# A Bayesian Network Analysis of Eyewitness Reliability: Part 1

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. Here I describe a highly idealized Bayesian network model of the relation between eyewitness behavior and trial verdict. In a companion 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 and are rarely brought to trial.

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 highly idealized Bayesian model

There is some correlation between whether a correct verdict is reached and whether witnesses correctly observe a VDE. If the witnesses accurately observe the VDE, the probability of a correct verdict is generally higher than if they don't observe the VDE accurately. Various estimates of the ratios of the probabilities in these two cases range from 2:1 to 3:1 ([3]).

A distinguishing feature of the relationship between the probability of a correct verdict and the accuracy of observation of the witnesses is that the probability of a correct verdict depends on the accuracy of the observation of the VDE.

This tells us that *conditionality* ([4], p. 23) is in play. In probability theory, conditionality can be captured as a *conditional probability*. A *conditional probability*, P(X|Y), "the probability of X, given Y" is a probability measure defined by ([4], Section 9.1)

$$P(X|Y) = P(X \cap Y) / P(Y) \quad \text{Eq. 1}$$

where

X and Y are sets
$\cap$ is set intersection.

We can model this problem in a simple Bayesian network (BN, [2]). The heart of a BN is Bayes Theorem ([4], p. 320), derivable from the probability axioms ([4], p. 23). In its simplest form, Bayes Theorem is

$$P(A|B) = P(B|A) \times P(A) / P(B) \quad \text{Eq. 2}$$

where

P(A) is the probability of A
P(B) is the probability of B
P(A|B) is the probability of A, given B
P(B|A) is the probability of B, given A

Bayesians view probability as *degree of belief*, and regard Eq. 1 as a statement about the relationship between the degree of belief in a hypothesis (A) and the degree of belief in evidence (B). In this view, the left-hand-side of Eq. 1 is the degree of belief in hypothesis A, given evidence B. P(B|A), to Bayesians, denotes the degree of belief in the evidence B, given the hypothesis and call P(B|A) in Eq. 1 a "prior probability", or more briefly, a "prior". I will use the term "prior" to denote the uninterpreted term P(B|A) in Eq. 1, but remaining agnostic about the Bayesian claim that probability is a measure of degree of belief.

We need only any three of the quantities in Eq. 1 to determine the remainder. For example, if we have values for each of

P(A) = unconditional probability of a correct verdict
P(B) = unconditional probability of an accurate observation
P(B|A) = the probability of an accurate observation, given a correct verdict

we can calculate P(A|B), the probability of a correct verdict, given an accurate observation.

A BN is a system of conditional probabilities conforming to Eq. 2, mapped onto a directed graph ([5]) of whose nodes represent random variables ([4], Section 3.1). A possible value of a random variable X is called a *state* of X. An edge between a node A and a node B means that the distribution of possible values (states) of B depends probabilistically on the distribution of possible values (states) of A.

For the sake of discussion, let's assume a highly idealized model of the relationship between eyewitness and verdict in which:

(SC)

- the verdict is determined solely by the testimony of three independent eyewitnesses who observe a VDE

- in the case of a correct verdict, each eyewitness has a probability of only 0.75 of correctly observing the VDE
- in the case of  an incorrect verdict, each eyewitness has a probability of 0.25 of correctly observing the VDE
- in the absence of any observation, the probability of a correct verdict is 0.5

Figure 1 shows a graphical user-view of  a BN, *SimpleThreeWitness* (*STW*), 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 *STW*, a variable can be toggled between a finding and a calculated value by a mouse-click.

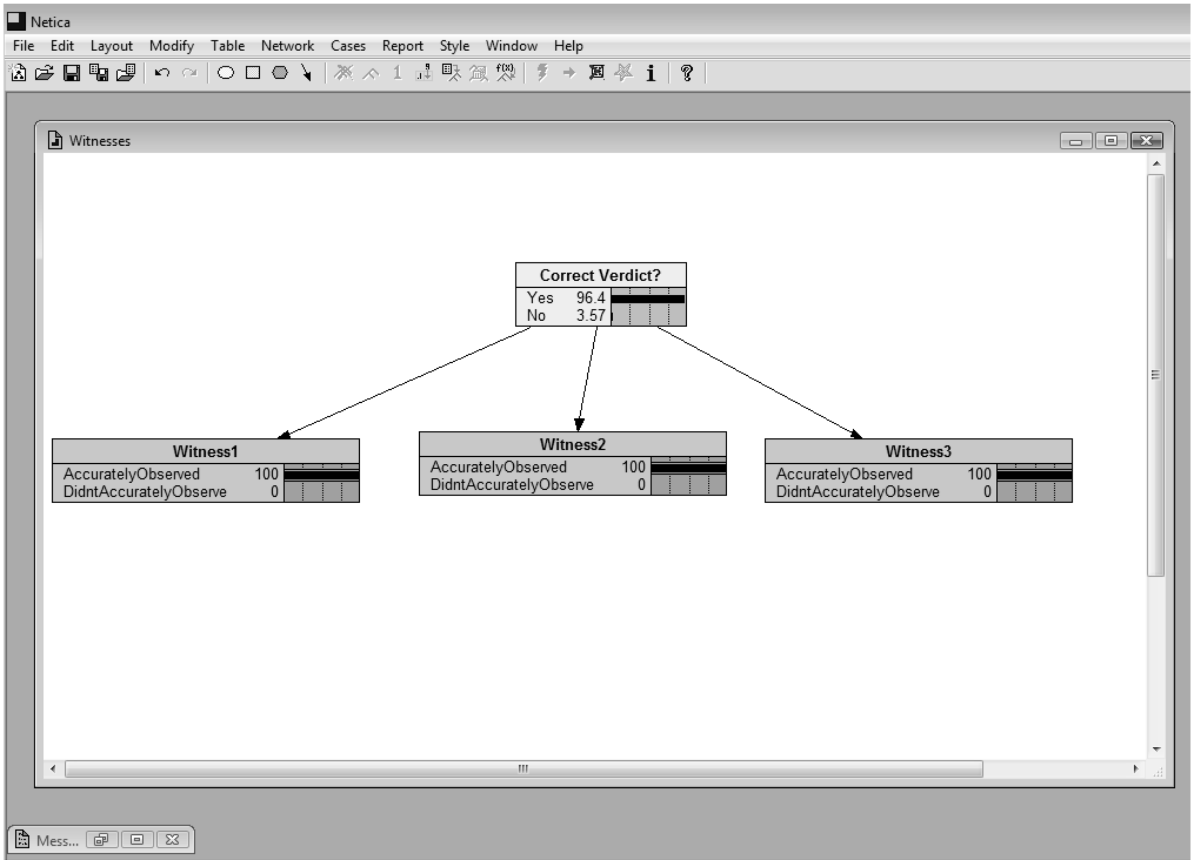


**Figure 1. User-view of *STW*, assuming all three witnesses correctly observe a VDE. In the configuration, the probability of a correct verdict is 0.96.**

In this example, if all three witnesses correctly observe the VDE, then the probability of a correct verdict is 0.96, even though the prior probability of any single eyewitness correctly observing the VDE is no greater than 0.5.

What happens in *STW* if only two of the three witnesses correctly observe the VDE? Figure 2 shows that the probability of a correct verdict is 0.75 -- even though the probability that any particular witness correctly observes the VDE ranges from 0.25 - 0.6.

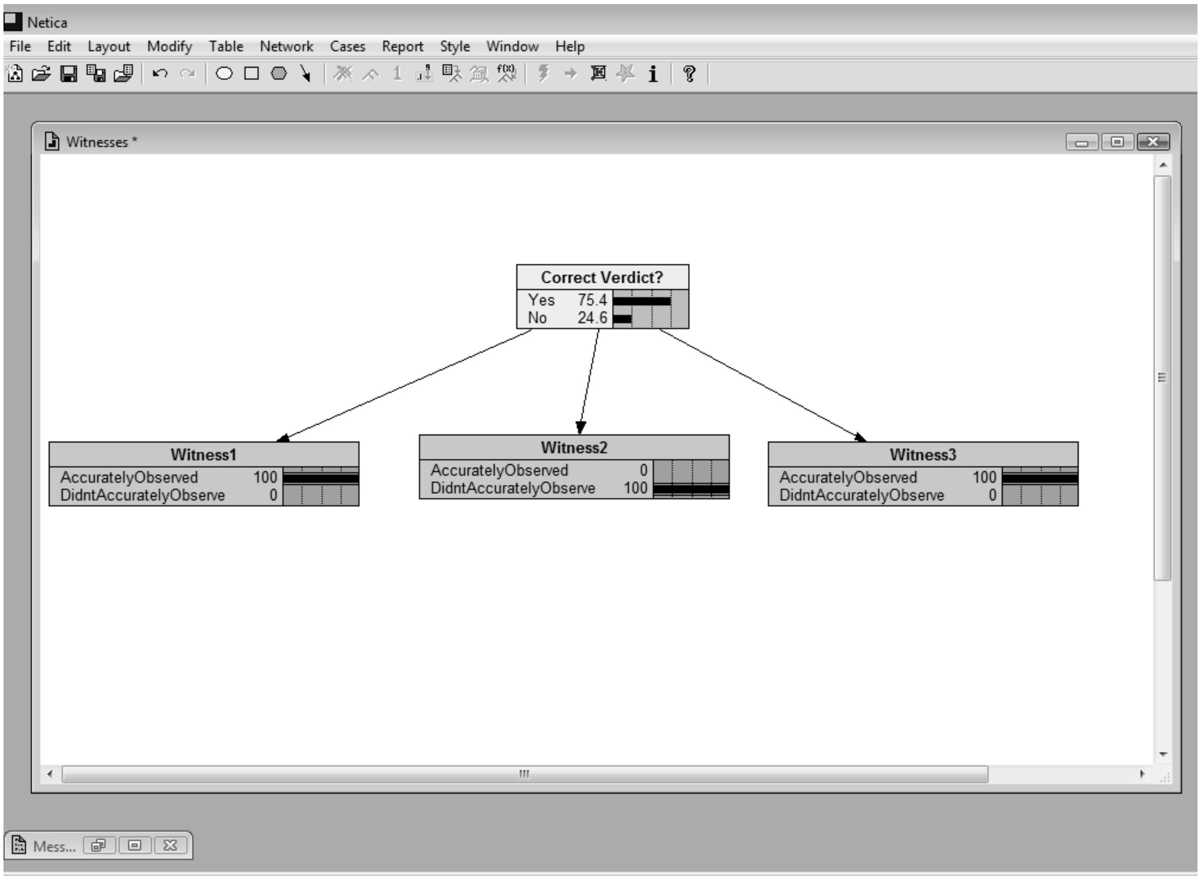


**Figure 2. User-view of *STW* when only two witnesses observe the VDE correctly. The probability of a correct verdict is 0.75.**

How sensitive is *STW* to the particular choice of priors? In general, answering this question requires analyzing a large set of cases. The case in which the probability of accurate observation is the same whether a correct verdict is achieved turns out to be especially interesting. In particular, let's assume the conditions of *STW*, but assume that probability of a witness accurately observing the VDE is 0.75 regardless of whether the verdict is correct. In such a case, we expect the probability of a correct verdict, given accurate observations by the witnesses, to be the same.

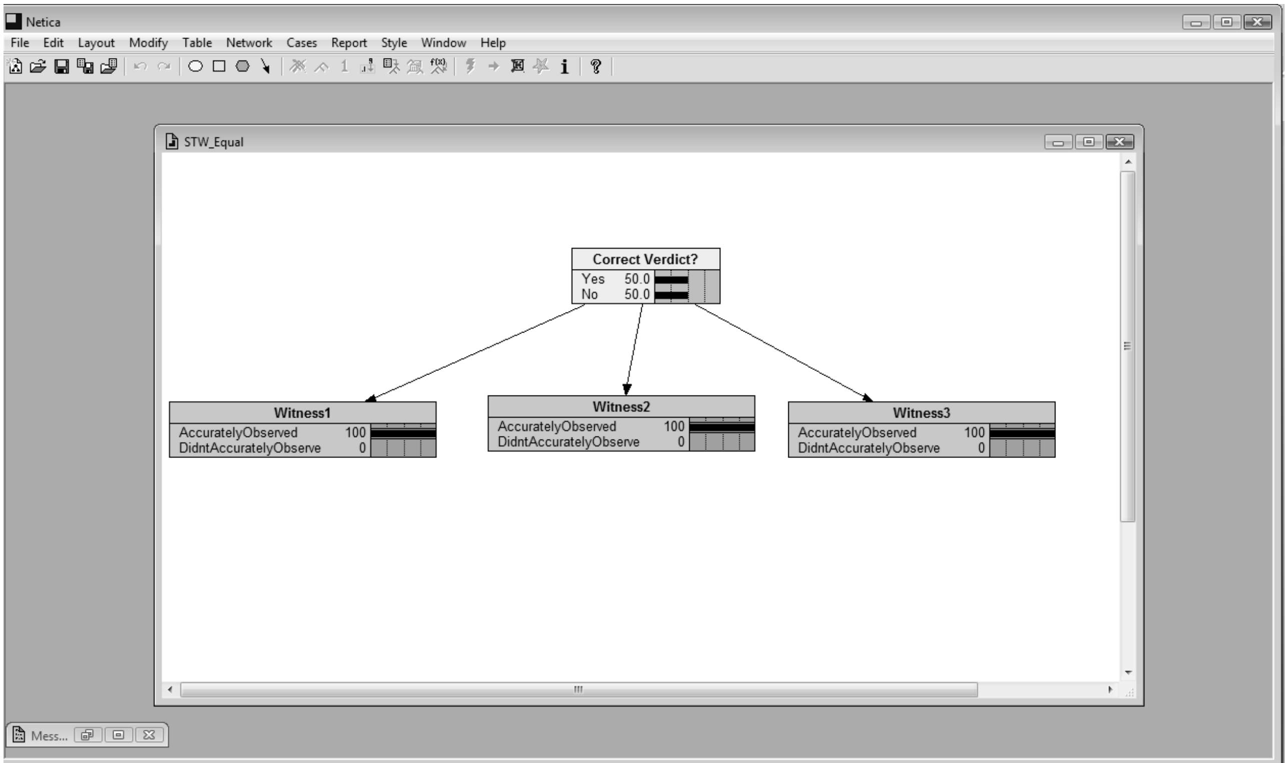


**Figure 3. User-view of *STW*, assuming all witnesses correctly observed the VDE.**

As expected, *STW_Equal* predicts the same distribution of probabilities for *Correct Verdict* if none of the witnesses correctly observe the VDE (Figure 3).

## 3.0 Discussion

The analysis above motivates several observations:

1. The technique shown here can be extend to an arbitrary number of witnesses, although the effect of more than three correctly observing witnesses 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 more refined model that shows the predictions of the model described in this paper are surprisingly informative.

## 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.

[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.