.

[13] eBay, *Trading Web Services – eBay Developers Program*, eBay Inc., Retrieved on September 18, 2009 from http://developer.ebay.com/products/trading/

[14] D. J. Hand, H. Mannila, and P. Smyth, *Principles of Data Mining*, The MIT Press, Cambridge, Massachusetts, 2001.

[15] R. Patel, H. Xu, and A. Goel, "Real-Time Trust Management in Agent Based Online Auction Systems," *Proceedings of the 19th International Conference on Software Engineering and Knowledge Engineering (SEKE 2007)*, July 9-11, 2007, Boston, USA, pp. 244-250.