A Cognitive Model of Social Influence

J. Gregory Trafton
(greg.trafton@nrl.navy.mil)
Naval Research Laboratory
Washington, DC

J. Malcolm McCurry
(malcolm.mccurry.ctr@nrl.navy.mil)
Harris
Washington, DC

Kevin Zish
(kzish@gmu.edu)
George Mason University
Fairfax, VA

Laura M. Hiatt
(laura.hiatt@nrl.navy.mil)
Naval Research Laboratory
Washington, DC

Sangeet Khemlani
(sangeet.khemlani@nrl.navy.mil)
Naval Research Laboratory
Washington, DC

Abstract

We describe two different cognitive process models of a well known experiment on social influence (Salganik, Dodds, & Watts, 2006). One model, the *social influence model*, reproduced the choices that participants took by modeling both the cognitive processes the participant engaged in and the social influences that the participant saw. The second model, the *pure cognitive model*, used only cognitive capabilities and did not model any social influences that the participant saw. Somewhat surprisingly, the two models showed no difference in quality of fit (the pure cognitive model actually fit slightly better than the social influence model), suggesting that social influence models should take cognitive functions into account in their theories.

Keywords: Cognitive models, social influence, cognitive architectures

Introduction

People are routinely influenced by other people: this is the crux of social influence. There are many factors that can impact social influence, including the popularity of others (Cialdini, Reno, & Kallgren, 1990; Latane, 1981), authority or expertise (Cialdini et al., 1990; Milgram, 1963), and culture (Milgram, 1963). While each of these factors can have a large impact in different situations, a fourth factor, visibility -- seeing what others have done or are doing -- seems to be among the most important (Cialdini et al., 1990).

Most models of social influence describe the effect in terms of social constructs (e.g., conformity, peer pressure, etc.) and/or networks of people (e.g., families or friends), and that cognition has a relatively minor explanatory role.

For example, MacCoun (2012, in press) has proposed a very successful model of social influence called the uBOP (unidirectional burden of proof). The model itself is a mathematical model in the form

$$p = m / (1 + \exp[c(S/T - b)])$$

where p is the probability that an individual chooses an option, m is a ceiling parameter, S is the number of people advocating one option, T is the number of advocates advocating a second option, b denotes where an individual is more likely to adopt the group's decision and c reflects the difficulty to make a decision (steepness). This model can successfully characterize the classic Milgram (1969) and Asch (1951) studies with few changes in parameters (MacCoun, 2012; MacCoun, in press).

Social network models have also been used to model social influences. Social networks consist of nodes (people) who are linked through some form of interdependency (family, friends, beliefs, etc.). Social networks have been very successful at differentiating the effects of social ties from other external influences and have been applied to explain phenomena as diverse as smoking (Christakis & Fowler, 2008) and obesity (Christakis & Fowler, 2007). Social network models typically use graph theoretic or network models or statistical models (e.g., structural modeling or autoregression).

Cognition in all of these models serves, at best, a purely functional role: people perceive others' actions or remember actions that others have performed, but the theoretical power comes from social or network constructs. The uBOP model can be used to describe both individuals and groups, but has nothing that could be considered a cognitive process. Social network models describe relationships and membership rather than an individual's cognitive activities. The fact that there are few (if any) cognitive processes in these models is, perhaps, not surprising: most of the existing models are not process models. We believe that cognition is a large component of most social behavior and will explore this issue by developing a cognitive process model of a well-known social influences study (Salganik et al., 2006; Salganik & Watts, 2009).

Salganik and colleagues investigated the effects of social influence in a cultural market with a novel paradigm

(Salganik et al., 2006; Salganik & Watts, 2009). Salganik et al. created an artificial music market where participants could listen and download previously unknown songs. Salganik and colleagues created independent instantiations, or worlds, where the markets could grow without influence from other worlds; this was a between subjects manipulation. Individuals in each world could only be influenced by individuals in their own world. This approach allowed the authors to explore how social influences develop over time in different situations. Across two experiments, they looked at two conditions, an independent world and a social influence world. Participants in the independent world made decisions about what songs to listen to based only on the names of the bands and the song, while in the social influence worlds, participants could also see how many times each song had been downloaded by previous participants. In their first experiment, they found a modest social influences effect, but in their second experiment they found a very strong social influences effect. We describe and model the second experiment.

Method (Salganik et al., 2006)

A complete description of the experiment can be found in Salganik et al. (2006) and Salganik and Watts (2009).

Participants

There were 7192 participants recruited from a music website (Bolt). There were approximately 700 participants in each of eight social influence worlds and approximately 1400 participants in the independent world.[1] Participants logged onto the website and various times over 83 days.

Setup and Procedure

48 Songs were presented in a single column and sorted by the number of downloads for the social influence worlds and in a random order for the independent world. The display was updated as every participant downloaded songs. Additionally, the social influence world displays contained information about the number of downloads each song had received; this information was dynamically updated as the experiment progressed; see Figure 1.

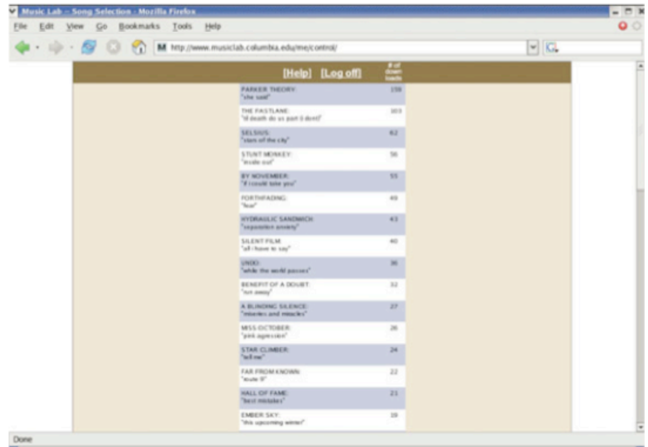

Figure 1: A screenshot of the social influence world, taken from Salganik and Watts (2009).

Participants were able to click any song on the list to listen to it. While the song was playing, participants were asked to rate the song on a 1-5 scale where 1 was "I hate it" and 5 was "I love it." After rating the song, participants could download the song and then could go back to the primary display so they could listen, rate, and download more songs if they chose. Participants were able to download as many of the 48 songs as they wished, but they had to listen to them and rate them before they could download each song.

Each of the social worlds began in a random state, so each social world could evolve based on the participants' behavior in that specific world.

Measures

There were several different variables that the authors coded in the data. The number of songs each participant listened to and the number of downloads that were made was recorded. The popularity of individual songs was also recorded. One of the most informative variables that was recorded was how often participants listened to a song at a specific rank (regardless of what song it was). When each participant examined songs, each song had a specific market rank (the song with the most downloads had a market rank of #1).

Results and Discussion

Salganik et al. reported that participants listened to an average of 3.6 songs and downloaded an average of 1.4 songs. Figure 2 shows the probability that a participant would listen to a song based on its rank market share. Note that all the social influence worlds will be combined for this and all further analyses, as reported in Salganik et al.

[1] The participant size was unbalanced because the original authors were concerned about unpredictability. For our purposes it should not impact our results.

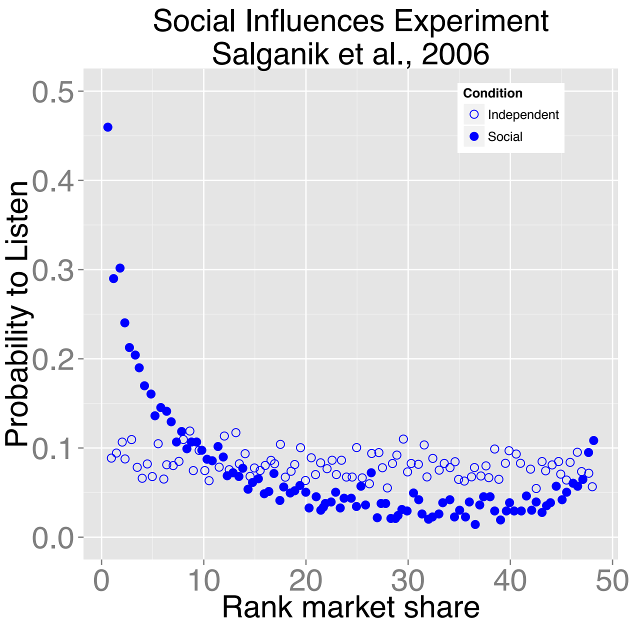

Figure 2: The probability that a participant would listen to a song based on its rank market share for both the independent and social influence conditions in Salganik et al.'s experiment.

As Figure 2 suggests, the independent condition was mostly flat, with no strong effect of either social influence or of song quality. However, the social influence conditions showed a very strong effect of social influence: the top ranked song had a 45% chance of being listened to, while the average song had only a 5% chance of being listened to. There is also an interesting "hipster" effect when the song that was ranked last got listened to a great deal more than average.

Salganik et al. suggest that these results "confirm that the popularity of the songs affected participants' choices and generally led them to listen to the more popular songs—a result that is consistent with the large literature on social influence and conformity" (Salganik & Dobbs, 2009, p. 447). They also show that while specific songs were considered better than others, the social influence condition had a substantial effect on the success of the songs.

As suggested earlier, most models of social influence are not cognitive process models. So we developed a cognitive process model of the individuals in this experiment in order to examine the effect of cognition and social influence. By developing a process model, we were able to create two slightly different models: a *social influence model* and a *pure cognitive model*. By developing the two models, we will be able to determine how much better the social influence model fits the data beyond the pure cognitive model and thus determine the importance of social influence over basic cognitive factors. The models were developed using the ACT-R architecture.

Architecture and Model Description

ACT-R is a hybrid symbolic/sub-symbolic production-based system (Anderson et al., 2004; Anderson, 2007). ACT-R consists of a number of modules, buffers, and a central pattern matcher. Modules in ACT-R contain a relatively specific cognitive faculty usually associated with a specific region of the brain. For each module, there are one or more buffers that communicate directly with that module as an interface to the rest of ACT-R. At any point in time, there may be at most one item in any individual buffer; thus, the module's job is to decide what and when to put a symbolic object into a buffer. The pattern matcher uses the contents of the buffer to match specific productions.

ACT-R uses if-then rules (productions) that will fire when their preconditions are met by matching the contents of the buffers. If there is more than one production that can fire, the one with the highest utility (production strength) will fire. Each production can change either internal state (e.g., buffer contents) or perform an action (e.g., click on a button).

ACT-R interfaces with the outside world through the visual module, the aural module, the motor module, and the vocal module. The architecture supports other faculties through intentional, imaginal, temporal and declarative modules.

High Level Description of the Social Influences Model

There are three components to each model: search, consideration, and decision-making. Each component has different productions that instantiate the specific goal. There model is a pure performance model: there is no learning in the model.

Search ACT-R has a theory about visual attention (Byrne & Anderson, 1998), which this model follows. In brief, the model searches for an unattended song, then moves its visual attention to that song and then encodes the information about the song. The model determines which song to search for in one of four ways:

(1) The model begins at the top of the display to search for an unattended song. This is typical ACT-R behavior for searching.
(2) The model finds a random song and attends to it.
(3) The model starts at the bottom of the display to search for an unattended song. This is the "hipster" component of the model.
(4) Stop searching completely and finish.

All four of these rules are in competition any time the model has a goal to look for a song. Note that if we were in a different culture where reading occurred in a different direction, the model would need to take those preferences into account. Also note that sometimes the model will have a goal to search for a new song and then give up; participants also began the study and stopped the experiment before listening to any songs.

Consideration After a song has been attended to and encoded, the model next determines whether that song should be listened to or not. It has three options:

(1) The model decides that the song "looks interesting" so it decides to listen to it. We assume that people have some preference for the name of the band or the name of the song; this is a simple version of that preference process.
(2) The model decides that the song "looks terrible" so it decides not to listen to it. Again, this is a simple way to model the preferences that people have.
(3) The song is listened to based on its rank. The probability is a very simple 1/rank. There are other, more sophisticated versions of selection based on group behavior (Mullen, 1983), but this simple version suffices for this model. Note that this is where social influence occurs in this model.

All three of these rules are in competition any time the model has a goal for considering whether to listen to the song. If the model decides not to listen to a song, it searches for another song.

Decision-making If the model does listen to a song, it must next decide whether to download it. The decision to download is very straightforward: there is a 50% chance the model will decide it should download the song. If it does download the song, the world is updated; if it does not download the song the world is not updated in any way, but the model then searches for another song.

A series of sample experimental model runs

For the following example, three models are run in the same world; the social influences model is being run in the social influences condition. We assume that each model corresponds to a single simulated participant. The world is updated based on what each model does in the world, and the world is displayed appropriately based on what others have done.

The first thing that the model does in an experiment is to search for a song. The first model in the experiment stops searching and no updating occurs.

The second model in the experiment starts at the top of the list of songs; the first song on the list is unattended, so it encodes it and considers whether to listen to it. The model decides that it will listen to it and then must decide if it should download it. There is a 50% chance the model will download it, which it does in this case. This song now becomes the most popular with rank 1 and for future participants it will show as the top song on the display. The model then searches for another song, again decides to search from the top and finds the second song, which is the top unattended song. The second song does not look interesting, so the model does not listen to it (and thus does not download it, either). The model next searches for another song, but then stops searching and this model is finished.

The third model sees the previously downloaded song in the first slot. The model, however, chooses a random song from the list and decides to listen and download it. The model next starts at the bottom of the display and looks for an unattended song. The model will listen to this song based on its rank, which is currently 48; so it has approximately a 2% chance of listening to the song. Luckily, for this run the model will listen to it, so that song now is tied for rank 1, and all future models will evaluate appropriately. After 2000 model runs, the simulation is stopped and the simulated world is reported.

The run just described was based on the *social influences model*. In this model, social influence occurs during the consideration stage.[2] The *pure cognitive model* was identical to the social influences model, except it did not have pay attention to any social influence. Without social influences, the model simply considers a song based on whether it is "interesting" or "terrible."

For all models, we kept most of the ACT-R parameter defaults. The parameters that were changed include a production noise parameter (.4, which is within a normal range for this parameter) to provide some stocasticity and the aforementioned 50% probability for downloading a specific song. Parameter fits were run using the social influence model and those same parameters were also used for the pure cognitive model.

Model fit

First, it is possible to examine how many listens and downloads each model performed and compare them to the experimental data. On average, there were 3.6 listens per participant in the experimental data; both models made 3.2 listens. Comparably, there were an average of 1.4 downloads per participant; both models had 1.6.

Figure 3 shows the fit of the independent condition; both models provide the same results. As Figure 3 suggests, the fit is quite good, with the model data overlapping a great deal with the experimental data. Calculating an R^2 is uninformative because both the data and the model are flat. RMSD for this model is .02, which demonstrates quite a good fit.

[2] Note that social influence also could have occurred at other places in the model (e.g., search). However, preliminary testing showed that the model actually fit *worse* when social influence occurred in a stage other than consideration.

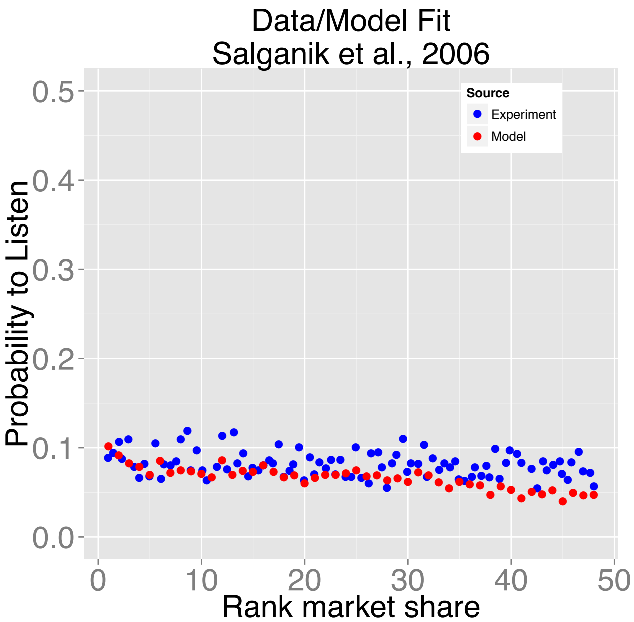

Figure 3: Data and model of the independent condition.

Figure 4 shows the fit for the social influence model and Figure 5 shows the fit for the pure cognitive model. As is evident from Figures 4 and 5, they both show quite a good fit; Table 1 shows the quantitative fit statistics. The social influence model had a very strong fit in both R^2 and RMSE. Somewhat surprisingly, however, the pure cognitive model had a slightly better R^2 fit and a comparable RMSE fit. We can conclude from these analyses that the social influence model does *not* fit better than the pure cognitive model.

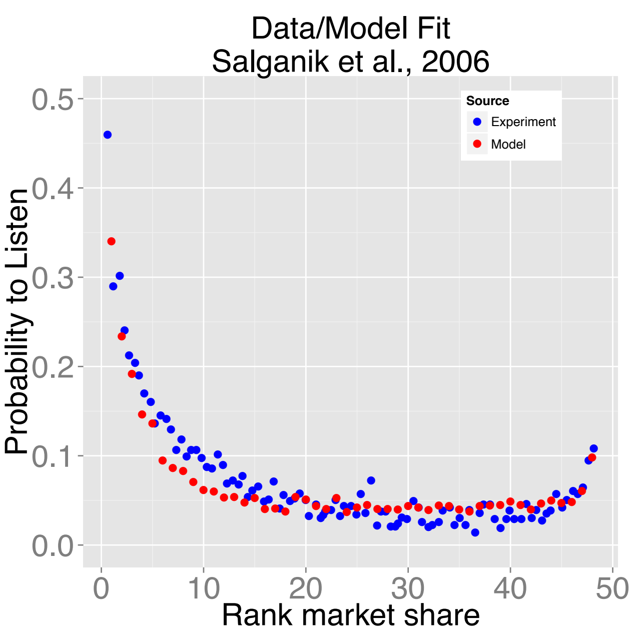

Figure 4: Experimental data from the social influence condition and the social influence model.

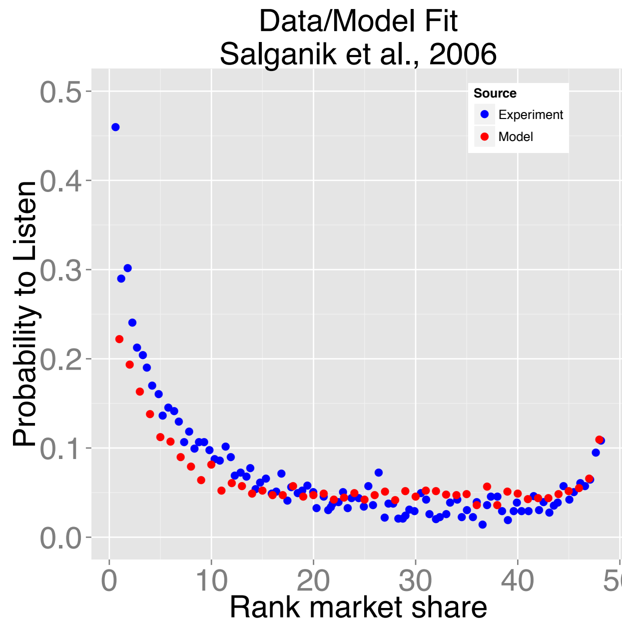

Figure 5: Experimental data from the social influence condition and the pure cognitive model..

Model	R^2	**RMSE**
Social Influence	.88	.023
Pure cognitive model	.92	.028

Table 1: Fit metrics for both the social influence model and the no social influence model. Both models were compared to the social influence condition of Salganik et al. (2006).

General Discussion

We described a process model of a well-known experiment on social influence (Salganik et al., 2006; Salganik & Watts, 2009). The experiment showed that when people had access to what others had done, it greatly influenced their behavior, consistent with current theories on social influence (Cialdini et al., 1990; MacCoun, 2012). We built two slightly different cognitive process models that perform the perceptual and cognitive steps in the experiment. Both the social influence model and the pure cognitive model fit the data extremely well. However, somewhat surprisingly, the pure cognitive model fit the experimental data slightly better than the social influence model. We interpret these results as showing that for this experiment, the effect of social influence is very small: a pure cognitive model was able to fit the data at least as well (if not slightly better) than the social influence model.

It was a bit surprising that the pure cognitive model and the social influence model shared so much overlap: this is almost assuredly one of the reasons for the similarity in the two models. This should not come as a big surprise, however: this type of task of searching and selecting objects on a computer screen is a classic cognitive task that has been investigated both experimentally and theoretically many times.

It could be argued that during the search phase, the scanning down rule is also a social influence rule since participants knew that songs were ranked from top to bottom in order of the number of downloads. However, we would argue that scanning from the top to the bottom of a list is more a cognitive and cultural function than a social influence function. Many other researchers have shown that people in the US search for objects approximately top-down and left-to-right on computer interfaces (Byrne, Anderson, Douglass, & Matessa, 1999; Norman, 1991; Schunn & Anderson, 1999).

Note that we are not saying that people are not influenced by social influence. There are many experiments and models that show the importance of social influence. For example, Cialdini et al. (1990) found that when there was evidence that other people had littered, individuals were more likely to litter than when there was evidence that people had not littered. Many other classic experiments have shown the importance of social influence (Asch, 1951; Milgram, 1963)

The model presented here does, however, highlight the importance of cognitive processes in explaining at least some social influence effects. We believe that providing a process level description of cognitive and social behavior will lead to a better understanding of how social influences impact people's behavior. Specificially, we can isolate those processes that may result from cognitive aspects of the task from those processes that result from social influence.

Acknowledgments

This work was supported by the Office of Naval Research to JGT. The views and conclusions contained in this document should not be interpreted as necessarily representing the official policies of the U. S. Navy. We thank Sunny Khemlani, Laura Hiatt, Alan Schultz, and Joe Kreke for their advice and comments on a previous draft.

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of mind. *Psychological Review*, *111*(4), 1036-1060.

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press, USA.

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Group leadership and men* (pp. 177-190). Pittsburgh, PA: Carnegie Press.

Byrne, M. D., & Anderson, J. R. (1998). Perception and action. In J. R. Anderson & C. Lebiere (Eds.), *Atomic Components of thought* (pp. 167-200). Mahwah, NJ:: Lawrence Erlbaum.

Byrne, M. D., Anderson, J. R., Douglass, S., & Matessa, M. (1999). *Eye tracking the visual search of click-down menus*. Proceedings from Proceedings of the computer human interaction conference.

Christakis, N. A., & Fowler, J. H. (2007). The Spread of Obesity in a Large Social Network Over 32 Years. *New England journal of medicine*.

Christakis, N. A., & Fowler, J. H. (2008). The Collective Dynamics of Smoking in a Large Social Network. *New England journal of medicine*.

Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A focus theory of normative conduct: Recyling the concept of norms to reduce littering in public places. *Journal of personality and Social Psychology*, *58*, 1015-1026.

Latane, B. (1981). The psychology of social impact. *American psychologist*, *36*, 343-356.

MacCoun, R. J. (2012). The burden of social proof: Shared thresholds and social influence. *Psychological review*, *119*, 345-372.

MacCoun, R. J. (in press). Computational models of social influence and collective behavior. In *Computational Models in social psychology.* New York: Psychology Press.

Milgram, S. (1963). Behavioral study of obedience. *The Journal of abnormal and social psychology*, *67*(4), 371.

Mullen, B. (1983). Operationalizing the effect of the group on the individual: A self- attention perspective. *Journal of Experimental Social Psychology*, *19*, 295-322.

Norman, K. L. (1991). *The psychology of menu selection: designing cognitive control of the human/computer interface.* Norwood: Ablex.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*, 854-856.

Salganik, M. J., & Watts, D. J. (2009). Web-based experiments for the study of collective social dynamics in cultural markets. *Topics in Cognitive Science*, *1*, 439-468.

Schunn, C. D., & Anderson, J. R. (1999). The generality/specificity of expertise in scientific reasoning. *Cognitive Science*, *23*(3), 337-370.