

A Castro Consensus: Understanding the Role of Dependence in Consensus Formation

Jarred Allen
AMISTAD Lab
Dept. of Computer Science
Harvey Mudd College
Claremont, CA, USA
jaallen@hmc.edu

Cindy Lay
AMISTAD Lab
Dept. of Mathematical Sciences
Claremont McKenna College
Claremont, CA, USA
clay22@cmc.edu

George D. Montañez
AMISTAD Lab
Dept. of Computer Science
Harvey Mudd College
Claremont, CA, USA
gmontanez@hmc.edu

Abstract

Consensus is viewed as a proxy for truth in many discussions of science. When a consensus is formed by the independent and free deliberations of many, it is indeed a strong indicator of truth. Yet not all consensuses are independent and freely formed. We investigate the role of dependence and pressure in the formation of consensus, showing that strong polarization, external pressure, and dependence among individuals can force consensus around an issue, regardless of the underlying truth of the affirmed position. Dependence breaks consensus, often rendering it meaningless; a consensus can only be trusted to the extent that individuals are free to disagree with it.

1 Introduction

Water freezes at 0 °C, the earth revolves around the sun, and cigarettes can cause lung cancer. You likely agree with these statements, as most do. Furthermore, sound scientific evidence supports each of these claims and many experts agree with that evidence. A *consensus* surrounds each of the claims, seemingly formed from the research and testimony of many independent expert parties. However, complete independence is rare. When independence is compromised, the reliability of consensus as a marker of truth suffers.

Consider cases of political consensus. In Fidel Castro's Cuba, superficially democratic elections unfailingly produced overwhelming support for the Communist Party of Cuba. However, a lack of choices left no other option. In both municipal and national elections, potential candidates had to pass specific requirements and secure backing from government-influenced organizations to be placed on a ballot (Shugerman, 2018; Foundation for Human Rights in Cuba, 2019). Thus, the overwhelming consensus supporting the Cuban Communist Party in elections resulted from political

pressure on voters and government restrictions on candidates. The consensus was neither independent nor meaningful.

Dependence can compromise consensus in other ways. Psychologists have found that eyewitness accounts can become distorted, often subconsciously, when witnesses unwittingly influence the testimony of other witnesses (Memon et al., 2008).

Using mathematical modeling, we explore the influence of dependence on the formation of consensus. We show that dependence among a population often leads to consensus regardless of the truth of the underlying position. We find that while a consensus of independent individuals strongly correlates with truth even under adverse conditions, dependence greatly weakens the evidential value of any consensus formed. Centuries ago, the Reverend Thomas Bayes derived his namesake theorem to answer the question of whether the testimony of a large number of independent eyewitness could ever establish the probable occurrence of an a priori highly unlikely event. He showed that even highly unlikely events become plausible once the testimony of many independent witnesses is taken into account. Bayes' theorem has found many real-world uses, including military and medical applications (Cepelwicz, 2016). In agreement with its historical purpose, we use Bayes' theorem to model how consensus within a population can provide evidence for or against a position. We extend our analysis to dependent models, allowing us to consider the effects of differing levels of dependence, external pressure, and hyper-polarization on the formation of consensus within our framework.

2 Related Work

Bayes' theorem is widely used in the data sciences (Lock and Dunson, 2013; Sambasivan et al.,

2018; Rudner and Liang, 2002), often under independent data assumptions, which may not be met in practice (Gammerrman and Thatcher, 1991; Domingos and Pazzani, 1996). Some have investigated use of Bayes’ theorem with dependent data, either by reformulating the data to be independent (Gammerrman and Thatcher, 1991) or adjusting the math to account for dependence (Brune and Pierce, 1974). Others have analyzed errors introduced into Bayes’ theorem by incorrect independent data assumptions (Domingos and Pazzani, 1996; Russek et al., 1983), finding that such approaches usually identify the correct hypothesis as the most likely, even if incorrectly computing posterior probabilities.

Researchers have investigated means by which witnesses can become dependent. Two approaches have been taken to explain this phenomenon: *peer pressure* and *the spread of misinformation*. Many peer pressure models exist. Some model peer pressure from axioms observed in the behavior of people (Estrada and Vargas-Estrada, 2013; Castillo-Garsow et al., 1997), while others derive peer pressure from considerations of utility maximization (Bishop, 2006; Yang et al., 2015). Misinformation spread has also been modeled, in the context of social networks (Kempe et al., 2003). More recent research uses sophisticated techniques to model the concurrent spread of both misinformation and correct information (Nguyen et al., 2012; Tambuscio et al., 2015; Abdullah et al., 2015), for both homogenous and heterogenous populations.

Decades ago, Gold discussed the dangers that *herd behavior* poses for scientists, potentially leading to an inertia-driven persistence of false consensus opinion within the sciences (Gold, 1989). Gold’s prescient essay stands as a warning. More recently, Kahneman highlighted the statistical advantages of averaging independent errors, and the dangers of dependent corroboration (Kahneman, 2011). Foreshadowing the mathematical analysis presented here, Kahneman states, “*However, the magic of error reduction only works well when the observations are independent and their errors uncorrelated. If the observers share a bias, the aggregation of judgments will not reduce it.*” (Kahneman, 2011).

3 Models

We introduce our models, first in the fully independent setting, followed by a series of modifi-

cations that allow us to account for polarization, time-varying dependence, external pressure, and partial dependence. We begin by reviewing when consensus performs as expected, allowing us to infer the probable truth of a position based on the strength of consensus affirming it. In what follows, we refer to individuals who **vote** to either affirm or deny a position as **witnesses**, referencing both our previous political voting discussion and Bayes’ original use of the theorem (while thoroughly mixing metaphors).

3.1 Independence: When Consensus Works

We begin with fully independent witnesses, where each witness has some probability p_f of fallibly affirming a hypothesis when it is not true. Let E represent our observed *evidence*, namely, the event of observing a given set of affirmations and denials. We then let $p_t := P(E|H)$ be the fixed probability of the evidence emerging if the hypothesis is true. We also let the hypothesis have some prior probability, p_h . Lastly, we let N be the number of witnesses who have reported.

Letting H denote the event that a hypothesis is true, and letting E be the event that there is an observed consensus of all N individuals affirming the hypothesis, we can use Bayes’ theorem to determine the posterior probability of the hypothesis being true given the affirming consensus, namely $P(H|E)$. Under our independence assumption we have $P(E|\neg H) = p_f^N$, and applying Bayes’s theorem, we obtain

$$\begin{aligned} P(H|E) &= \frac{P(E|H)P(H)}{P(E|H)P(H) + P(E|\neg H)P(\neg H)} \\ &= \frac{p_t p_h}{p_t p_h + p_f^N (1 - p_h)} \end{aligned}$$

with $P(\neg H|E) = 1 - P(H|E)$ as expected.

Let us suppose that a consensus is unlikely to emerge when the hypothesis holds ($p_t = 0.1$) and that people are each often mistaken ($p_f = 0.95$). Additionally, let us assume the prior probability of the hypothesis is low ($p_h = 10^{-9}$) and that the number of people is small ($N = 500$). We find that $P(H|E) \approx 0.932$. Thus, despite highly unfavorable assumptions, an independently-formed consensus still provides strong evidence that the hypothesis is true, and a consensus is unlikely to emerge around a false claim. When individuals are fully independent, $P(H|E)$ rapidly approaches 1 as we increase the number of individuals, imply-

ing the likely truth of an affirmed position given a consensus around it.

3.2 Dependence: Compromising Consensus

While a consensus of fully independent individuals can provide strong evidence in favor of the affirmed position, a forced consensus is epistemically much less meaningful; it can result from either truth or coercion, and we have no way of knowing which is the case based on the strength of the consensus alone. We next explore models of increasing complexity that allow us to control the level of dependence and observe its effect on consensus formation.

3.3 Polarized Majority Vote Model

We can view eyewitness accounts as coming in sequentially, with people being aware of the total fraction of prior testimonies which affirm or reject the hypothesis. Witness i receives a numerical representation of this fraction:

$$M_i = \frac{n_i + 1}{N_i + 2} \quad (1)$$

where M_i is the witness's perception of the prior accounts (which could be viewed as a proportion of "votes" in favor, hence our name for this model), n_i is the number of witnesses who affirmed the hypothesis before the i -th witness, and N_i is the total number of witnesses who either affirmed or rejected the hypothesis before the i -th witness. We use Laplace smoothing (i.e., pseudocounts) to ensure that M_i is defined for all witnesses and to prevent the second witness from seeing a pure consensus around whichever position the first witness happens to take.

Next, we define the probability of a witness incorrectly affirming a false hypothesis. Let

$$f(\theta, M_i) = \frac{1}{1 + e^{-\theta(M_i - 0.5)}}$$

where $\theta > 0$ is a *polarization* parameter controlling the sensitivity of witnesses to slight majorities. Letting X_i denote the random outcome of witness i and defining $\gamma_{\theta, \min} := f(\theta, 0)$ and $\gamma_{\theta, \max} := f(\theta, 1)$, the probability that the i -th witness affirms a false hypothesis can be modeled by

$$g_\theta(M_i) = P(X_i = 1 | M_i) := \frac{f(\theta, M_i) - \gamma_{\theta, \min}}{\gamma_{\theta, \max} - \gamma_{\theta, \min}}.$$

Note we have rescaled g_θ to the range $[0, 1]$, ensuring that only the sensitivity of each witness to a

change in the majority response among witnesses is adjusted as the value of θ changes. Figure 1 shows the effect of θ on the shape of the g_θ .




Figure 1: Effect of polarization parameter θ on the probability of affirming a false hypothesis versus M_i .

Once witness i either affirms or rejects the hypothesis, the updated value M_{i+1} is computed, and then the cycle repeats for witness $i + 1$. In modeling $g_\theta(M_i)$ as we have, we have assumed that M_i is sufficient for determining the conditional probability of X_i given $X_1, \dots, X_{i-1}$. Within our Bayesian framework from Section 3.1, we can then model $P(E|\neg H)$ using $g_\theta(M_i)$ as

$$\begin{aligned} P(E|\neg H) &= \prod_{i=1}^n P(X_i | X_1, \dots, X_{i-1}) \\ &= \prod_{i=1}^n P(X_i | M_i) \\ &= \prod_{i=1}^n g_\theta(M_i)^{X_i} (1 - g_\theta(M_i))^{1-X_i}. \end{aligned}$$

For a unanimous consensus, this simplifies to $P(E|\neg H) = \prod_{i=1}^n g_\theta(M_i)$. Note that because M_i is affected by all prior outcomes, dependence among witnesses is introduced through the chained conditioning on statistic M_i .

3.3.1 Polarized Majority Vote Model Results

As shown in Figure 2, witnesses are likely to reach an agreement about their claims of whether or not the event happened, at a rate which depends on the polarization θ . Thus, under this model a consensus usually emerges. Because this model is symmetric under switching affirming and denying, a consensus is as likely to form affirming an idea as it is to form denying it. In this case, the amount of




Figure 2: The probability of consecutive witness agreement over time, estimated through simulation using the polarized majority vote model. The shaded regions indicate 95% confidence intervals.

evidence gained by observing a consensus among witnesses is greatly diminished, since the formation of a consensus is not strongly affected by the truth or falsity of the position affirmed.

3.3.2 Polarized Majority Model Limitations

There are a few limitations of this model. First, the behavior of witnesses is completely symmetric to the hypothesis being true or false. This does not reflect many real-world scenarios, where there may be external pressures to bias responses, such as the aforementioned political coercion present in one-party states, which severely punish dissenters. We address this issue in Section 3.5.

Second, this model is also slow to adapt to a changing consensus opinion. If there is a change in the local majority opinion, then it is likely that people would perceive the new local majority as the actual majority. However, in this model, the shift to the new local majority being perceived as the majority will not occur until after it reaches an absolute majority, which may take much longer. We address this issue in Section 3.4.

Third, while our model controls for polarization and introduces a level of dependence among witnesses, it does not allow us to fine-tune for partial dependence or interpolate between a fully independent and fully dependent model. We explore such an extension in Section 3.6, though the model presented there is not exhaustive and leaves room for future work.




Figure 3: Plot of $(1-r)$ vs. number of witnesses needed to overcome an existing consensus, if $W_1 = 0$ and all subsequent witnesses affirm the hypothesis. As $(1-r)$ increases, indicating a stronger reliance on historically entrenched outcomes, the number of witnesses needed to overturn the existing consensus surges.

3.4 Recent Majority Vote Model

In the polarized majority vote model, M_i was taken as the smoothed average of the proportion of witnesses who affirmed a position prior to i , with each outcome equally weighted. Instead of computing the simple uniform average, we can introduce a notion of time-varying dependence by weighting more recent votes more heavily. We do so by computing a new proportion statistic, W_i , defined recursively as

$$W_i = (1-r) \cdot W_{i-1} + r \cdot x_{i-1} \quad (2)$$

where W_1 is the perception of the initial witness (with $W_1 = 0.5$ representing an even split, $W_1 = 0$ representing a consensus against the hypothesis, and $W_1 = 1$ representing a consensus affirming the hypothesis), r is a parameter of the system which represents the strength of the bias towards recent responses, and x_i is 1 if witness number i affirmed the hypothesis and 0 if witness number i rejected it.

Equation 2 is equivalent to the summation expression

$$W_i = \sum_{j=0}^{i-1} w(r, i-j) \cdot x_j, \quad (3)$$

$$w(r, \delta) = (1-r)^{\delta-1} r$$

with x_0 defined equal to W_1/r , and all other x_i defined the same as in Equation 2. Plotted in Figure 4, the weighting function $w(r, \delta)$ determines the




Figure 4: Weight versus temporal distance for various values of r . Larger values of r assign higher weight to more recent witnesses.

contribution of previous witnesses based on temporal proximity, where r is a parameter controlling the strength of temporal decay and δ is the temporal distance between the two witnesses. Using $w(r, \delta)$, Figure 3 plots the number of new witnesses needed to overcome an existing consensus as a function of r , showing that as temporal proximity matters less and more weight is placed on older outcomes (namely, as r decreases), the number of contrary witnesses needed to overcome an existing consensus grows exponentially. In such cases, consensus around a false position becomes extremely difficult to dislodge; unless one is open to new ideas and giving nonzero weight to contrary opinions, they might find themselves trapped in a false consensus for an extremely long time. Such a dynamic is closely related to herd behavior as noted by Gold (Gold, 1989).

3.5 External Pressure Model

The three models defined thus far do not consider external pressures placed on witnesses, as often arise in repressive political situations. To model this behavior, we add two new parameters, α and β with $0 \leq \alpha \leq \beta \leq 1$, to encode external pressures influencing witness responses. We then define the probability of a given witness affirming the hypothesis, $g_{\theta\alpha\beta}(W_i)$, as

$$g_{\theta\alpha\beta}(W_i) = (\beta - \alpha)g_{\theta}(W_i) + \alpha. \quad (4)$$

Thus, we rescale the function to output a value between the parameters α and β . The Recent Majority Vote Model presented in Section 3.4 is a special case of this model with $\alpha = 0, \beta = 1$.

W_i is determined by the method defined in Equation 2. Thus, this model resolves two issues with the Polarized Majority Vote Model presented in Section 3.3, by considering external pressures and adapting to a recent majority.

To model an external pressure against affirming, we can set β less than one, to decrease the likelihood of affirmative responses. Likewise, to model an external pressure towards affirming, α can be set greater than zero, which decreases the likelihood of negative responses.

3.5.1 External Pressure Model Results

Figure 5 shows the influence of α on the likelihood of witnesses affirming. This result, along with the setup being symmetric around affirming vs. rejecting the hypothesis, demonstrates that external pressure will likely push the group to a consensus on whichever position they are being pressured towards. However, in the presence of a weak external pressure, it is possible, albeit unlikely, that a near-consensus may emerge against that position.

Accordingly, a consensus position towards which witnesses are pressured supplies little evidence in favor of the consensus position, as such a consensus is highly likely to emerge regardless of the truth of the position around which it emerges.

3.6 A Spectrum of Dependence

Finally, we introduce a model which allows us to modify the level of dependence among witnesses. We add two parameters to the model, λ and p_f , where $\lambda \in [0, 1]$ controls the dependence of each witness on others and $p_f \in [0, 1]$ is the probability that a witness will fallibly affirm a false hypothesis, independent of the influence of other witnesses. We then define the probability that a given witness will incorrectly affirm a hypothesis as

$$g_{\theta,\lambda}(W_i) = (1 - \lambda)p_f + \lambda g_{\theta}(W_i). \quad (5)$$

The independent model from Section 3.1 is a special case of this model, with $\lambda = 0$. The dependent model from Section 3.4 is another special case, with $\lambda = 1$. Furthermore, this model can be emulated by Equation 4 from Section 3.5, by setting α and β as

$$\begin{aligned} \alpha &= (1 - \lambda)p_f \\ \beta &= p_f + (1 - p_f)\lambda. \end{aligned} \quad (6)$$

In the simulations for this model, we run the model until it reaches an equilibrium of W_i val-





Figure 5: Evolution of the probability of a witness affirming for various values of α . Shaded regions denote 95% confidence intervals. Simulations were run with $\theta = 10, \beta = 1, r = 0.035$. In (a), $W_1 = 0.5$. In (b), W_1 is chosen such that the first witness has a 50% chance of affirming.

ues, and then observe the probability of a consensus among the N subsequent witnesses. We are concerned with factors that lead to the eventual formation of a consensus, rather than only considering consensus that emerge immediately.

3.6.1 A Spectrum of Dependence Results

Considering the case when individual witnesses are each highly likely to mistakenly affirm a false hypothesis (as in Section 3.1), Figure 6 shows the result of applying Bayes' theorem to find the posterior probability that a hypothesis is true, given that a consensus has formed affirming it. It shows that the strength of evidence from a consensus decreases rapidly as independence among the witnesses becomes compromised, and the consensus no longer protects against the fallibility of individual witnesses, in stark contrast to the full independence case from Section 3.1.

From Figure 7, we can see that the probability of a consensus affirming a hypothesis grows exponentially on the level of dependence, reaching a probability of half at full dependence. We only show data for one set of parameters, because all values of the parameters we investigated resulted in similar plots.

Figure 8 shows the results of the same calculation, but allowing for a few dissenters from the clear consensus position, for two different sets of parameters. Figure 8a shows that, if individual witnesses are likely to be reliable, small amounts of dependence do not change the probability of consensus substantially, but larger values of de-




Figure 6: The posterior probability of a hypothesis given different values of λ and different priors, with $p_f = 0.95, p_t = 0.5, N = 150, \theta = 10, r = 0.035$.




Figure 7: The probability of a consensus affirming for different values of λ , found by simulation with parameters $p_f = 0.5, N = 13, \theta = 10, r = 0.035$, averaged over 8,000 independent trials. Note the log scale. 95% confidence intervals are plotted, but are not visible.




Figure 8: The probability of a near-consensus affirming for values of λ , found by simulation with $\theta = 10$, $r = 0.035$. (a) was run with $p_f = 0.35$, $N = 20$, and a consensus was defined as $\geq 90\%$. (b) was run with $p_f = 0.5$, $N = 30$, and a consensus was defined as $\geq 90\%$ agreement on the hypothesis. The shaded regions indicate 95% confidence intervals.

pendence result in exponential growth of the probability, causing the evidence from the consensus to become compromised. Figure 8b shows that, when an individual witness is equally likely to be right as wrong, the probability of a near-consensus forming grows exponentially on the value of λ .

3.6.2 Overall Consensus Probability

Section 3.6.1 presents the probability of a consensus forming affirming a given hypothesis. If one observes a consensus around a given position, that probability would be useful to determine the strength of the evidence for that position provided by the observed consensus. However, it may also be useful to know the overall probability of any consensus forming (for or against a hypothesis).

Figure 9 shows the overall probability of a consensus in either direction, based on λ and p_f . It shows that the probability of consensus formation increases with λ , equaling 1 when $\lambda = 1$, and it decreases as p_f approaches 0.5, being symmetric around individuals being either reliable or error-prone. With strong dependence a consensus is almost certain to arise, whether or not the position being held is true. Thus, observing a consensus in highly dependent cases tells us almost nothing concerning the truth status of the consensus position. Such a consensus lacks any real meaning.

4 Conclusion

A *Castro Consensus* is a near-unanimous show of agreement brought about by means other than

the honest and uncoerced judgements of individuals. Using mathematical modeling, we demonstrate how dependence, polarization, and external pressure compromise the relation between truth and consensus. When individuals are fully independent, even under highly unfavorable circumstances a consensus provides strong evidence for the correctness of the affirmed position. This no longer remains the case once dependence, polarization, and external pressure are introduced. With such interventions, the probability of a false consensus increases dramatically.

All models except the one presented in Section 3.6 tend to approach a consensus around one position. In Section 3.5, we find that strong pressure against a position will eventually result in a consensus against that position. This suggests that in cases where holding certain contrary positions are punished, consensus becomes an unreliable indicator of truth. In Section 3.6, we find that the probability of a consensus or near-consensus forming grows rapidly as dependence among the witnesses increases. We find that even slight amounts of dependence between witnesses can greatly decrease the strength of evidence provided by a consensus for the truth of a position.

There are many ways to model agreement and consensus formation, of which our chosen models only represent one small set of possible options. Let us be clear that we make no sociological claims concerning our models' ability to fully capture and adequately model the dynamics of hu-




Figure 9: The probability of a near consensus forming, found via simulation with $\theta = 10$, $r = 0.035$, $N = 20$. A near-consensus was defined as $\geq 90\%$ agreement on the majority position. Subfigure (a) plots the consensus probability as a function of p_f for several values of λ , while (b) plots the consensus probability as a function of λ for several values of p_f . The shaded regions indicate 95% confidence intervals.

man agreement and testimony. The primary value of this work is not in presenting any particular set of models *but in demonstrating how those models change* in response to external factors like dependence and pressure. As such, we present an initial model for which consensus **does** strongly and reliably give evidence of the truth of the affirmed position, and proceed to show how the introduction of dependence, polarization, and pressure destroys the evidential value of the consensus in that same model. We demonstrate how dependence, pressure, and polarization can force a consensus, making reliance on consensus as an indicator of truth unreliable. As a result, a consensus can only be trusted to the extent that individuals are free to disagree with it, without repression or reprisal. Similarly, when strong incentives favor affirmation of a position, a consensus affirming it becomes almost inevitable, and therefore all but meaningless.

Acknowledgements

The authors would like to thank Caroline Chou, Jessica Winssinger, and Corbin Bethurem, whose early classroom collaboration on a similar project helped inspire the work presented here. This research was supported in part by the generous support of Harvey Mudd College and the National Science Foundation under Grant No. 1950885. Any opinions, findings or conclusions expressed are the authors' alone, and do not necessarily reflect the views of Harvey Mudd College or the National Science Foundation.

References

- Nor Athiyah Abdullah, Dai Nishioka, Yuko Tanaka, and Yuko Murayama. 2015. User's action and decision making of retweet messages towards reducing misinformation spread during disaster. *Journal of Information Processing*, 23(1):31–40.
- John Bishop. 2006. Drinking from the fountain of knowledge: Student incentive to study and learn—externalities, information problems and peer pressure. *Handbook of the Economics of Education*, 2:909–944.
- HD Brune and Donald A Pierce. 1974. Estimation of discrete multivariate densities for computer-aided differential diagnosis of disease. *Biometrika*, 61(3):493–499.
- Carlos Castillo-Garsow, Guarionex Jordan-Salivia, Ariel Rodriguez-Herrera, et al. 1997. Mathematical models for the dynamics of tobacco use, recovery and relapse. *Public Health*, 84(4):543–547.
- Jordana Cepelewicz. 2016. [How a Defense of Christianity Revolutionized Brain Science](https://bit.ly/2ZFWP5u). <https://bit.ly/2ZFWP5u>. Accessed: 2020/05/28.
- Pedro Domingos and Michael Pazzani. 1996. Beyond independence: Conditions for the optimality of the simple Bayesian classifier. In *Proc. 13th Intl. Conf. Machine Learning*, pages 105–112.
- Ernesto Estrada and Eusebio Vargas-Estrada. 2013. How peer pressure shapes consensus, leadership, and innovations in social groups. *Scientific Reports*, 3:2905.
- Foundation for Human Rights in Cuba. 2019. Elections in Cuba: What's the Point in Voting? <https://bit.ly/3fRTvuX>. Accessed: 2020/06/17.

- Alex Gammerman and AR Thatcher. 1991. Bayesian diagnostic probabilities without assuming independence of symptoms. *Methods of information in medicine*, 30(01):15–22.
- Thomas Gold. 1989. New Ideas in Science. *Journal of Scientific Exploration*, 3(2):103–112.
- Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Macmillan. Chapter 7, p. 84.
- David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146.
- Eric F Lock and David B Dunson. 2013. Bayesian consensus clustering. *Bioinformatics*, 29(20):2610–2616.
- Amina Memon, Serena Mastroberardino, and Joanne Fraser. 2008. Münsterberg’s legacy: What does eye-witness research tell us about the reliability of eye-witness testimony? *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 22(6):841–851.
- Nam P Nguyen, Guanhua Yan, My T Thai, and Stephan Eidenbenz. 2012. Containment of misinformation spread in online social networks. In *Proceedings of the 4th Annual ACM Web Science Conference*, pages 213–222.
- Lawrence M Rudner and Tahung Liang. 2002. Automated essay scoring using Bayes’ theorem. *The Journal of Technology, Learning and Assessment*, 1(2).
- Estelle Russek, Richard A Kronmal, and Lloyd D Fisher. 1983. The effect of assuming independence in applying Bayes’ theorem to risk estimation and classification in diagnosis. *Computers and Biomedical Research*, 16(6):537–552.
- Rajiv Sambasivan, Sourish Das, and Sujit K Sahu. 2018. A Bayesian Perspective of Statistical Machine Learning for Big Data. *arXiv preprint arXiv:1811.04788*.
- Emily Shugerman. 2018. Cuba election: When is it, who is voted in, and what does it mean for Raul Castro and the presidency? <https://bit.ly/31UMNzZ>. Accessed: 2020/06/15.
- Marcella Tambuscio, Giancarlo Ruffo, Alessandro Flammini, and Filippo Menczer. 2015. Fact-checking effect on viral hoaxes: A model of misinformation spread in social networks. In *Proceedings of the 24th international conference on World Wide Web*, pages 977–982.
- Han-Xin Yang, Zhi-Xi Wu, Zhihai Rong, and Ying-Cheng Lai. 2015. Peer pressure: Enhancement of cooperation through mutual punishment. *Physical Review E*, 91(2):022121.