Fake News and False Corroboration: Interactivity in Rumor Networks

Michael J. Spivey (spivey@ucmerced.edu)
Cognitive and Information Sciences
University of California, Merced
Merced, CA 95340 USA

Abstract

Rumors inundate every social network. Some of them are true, but many of them are false. On rare occasions, a false rumor is exposed as the lie that it is. But more commonly, false rumors have a habit of obtaining apparent verification, by corroboration from what seems to be a second independent source. However, in complex social networks, the connectivity is such that a putative second source is almost never actually independent of the original source. In the present work, rumor network simulations demonstrate how remarkably easy it is for a node in the network to be fooled into thinking it has received independent verification of a false rumor, when in fact that "second source" can be traced back to the original source. By developing a theoretical understanding of the circumstances under which the spread of false rumors, "alternative facts," and fake news can be controlled, perhaps the field can help prevent them from ruining elections and ruining entire nations.

Keywords: Networks, Social Networks, Interaction

Introduction

The interactivity that exists among the subsystems that form a cognitive system has powerful and lasting consequences. In the human brain, the interactivity among the neural subsystems that form the language comprehension network is what allows phonetics to influence syntactic processing (Farmer, Christiansen, & Monaghan, 2006) and semantics to influence speech perception (Gow & Olson, 2015; Spivey, 2016). In the human brain, the interactivity among the neural subsystems that form the visual perception network is what allows depth perception to influence motion discrimination (Trueswell & Hayhoe, 1995) and attention to influence visual perception (Gandhi, Heeger, & Boynton, 1999; Spivey & Spirn, 2000). In the human brain, the interactivity between the language comprehension network and the visual perception network is what allows visual context to influence spoken word recognition (Allopenna, Magnuson & Tanenhaus, 1998; Spivey-Knowlton, 1996), and linguistic input to influence visual perception (Lupyan & Spivey, 2010; Lupyan & Ward, 2013). These examples form just a tiny subset of the many consequences of interactivity in the human brain.

Outside the human brain, interactivity in a social network has powerful consequences for group behavior. When two people cooperate on a shared task, or even just have a conversation, they often exhibit real-time motor coordination in their postural sway (Shockley, Santana, & Fowler, 2003; M. Richardson, Marsh & Schmidt, 2005), their eye movements (D. Richardson, Dale & Kirkham, 2007), their gestures (Paxton & Dale, 2013), and their language use (Louwerse, Dale, Bard, & Jeuniaux, 2012). It has even been shown that behavioral and neural responses of two participants cooperating on a task exhibit the signatures of competition between the two subtasks, even though each person is in charge of only one of those subtasks (Sebanz, Knoblich, Prinz, & Wascher, 2006). Essentially, each person is doing some of the thinking for the other person. When these mechanisms of coordination are optimized between two people, they can even perform a joint perceptual task at a level that is better than either of them alone (Fusaroli et al., 2012).

When people share information with each other, they tend to self-organize into a larger cognitive system (Goldstone & Gureckis, 2009). Much like how cognition may be an emergent property of billions of neurons interacting with one another in a brain (Kello, Beltz, Holden & Van Orden, 2007), group cognition may also be an emergent property of multiple people interacting with one another in a shared context (Thiener, Allen, & Goldstone, 2010). Due to the continuous fluid flow of information throughout the network, every node (be it a neuron or person) is richly interdependent with every other node, at least indirectly. Not only can positive influences spread throughout such a network, as when two brains show improved performance on a shared perceptual task (Fusaroli et al., 2012), but negative influences can also spread throughout the network and infect nearly every component. Network simulations of rumor-spreading have recently begun to analyze this process of false information infecting a social network (Roshani and Naimi, 2012).

Traditional studies of rumor transmission tended to focus on linear sequential transfer of a rumor, and how the content can often become accidentally modified after several transmissions (Allport & Postman, 1947). Sometimes this is referred to as the "telephone game." However, more recent studies of rumor transmission have used network theory to examine how non-linear transmission of rumors happens in complex social networks that are richly interconnected (Del Vicario et al., 2016). For example, when the network has islands of homogeneity, tight-knit like-minded enclaves that connect mostly just to their own group, these subnetworks can become "echo chambers" that reinforce false narratives and conspiracy theories within their walls. Alternatively, when the connectivity of a social network is scale-free (neither random nor homogenous) – much like the brain's connectivity (Kello, 2013; Sporns, 2010) – then almost any rumor can be expected to spread throughout the entire network, irrespective of whether it is true or false (Nekovee, Moreno, Bianconi, Marsili, 2007). What has not been explored yet in this small cottage industry of research is how easily a false rumor can obtain independent verification via an apparent second source, even when that "second source" actually has the original source as its origin.

When an interactive system (be it a brain or a group of people) spends any amount of time sending signals back and forth among its subcomponents, it quickly becomes difficult to trace the source of a signal and determine whether a given signal is *afferent* (recently coming from an external source) or *efferent* (better described as generated endogenously). Under these circumstances, following the trail of a rumor in a social network is extremely difficult. The journalistic practice of "corroborating the story" can become quite complicated. A common method of fact-checking is to find a second source for the same story. If the second source is independent of the first source, and says essentially the same thing, then it adds veracity to the report. Even naïve experimental participants tend to use this tactic (Kim et al., 2008). However, in an interconnected network of people sharing information, almost no one is actually independent of anyone else. Frequently, an apparent second source, which gets used as verification of the rumor, actually acquired its information indirectly from the original source.

One concrete real-world example of such *false corroboration* is the U.S. Pentagon's case for Saddam Hussein stockpiling weapons of mass destruction (WMD) at the beginning of the 21st century. It has now been well-established that U.S. leaders were proactively seeking justification for a pre-existing plan to invade Iraq and depose its leader (Dreyfuss & Vest, 2004; Ryan, 2006). It turned out to be all too easy for information gatherers to fool themselves into thinking they had corroborated reports of WMD, when in fact the corroboration was actually a duplicate of the original false rumor. The CIA, British intelligence services, and the New York Times all collected reports of WMD in Iraq, and carefully sought independent verification. Each of these entities received fallacious reports from the same Iraqi defector, codenamed "Curveball" by the CIA. And what's more, each of them used un-sourced reports from one another as corroboration of their own report. What they each did not realize at the time was that the "second source" to corroborate their report from Curveball was actually just someone else's report from Curveball (Bamford, 2005; Prados, 2004).

False rumors, "alternative facts," and fake news have become an everyday occurrence recently, where too many people obtain their news reports on social network sites and blogs, where "news" is provided that has not been vetted by policies of ethical journalism. For example, in January of 2016, journalist and author, Fareed Zakaria, was "trolled" on the internet with a fake report of him calling for "jihad rape of white women to depopulate the white race." Some people believed this false rumor so strongly that they made threats on Zakaria's life, and frightening phone calls to his daughters in the middle of the night (Zakaria, 2016).

Similarly, in the fall of 2016, fake news reports were disseminated widely on Facebook about presidential candidate Hillary Clinton being involved in a child sex-trafficking ring based at a particular pizza shop in Washington, D. C. One man believed that false rumor so strongly that he felt compelled to travel across state lines to visit that pizza shop with an assault rifle in his hands and fire a shot to let them know he was there to save the children. The U. S. Department of National Intelligence has recently determined that many such fake news stories about Hillary Clinton were fabricated and disseminated via social networks specifically with the intent of influencing the results of the 2016 U. S. election (DNI Report, 2017).

In Del Vicario et al.'s (2016) computational analysis of conspiracy theories on the internet, they concluded that, "many mechanisms cause false information to gain acceptance, which in turn generates false beliefs that, once adopted by an individual, are highly resistant to correction." In the following rumor network simulations, the results suggest that *false corroboration* may be one of those many mechanisms.

Random Rumor-Net Simulations

In this first group of simulations, a 100-node network was constructed and given random placement of bi-directional connections, excluding self-connections. In one set of 100 simulations, the network was given 10% connectivity, such that each node on average was connected to about 10 of the possible 99 other nodes (i.e., average node degree=10). The average clustering coefficient for this network (which shows how interconnected each node's friends are) was .10. Another set of 100 simulations used a clustering coefficient of .33, and Figure 1 shows an example degree distribution from one of those networks. Another set of 100 simulations used a clustering coefficient of .5, and a fourth set used a clustering coefficient of .67.

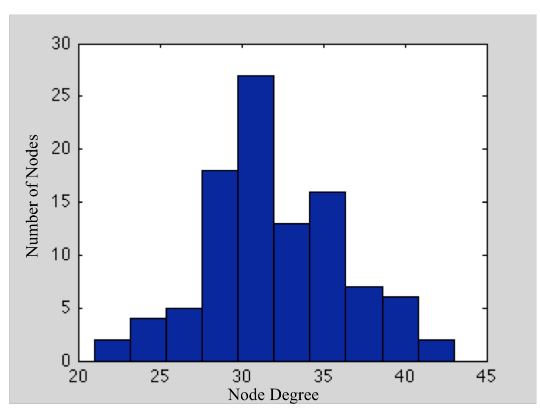

Figure 1: Degree distribution from a 100-node random network in which, on average, most nodes are connected to about 33 other nodes.

To begin a simulation, node #1 was infected with a rumor by flipping its state from zero to 1.0. This is the one-and-only origination of the rumor in this network. It could be true or false, but for the purpose of testing its evolution into "fake news," the rumor is treated as false. For every instance of transmitting the rumor, a randomly chosen infected node would select randomly among its connections to spread the rumor with one other node. After spreading, that bidirectional connection was erased in the network to prevent it from being used again in the future. (The simulation assumes that if the same rumor were shared again between the same two people, it would not count as a transmission.) For that very first transmission, this obviously involved node #1 sharing the rumor with one of the nodes connected to it. At which point there would then be two nodes that have been exposed to the rumor. Then one of those nodes was randomly selected to spread the rumor again. After 100 transmissions of the false rumor, some of the nodes had still never been exposed, some had been exposed once, and some had heard the rumor from two or more different connections. This latter case counts as people who had heard the rumor corroborated by what would seem to be a second source. However, the simulation actually has only one source of the rumor: node #1. For example, node #1 might spread the rumor to node #47, who then spreads the rumor to node #23. Next, node #47 might share the rumor again, this time with node #87, who shares it with node #18, who then shares the rumor with node #23. In that scenario, node #23 could easily be fooled into believing that it had received independent corroboration (from node #18) of the rumor it first heard from node #47.

In this first group of simulations, the number of nodes that received this false corroboration was recorded for low-, medium-, high- and very high-connectivity networks (i.e., clustering coefficients of .1, .33, .5, and .67). Interestingly, after 100 transmissions of the rumor, there were no differences across these four different types of random networks (results averaged across the 100 simulations in each case). In all simulations, irrespective of how densely interconnected the network was, around 26 of the 100 nodes had heard the rumor from two or more sources (Table 1). This insensitivity to network density is likely due to the fact that a rumor-spreader is randomly selected each time (among nodes that know the rumor), and its relative likelihood of spreading the rumor to a knowing node or an unknowing node is unchanged by how well-connected it is.

Table 1: Random networks with different numbers of connections show about the same number of nodes hearing false corroboration of the rumor (2+ times).

Avg Node Degree	Clustering Coefficient	Never Heard	Heard Once	Heard 2+ times
10	.10	33.6	40.1	26.3
33	.33	34.6	38.8	26.6
50	.50	34.8	38.4	26.8
67	.67	34.9	38.4	26.7

With 200 nodes and 200 rumor transmissions (or 500 nodes and 500 rumor transmissions), again about one-quarter of the nodes obtain false corroboration – irrespective of how densely or sparsely connected the network is. With half as many transmissions as there are nodes, about 10% of the nodes obtain false corroboration. And with twice as many transmissions as nodes, about 60% of the nodes obtain false corroboration. Based on these initial simulations, it appears that false corroboration of a rumor may be remarkably easy to obtain in a social network.

Scale-Free Rumor-Net Simulations

Most real-world networks, including social networks, are not at all random in their connectivity. Instead, social networks tend to have a scale-free pattern of connectivity, meaning that most nodes have a smallish number of connections (node degree), while a few nodes have a very large number of connections. Using a version of Barabasi and Albert's (1999) preferential attachment process, a group of scale-free rumor networks were designed that show a power-law in their degree distribution (Figure 2).

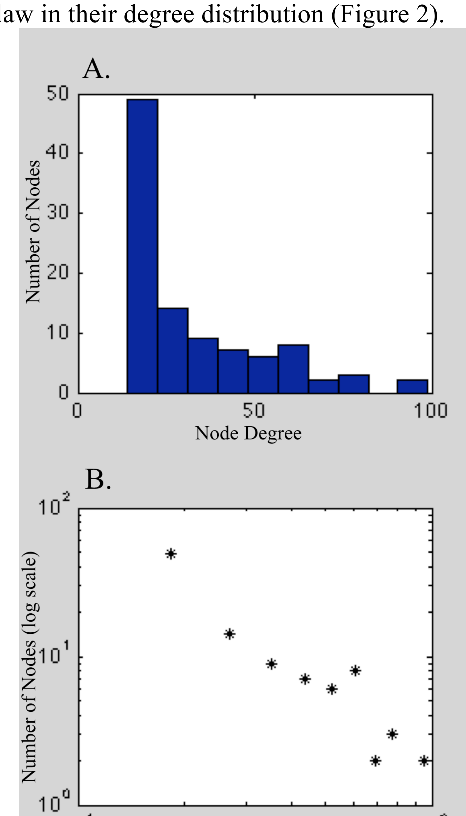

Figure 2: (A) Degree distribution from a 100-node scale-free network where the mean number of connections per node is 33, but most nodes have <25 connections and a few nodes have >75 connections. (B) On log-log coordinates, the degree distribution forms a relatively straight line with a slope of -1.3.

By contrast to a scale-free network, in a random network the proportion of connections each node has generally corresponds to the clustering coefficient as well. That is, if each node in a random network has about 10% of the possible connections, then the clustering coefficient (showing what proportion of each node's friends are connected to each other) will also tend to be around .10. However, in a scale-free network, the clustering coefficient (.62, in Figure 2) tends to be substantially higher than the average proportion of connections the nodes have (.33, in Figure 2). That is, in a scale-free network, most nodes have relatively few friends, but a sizeable proportion of those friends know each other.

In these next simulations, a hundred 100-node scale-free networks were designed that had an average of 10 connections per node, along with another hundred networks that had an average of 17 connections per node, then another hundred with 25, and another hundred with 33 connections per node. (In a scale-free network, when the average number of connections approaches 50% of the possible connections, its degree distribution can become bimodal and no longer adheres to a scale-free power law. Therefore, the highest node degree used here was 33.)

Each rumor-spreading simulation with these scale-free networks was carried out in a fashion similar to those with the random networks, except that the first rumor-infected node could not be an arbitrary choice because some nodes were substantially more connected than others. To test the limiting case, the least-connected node in each scale-free network was selected as the first node to spread the rumor. After that starting point, 100 transmissions of the rumor took place exactly as it did with the random networks.

Table 2: Scale-free networks with different numbers of connections show about the same number of nodes hearing false corroboration of the rumor (2+ times).

Node Degree	Cluster Coeff.	loglog slope	Never Heard	Heard Once	Heard 2+ times
10	.22	-0.71	40.1	34.6	25.3
17	.39	-1.05	40.4	34.3	25.3
25	.56	-1.25	40.8	33.9	25.3
33	.68	-1.27	40.8	33.6	25.6

In these scale-free rumor networks, a slightly larger proportion of the people never hear the rumor (about 40%) compared to that in the random networks (about 35%). However, remarkably, approximately the same number of false corroborations is observed (~25) as that seen with the random networks (compare Tables 1 & 2). As was tested with the random networks, this 25% false corroboration rate replicates for scale-free networks with 200 nodes and 200 rumor-transmissions. When there are 3-4 times as many transmissions as nodes, almost every node will have heard the rumor, and about ¾ of them will have heard it more than once (irrespective of network density). Not surprisingly, in these scale-free networks, it is usually the well-connected nodes that first obtain these false corroborations.

When a False Rumor Becomes Fake News

Based on all these simulations, when there are as many rumor-transmissions as there are nodes, then almost 2/3 of them will hear the rumor, and about 1/4 of them will obtain a false corroboration of the rumor – even though it never actually had any independent secondary source. This is true for both random rumor networks and for scale-free rumor networks. However, when one of the people in the network is a reporter for a news agency, who will broadcast the story to everyone if they obtain apparent corroboration, then it turns out that the type of connectivity does, in fact, matter. If one assumes that the reporter is among the most widely-connected people in the network, then the different degree distributions for random networks and for scale-free networks (Figures 1 and 2) make for substantially different reporters. In a random network, the most-connected node (i.e., the reporter) will have a number of connections that is greatly influenced by the density of the network's connectivity (its average node degree). However, in a scale-free network, the most-connected node is often connected to >85% of the other nodes, irrespective of the average node degree. Therefore, a reporter in a random network will only occasionally obtain a false corroboration, and thus publish the story (Table 3). However, in a scale-free network, a reporter (who is massively well-connected) will almost always obtain false corroboration, and therefore publish the rumor (Table 4). If that rumor is false, then its publication qualifies as fake news.

Table 3: In random rumor-nets, false corroboration sometimes leads to the publication of fake news.

Node Degree	Clustering Coefficient	Reporter-node Publishes Fake News
10	.10	58%
33	.33	42%
50	.50	38%
67	.67	35%

Table 4: In scale-free rumor-nets, false corroboration almost always leads to the publication of fake news.

Node Degree	Clustering Coefficient	loglog slope	Reporter-node Publishes Fake News
10	.22	-.071	92%
17	.39	-1.05	93%
25	.56	-1.25	93%
33	.68	-1.27	87%

Surprisingly, with random networks, denser connectivity leads to a reduced likelihood of the reporter-node obtaining false corroboration and publishing the rumor. Upon closer examination, this makes sense given the parameters of the simulation. In a random network with a small average node degree (sparse connectivity), whenever a rumor-infected node is about to spread the rumor, it has a small number of friends to choose among. If one of them happens to be the reporter, which is somewhat likely since the reporter is the most connected node, then the reporter might hear the rumor. And if that happens a second time, then a (false) corroboration has taken place, and the story gets broadcasted. By contrast, in a random network with a large average node degree (dense connectivity), whenever a rumor-infected node is about to spread the rumor, it has a large number of friends to choose among. One of them is probably the reporter, but a random selection of to whom the rumor will be spread leaves the reporter with a slim chance. In many of these random rumor-net simulations, the reporter never even heard the rumor once.

The situation is very different in a scale-free network. In a scale-free rumor-net, most nodes have fewer connections than they would in a comparable random network. Therefore, when a rumor-infected node is randomly selected to spread the rumor, it is usually one that has a smallish number of friends to choose among, and one of them is almost certainly the well-connected reporter (see also Doerr, Fouz, & Friedrich, 2012). Thus, almost every time the rumor is transmitted, the reporter has a reasonable chance of being its recipient. As a result, the reporter-node in such a network is highly likely to hear the rumor, and also highly likely to obtain a false corroboration of this rumor, even though the rumor actually has only one source.

Conclusion

Interactivity in a network is usually a good thing. Ambiguities or uncertainties present in one part of the network will often be resolved by strongly biasing information present in another part of the network (e.g., Kawamoto, 1993; MacDonald, Pearlmutter, and Seidenberg, 1994; McRae, Spivey-Knowlton, & Tanenhaus, 1998). However, when that strongly biasing information is objectively false, the interactivity within a network can compromise its ability to align itself with reality.

The present network simulations do not specifically distinguish between objectively false rumors and true rumors, but a recent analysis of 330 rumor threads on Twitter does. For a false rumor, the time between rumor onset and debunking can be as much as *seven times longer* than the time between rumor onset and verification for a true rumor (Zubiaga, Liakata, Procter, Hoi, & Tolmie, 2016). That is, it takes much longer to debunk a false rumor than it does to verify a true rumor. Therefore, if a long-standing uncertain rumor has not been verified as true, then the odds are steadily increasing every day that it is a false rumor (that just hasn't been debunked yet). Most true rumors get verified very quickly.

However, the nature of this verification process comes into question when considering the present rumor simulations. If the apparent verification comes in the form of a seemingly independent source that corroborates the original rumor, it may be illusory. The interactivity inherent in social networks can all too easily make a false corroboration (i.e., an echo from the echo chamber) appear as genuine independent corroboration.

One potential solution to this problem is for reporters to make better efforts at tracing the lineage of a report, so that two reports from the same source might be identified as such. A more reliable solution would be for journalism practices to avoid using secondary-source corroboration on its own as sufficient evidence to disseminate a story. These rumor network simulations demonstrate that it is simply too easy to obtain such corroboration in a fraudulent manner. Instead, the criterion for publication of a story might ought to include evidence that cannot easily be faked, such as photos, video, audio recordings, and documents whose source can be reliably determined. For example, if the report is that a public figure made sexist comments, or mocked a disabled person, or told the public a brazen lie, simply relying on two seemingly-independent sources to publish such a story may be insufficient. If the comments or mocking are evident in a video clip of the public figure, or if the lie is present in a verifiably-sourced tweet from the public figure, then those pieces of evidence should be what are repeatedly disseminated in reporting the story. Reports without such concrete evidence should be taken with a grain of salt, or perhaps not published in the first place.

It has been proven time and time again in everyday life, as well as in high-stakes politics, that the dissemination of false rumors can ruin lives, ruin elections, and even ruin entire nations. Understanding the mechanisms that allow, and exacerbate, the spread of misinformation in a social network of any kind may help with efforts to curtail and minimize the damage that can be done. The present simulations of a false rumor spreading throughout a network show convincingly that, even in a sparsely connected network, the "apparent corroboration" of a story often comes from a source whose own source can be traced back to the originator of the story, and thus should not actually count as independent corroboration. To quote Fareed Zakaria, "No matter how passionate people are, no matter how cleverly they can blog or tweet or troll, no matter how viral things get, lies are still lies."

References

Allopenna, P. D., Magnuson, J. S., & Tanenhaus, M. K. (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language*, *38*(4), 419-439.

Allport, G. W. & Postman, L. (1947). *The psychology of rumor*. NY: Holt & Co.

Bamford, J. (2005). *A pretext for war: 9/11, Iraq, and the abuse of America's intelligence agencies*. Anchor Publishing.

Barabási, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509-512.

Del Vicario, M., et al. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, *113*(3), 554-559.

DNI report (2017). *Background to "Assessing Russian Activities and Intentions in Recent US Elections": The*

Analytic Process and Cyber Incident Attribution. https://www.dni.gov/ files/documents/ICA_2017_01.pdf

Doerr, B., Fouz, M., & Friedrich, T. (2012). Why rumors spread so quickly in social networks. *Communications of the ACM*, *55*(6), 70-75.

Dreyfuss, R. & Vest, J. (2004, February). The lie factory. *Mother Jones*, pp. 34–41.

Farmer, T. A., Christiansen, M. H., & Monaghan, P. (2006). Phonological typicality influences on-line sentence comprehension. *Proceedings of the National Academy of Sciences*, *103*(32), 12203-12208.

Fusaroli, R., Bahrami, B., Olsen, K., Roepstorff, A., Rees, G., Frith, C., & Tylén, K. (2012). Coming to terms quantifying the benefits of linguistic coordination. *Psychological Science*, *23*(8), 931-939.

Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, *96*(6), 3314-3319.

Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, *1*(3), 412-438.

Gow Jr, D. W., & Olson, B. B. (2015). Sentential influences on acoustic-phonetic processing: A Granger causality analysis of multimodal imaging data. *Language, Cognition and Neuroscience, 31*(7), 841-855.

Kawamoto, A. H. (1993). Nonlinear dynamics in the resolution of lexical ambiguity: A parallel distributed processing account. *Journal of Memory and Language*, *32*(4), 474-516.

Kello, C. T. (2013). Critical branching neural networks. *Psychological Review*, *120*(1), 230-245.

Kello, C. T., Beltz, B., Holden, J., & Van Orden, G. (2007). The emergent coordination of cognitive function. *Journal of Experimental Psychology: General*, *136*(4), 551-568.

Kim, N. S., Yopchick, J. E. E., & De Kwaadsteniet, L. (2008). Causal diversity effects in information seeking. *Psychonomic Bulletin & Review*, *15*(1), 81-88.

Louwerse, M. M., Dale, R., Bard, E. G., & Jeuniaux, P. (2012). Behavior matching in multimodal communication is synchronized. *Cognitive Science*, *36*(8), 1404-1426.

Lupyan, G., & Spivey, M. J. (2010). Making the invisible visible: Verbal but not visual cues enhance visual detection. *PLoS One*, *5*(7), e11452.

Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences*, *110*(35), 14196-14201.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, *101*(4), 676-703.

McRae, K., Spivey-Knowlton, M., & Tanenhaus, M. (1998). Modeling the effects of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, *37*, 283-312.

Nekovee, M., Moreno, Y., Bianconi, G., & Marsili, M. (2007). Theory of rumour spreading in complex social networks. *Physica A: Statistical Mechanics and its Applications*, *374*(1), 457-470.

Paxton, A., & Dale, R. (2013). Frame-differencing methods for measuring bodily synchrony in conversation. *Behavior Research Methods*, *45*(2), 329-343.

Prados, J. (2004). *Hoodwinked: The documents that reveal how Bush sold us a war*. New Press.

Richardson, D. C., Dale, R., & Kirkham, N. Z. (2007). The art of conversation is coordination common ground and the coupling of eye movements during dialogue. *Psychological Science*, *18*(5), 407-413.

Richardson, M. J., Marsh, K. L., & Schmidt, R. C. (2005). Effects of visual and verbal interaction on unintentional interpersonal coordination. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(1), 62-79.

Roshani, F., & Naimi, Y. (2012). Effects of degree-biased transmission rate and nonlinear infectivity on rumor spreading in complex social networks. *Physical Review E*, *85*(3), 036109.

Ryan, M. (2006). Filling in the 'unknowns': Hypothesis-based intelligence and the Rumsfeld commission. *Intelligence and National Security*, *21*(2), 286-315.

Sebanz, N., Knoblich, G., Prinz, W., & Wascher, E. (2006). Twin peaks: An ERP study of action planning and control in coacting individuals. *Journal of Cognitive Neuroscience*, *18*(5), 859-870.

Shockley, K., Santana, M. V., & Fowler, C. A. (2003). Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 326-332.

Spivey, M. J. (2016). Semantics influences speech perception. *Language, Cognition, and Neuroscience, 31, 856-859.*

Spivey, M. J., & Spirn, M. J. (2000). Selective visual attention modulates the direct tilt aftereffect. *Perception & Psychophysics*, *62*(8), 1525-1533.

Spivey-Knowlton, M. J. (1996). *Integration of visual and linguistic information: Human data and model simulations*. PhD Dissertation, U. Rochester.

Sporns, O. (2010). *Networks of the Brain*. MIT press.

Sunstein, C. R. (2016). How Facebook makes us dumber. *Bloomberg View*, Jan. 8.

Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research*, *11*(4), 378-395.

Trueswell, J. C., & Hayhoe, M. M. (1993). Surface segmentation mechanisms and motion perception. *Vision Research*, *33*(3), 313-328.

Zakaria, F. (2016). Bile, venom and lies: How I was trolled on the internet. *The Washington Post*, Jan. 14.

Zubiaga, A., Liakata, M., Procter, R., Hoi, G. W. S., & Tolmie, P. (2016). Analysing how people orient to and spread rumours in social media by looking at conversational threads. *PLoS One*, *11*(3), e0150989.