A Bayesian Network Analysis of Eyewitness Reliability: Part 2

Jack K. Horner
PO Box 266
Los Alamos NM 87544
jhorner@cybermesa.com

ICAI 2014

Abstract

In practice, many things can affect the verdict in a trial, including the testimony of eyewitnesses. Eyewitnesses are generally regarded as questionable sources of information in a trial setting: cases that turn on the testimony of a single eyewitness almost never result in a guilty verdict. Multiple eyewitnesses can, under some circumstances, collectively exhibit more robust behavior than any witness individually does. But how reliable, exactly, are multiple eyewitnesses? The legal literature on the subject tends to be qualitative. In a companion paper, I describe a highly idealized Bayesian network model of the relation between eyewitness behavior and trial verdict. In this paper, I describe a more refined Bayesian model of the same setting. It turns out that the highly idealized model provides nearly as much information as the more refined one does.

Keywords: eyewitness, Bayesian network

1.0 Introduction

In practice, many things can affect the verdict in a trial -- procedural conventions, material evidence, the psychology of the jurors, the persuasive power of the attorneys, and often, the testimony of eyewitnesses. Eyewitnesses are generally regarded as questionable sources of information in a trial setting ([3]): cases that turn on the testimony of a single eyewitness almost never result in a guilty verdict.

Multiple eyewitnesses can, under some circumstances, collectively exhibit more robust behavior than any witness individually does. But how reliable, exactly, are multiple eyewitnesses? The outcome of the recent trial of George Zimmerman, accused of second-degree murder or manslaughter of a teenager, rested heavily on the answer to this question ([10]). The legal literature on the subject tends to be qualitative (see, for example, [3]). A quantitative model is required.

Throughout, I will use the term *correct verdict* to mean a verdict that agrees with what actually happened, independently of the trial. I will use the term *verdict-determining-event* (VDE) to mean an event that could be witnessed by an eyewitness or that could contribute to a verdict.

2.0 A more refined Bayesian model

In a companion paper, I described a highly idealized model of a three-eyewitness domain. In many trials, witnesses observe only disjoint, or only partially overlapping, segments of an possible event timeline that is probabilistically related to a correct verdict. In the George Zimmerman trial ([10]), for example, witness saw or heard various, largely disjoint, parts of a hypothesized timeline of events.

To capture at least some of the more important features of cases of this kind, let's consider a model, *WitnessTimeline* (*WT*), in which:

(SC)

- There are three eyewitnesses
- There are three possible events which lie on a hypothesized timeline that occurred with high probability in the case of a correct a correct verdict. A possible event can occur or not occur.
- Each witness makes an observation of exactly one possible event. No two witnesses observed/didn't observe any one of the events. That is, there is a one-to-one onto mapping between witnesses and possible events. Any witness can be mistaken about whether he/she observed his/her respective event.
- If a possible event occurred, then the probability that the corresponding witness observed that possible event is 0.75. If a possible event didn't occur, then the probability that the corresponding witness observed that possible event is 0.25.
- If the hypothesized timeline occurred, then the probability that each of the possible events occurred is 0.9
- In the case of a correct verdict, the probability that the hypothesized timeline occurred is 0.9
- in the absence of any observation, the probability of a correct verdict is 0.5

Figure 1 shows a graphical user-view of more refined model, *Witness Timeline (WT*), that satisfies (SC) and is implemented in [1]. Each box in Figure 1 represents a random variable of a system. An arrow from a box A to a box B signifies that the distribution of the values of B depends probabilistically depends on the distribution of the values of A (and by Eq. 2, conversely). Thus, for example, in Figure 1 the probability that Witness3 correctly observed a VDE depends probabilistically on whether a correct verdict was delivered.

The prior probabilities of *STW* are defined in tables (not shown) to be the probabilities in the (SC).

Each box in Figure 1 has three regions, delimited by horizontal borders.

The top region of a box contains the name of a (random) variable of interest, e.g., *Correct Verdict.*

The middle region of a box consists of three elements (read horizontally):

i. a textual value-range for the random variable named in the top region of the box. For example, the top box in Figure 1 represents the random variable *Correct Verdict*.

ii. to the right of (i), a numerical literal (expressed as a percentage) indicating the probability that the variable of interest has a value lying in the value-range. For example, in Figure 1, the variable *Correct Verdict* has a probability of 96.4% of being true.

iii. to the right of (ii) a (segment of a) a histogram representation of the probability that the variable of interest has a value lying in the value-range denoted by (ii). Taken as a whole, the histogram spanning the middle region of the box represents the probability distribution for the variable named in (i), conditional on the variables at the tails of the arrows whose heads touch the box.

In Figure 1, the "Correct Verdict?" box has a pink background; the bottom row of boxes, a grey background. A box with a grey background means the variable corresponding to that box is intended as an "input" (also called an "asserted-value" or "finding") variable. Input variables represent information that is posited as given. A box with a pink background means the variable corresponding to that box is intended as an "output" (also called a "calculated") variable.

In *WT*, a variable can be toggled between a finding and a calculated value by a mouse-click.

Now suppose we knew nothing about what the witnesses observed. This situation is depicted Figure 5. As expected ([12], pp. 65-68), *WT* predicts that the probabilities of all possible states of the system are equal (0.5).

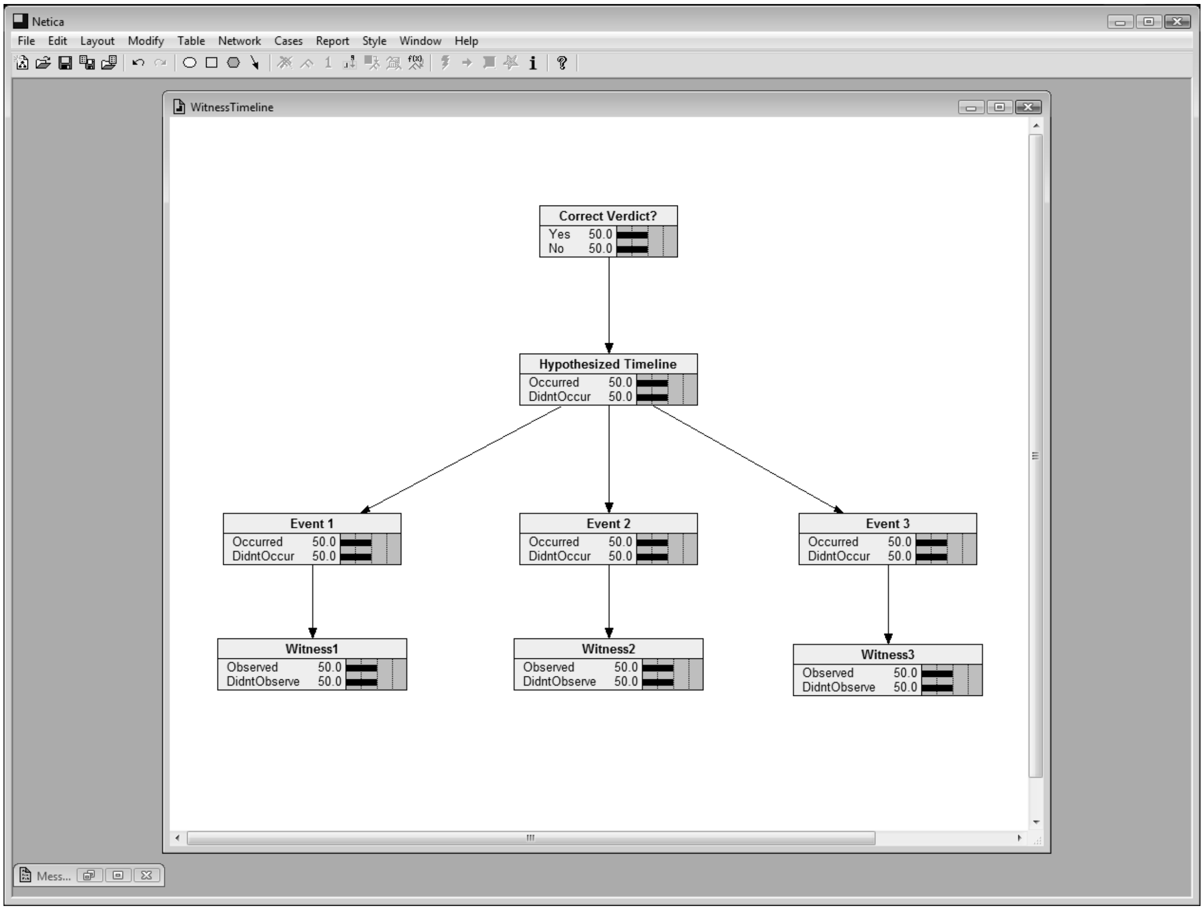

Figure 1. User-view of *WT*, given no data about the observations of the witnesses. In this case, all states of the network are equally probable.

Now let's suppose that all three witnesses observed their respective possible events and that the events in fact occurred. This situation is depicted in Figure 2. Then under the configuration described above, *WT* predicts that the probability of a correct verdict is ~0.84.

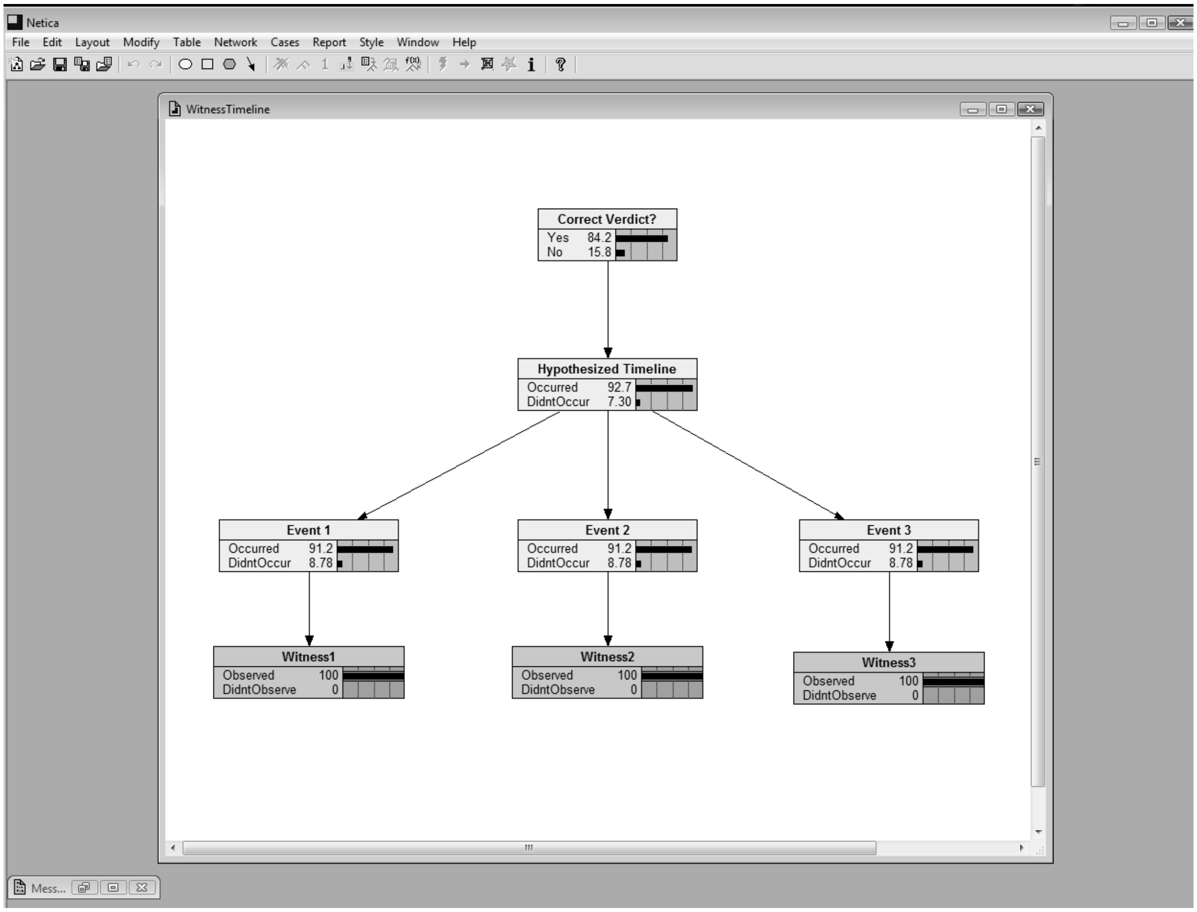

Figure 2. User-view of *WT*, a BN in which there is a one-to-one onto mapping between a set of three witnesses and a set of three events that lie on a hypothesized timeline. By assumption, the hypothesized timeline occurs with high probability in the case of correct verdict.

What happens to Correct Verdict in *WT* if one of the witnesses does not observed his/her event? This situation is depicted in Figure 7. *WT* predicts that, as expected, the probability of a correct verdict is lower than the configuration depicted in Figure 3.

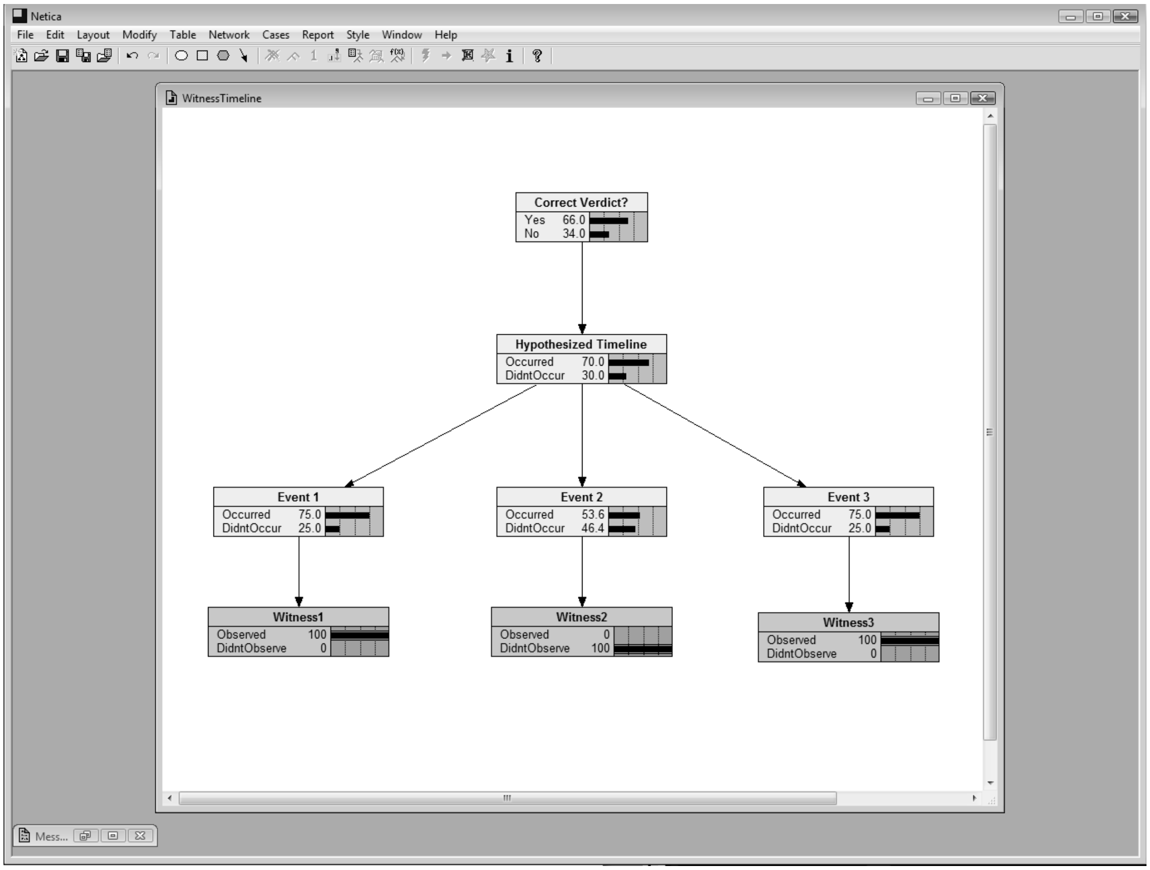

Figure 3. User-view *WT*, assuming two witnesses observe his/her possible event, but the third witness does not observe his/her event. The Correct Verdict occurs with probability 0.66.

3.0 Discussion

The analysis above motivates several observations:

1. The rubric shown here can be extend to an arbitrary number of witnesses, although the effect of more than three correctly witnesses, all other suffering being the same, contributes little.

2. The effect of one inaccurate witness is significantly mitigated by at least two accurate witnesses of a VDE. Adding the testimony of more than two accurate witnesses has decreasing returns. In addition, adding witnesses always runs the risk of introducing a witness whose testimony could raise doubt about the testimony of the rest. From the prosecution's point of view, this risk may not be negligible.

3. A companion paper describes a highly idealized model, *STW*, of the three-eyewitness domain. The current paper shows the predictions of *STW* are surprisingly informative. As expected, because *WT* has more probability distributions between the eyewitnesses and the verdict than *STW* does, whether a correct verdict is obtained depends less on the eyewitnesses in *WT* than in *STW*.

4.0 Acknowledgements

This work benefited from discussions with Tom Rudkin, Brent Boerlage, Tony Pawlicki, Ron Giere, and Wesley Salmon. For any errors that remain, I am solely responsible.

5.0 References

[1] Norsys Software, Inc. *Netica* v4.08. 2008. http://www.norsys.com.

[2] Pearl J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Revised Second Printing. Morgan Kaufmann. 1988.

[3] Steblay NK and Loftus EF. Eyewitness identification and the legal system. In Shafir E, ed. *The Behavioral Foundations of Public Policy*. Princeton. 2013.

[4] Chung KL. *A Course in Probability Theory*. Third Edition. Academic Press. 2001.

[5] Diestel R. *Graph Theory*. Springer. 1997.

[6] Jensen FV. *Bayesian Networks and Decision Graphs.* Springer. 2001.

[7] Pearl J. *Causality: Models, Reasoning, and Inference.* Second Edition. Cambridge. 2009.

[8] Jensen FV. *An Introduction to Bayesian Networks*. Springer. 1996.

[9] Horner JK A Bayesian network analysis of eyewitness reliablility: Part 1. Submitted to the *Proceedings of the 2014 International Conference on Artificial Intelligence.* CSREA Press.

[10] State of Florida v Zimmerman, 2012-CF-001083-A.

[12] Salmon WC. *The Foundations of Scientific Inference.* Pittsburgh. 1967.

[13] Chang CC and Keisler HJ. *Model Theory*. North Holland. 1990.

[14] Halmos PR. *Measure Theory*. D. Van Nostrand Reinhold. 1950.

[15] Bernays P. *Axiomatic Set Theory*. North-Holland, 1968. Dover reprint, 1991.