The Spread of a Meme Across a Social Network

Kristen Amaddio
EECS 372 – Agent Based Modeling
8 June 2015

Overview

Throughout human history, trends have played a large role in determining in- and out-groups in society. Trends could be anything from fashionable clothes to popular vernacular. Trends pervade every facet of human life, determining which style of music or other interests are socially acceptable. Trends not only affect the ways humans express themselves, but also the products they buy. Trends come in many types, and humans are exposed to them by interacting with friends, observing strangers, consuming media, and browsing the Internet.

Today, as people become more and more connected over the Internet, and as they begin to live significant parts of their lives on the web, Internet memes have become a huge category of trend. They spread across these networks from person to person, and an individual often sees the same meme several times from different sources. A meme could be anything from an inspirational video to a cute cat photo to a silly photo with modifiable text over it. Memes have become part of the human culture, and, in their spread across a social network, memes have transformed into a variety of "in-group" inside jokes shared across the connected culture of humans on the Internet.

This model simulates the spread of a meme across a social network. It explores which factors are the most influential in the success of the meme (i.e. how long it takes the meme to spread): social interactions, media consumption, popularity of the trend-setter, and the inherent subject matter of the meme.

Motivation and Rationale

While it could be modified to incorporate about any type of trend, this model focuses specifically on the spread of a meme. Since Internet culture is still a relatively new phenomenon, it is interesting to concentrate exclusively on the behavior it fosters. Humans behave differently on the Internet than they do in real life; not only does it allow them to express themselves more freely, but it also allows them to keep in touch with many different people at any time.

Since the model simulates a person's Internet interactions, it is most realistic and reasonable to construct a social network structure. The use of agent based modeling in this case will show how each person in the network will be influenced by their friends and the media. Each individual will be interested in a single category of subject matter, and each meme will fall into one of those categories. Based on his or her own interests, a person will make decisions about which meme he or she will want to take on.

What Can Be Learned

Since this model is a simplified simulation of the real-world phenomenon of a spreading meme, it is not intended to fully explain the phenomenon, nor does it claim to be a ruling authority in the evidence of the culture of an actual social network. The purpose of this model is to explain some of the basic factors that could influence the spread of a meme across individuals in a social network. Specifically, it hopes to determine which of the following factors affect the spread most strongly:

- Social interaction with friends
- Media influence and frequency
- Popularity of the "trend-setter", the person who starts the meme
- Inherent subject matter of the meme compared to a person's interests

Model Implementation

This model consists of turtle, link, and patch agents. Turtles represent people, links represent the relationships between the people, and patches represent the media. When the network is first generated, each person is randomly assigned an "interest-category", an integer ranging between 0 and 9, that represents the type of subject matter in which that person is interested. Each person is also initialized as blue to show that they are not currently following the meme (i.e. they are not "trendy"). A person's "popularity" is the exact value of their degree, or how many links, or relationships, they have in the social network.

Setting Up the Network

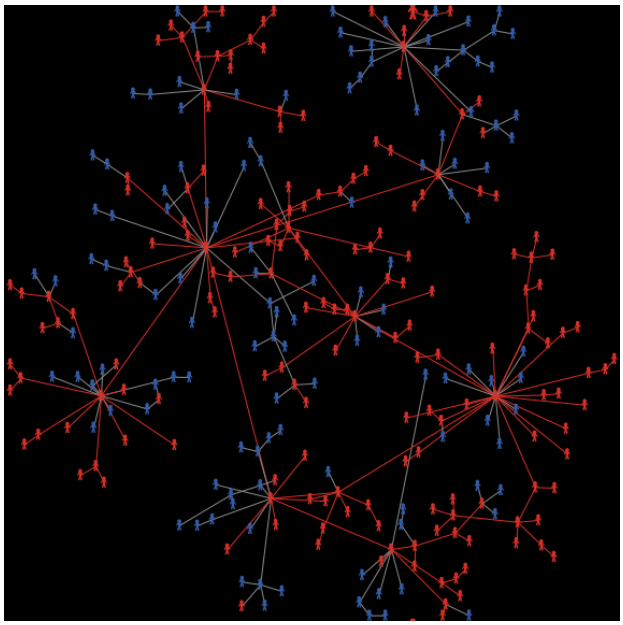

Figure 1: An example of the preferential attachment network used in this model. Notice how most nodes have a degree of one, while a few nodes have very large degrees.

First, a social network must be created. The "preferential attachment"[1] method is used to most closely model this network. This method begins as two nodes connected by an edge. Then, at each step, a node is added. The new node chooses to connect to an existing node randomly, but with a bias proportional to the number of connections, or degree, an existing node already

has. A new node prefers to connect to one of the most connected existing nodes. So the higher the degree of an existing node, the more likely it is to receive more connections. This process continues until the set "population" number of nodes have been created.

Seeding the Meme

Next, it is time to start a meme. A random person is seeded as the "trend-setter", or the designated person to begin the spread of the meme. This person turns red, a visual indication that they follow the meme (i.e. they are trendy). A meme has a "trend-category", an integer ranging from 0 to 9 that directly corresponds with the "interest-category" values of people. The current meme's "trend-category" will be he exact value of the trend-setter's "interest category" because the system works under the assumption that a person would create a meme in the same category as their own interests.

Spreading the Meme

Finally, the meme must spread across the social network. The first-person pseudo-code for the rules for each person is as follows:

At each clock tick:

1. If I am "trendy", then I will try to spread the meme to one of my linked neighbors.
2. If I am not "trendy", and one of my linked neighbors is trying to spread the meme to me, I will have a probability of accepting the meme that is proportional to the difference between the "trend-category" and my own "interest-category".
3. If I accept the meme, I become "trendy" and change color to visually represent this.
4. If both my neighbor and I are "trendy", the link between us changes color to express this.

As mentioned in Step 2 above, the success of the spread of a meme to a person depends on the closeness of the trend's "trend-category" to the person's "interest-category". This means

that the closer the “trend-category” is to the “interest-category”, the more likely the person is to follow the meme and become trendy at that time. So realistically, if a meme is about a subject matter in which the person is interested, the person will be more open to accepting and spreading it.

Media Influence

The interface includes a “media?” switch that can be toggled on or off. If the switch is turned off, then the spread of the meme is based purely on the social interactions between the people. If the switch is turned on, then the user can select a “media-frequency” value that determines how frequently the media flashes in the social network. The rules for the media aspect of the code are given below, from the person’s point of view:

1. If “media?” is turned on, every “media-frequency” ticks, a random patch with people on it flashes white, becoming a “TV”.
2. The TV patch will select one of the people on it as its target.
3. If I am the target of the TV patch, and I am already “trendy”, then nothing will happen.
4. If I am the target of the TV patch, and I am not “trendy”, then I will accept the meme, become “trendy”, and change color to visually represent this. I will then follow the typical person behavior outlined in the rules in the previous section.

Analysis

It is important to note that, due to the single-component network structure, the meme will always inevitably spread to all people at the model’s time of completion. This is because every person is connected to at least one other person in the social network, and while people may initially turn down the meme, with enough tries and influence from their friends, they will eventually accept it. So rather than measuring the percentage of the population that follows the

meme in the end, it will be more telling to measure the most influential source of the meme for each individual, as well as the time it takes for the meme to spread across the entire network.

The Influence of Social Versus Media Factors

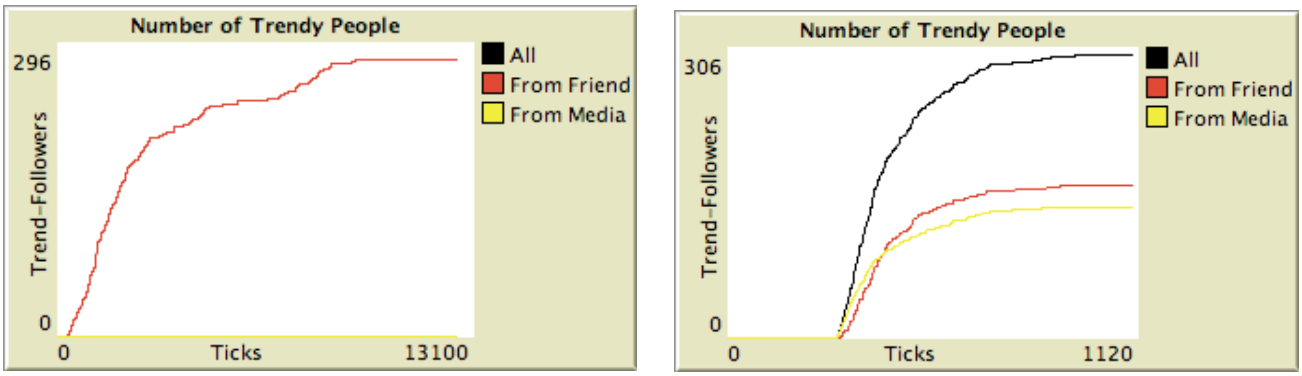

Figure 2: The spread time of a meme in a network with no media influence (left) and a network with a media frequency of once every 2 ticks (right).

Figure 2 shows how long, in ticks, it takes for a meme to spread across the entire network when there is no media influence versus when there is the highest frequency of media influence. In both cases, the population was 300 people, and the popularity of the trend-setter was 3 degrees. With all other factors kept equal, it is clear that the media frequency expedited the spread of the meme. This makes sense because the media influence allows new sources of the same meme to be spread at random parts of the network, rather than the meme having to make its path from one single trend-setter node.

Figure 3 also appears to support this claim, showing that as media frequency decreases, the amount of time taken for the meme to completely spread across a network generally increases. The findings in Figure 3

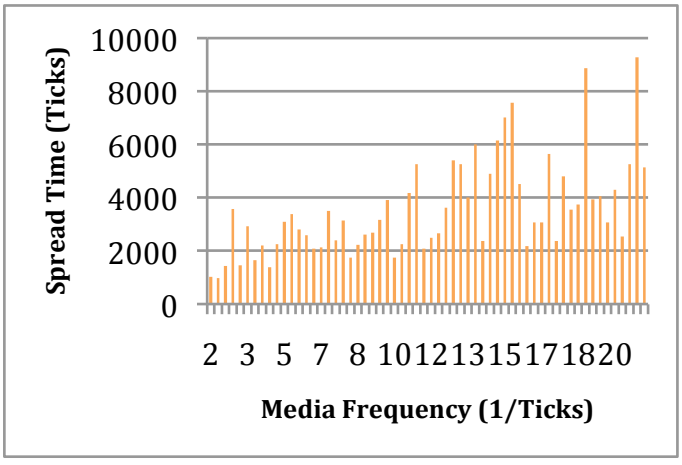

Figure 3: The effect of media frequency on the meme spread time completion in a network with a population of 300 people.

are not necessarily completely definitive because there are some cases where a run with a lower media frequency results in lower ticks to completion than a run with a higher media frequency, but it can be hypothesized that with all other things equal, such as the popularity of the trend-setter and the layout of the network, a higher media frequency will generally result in a faster spread of the meme.

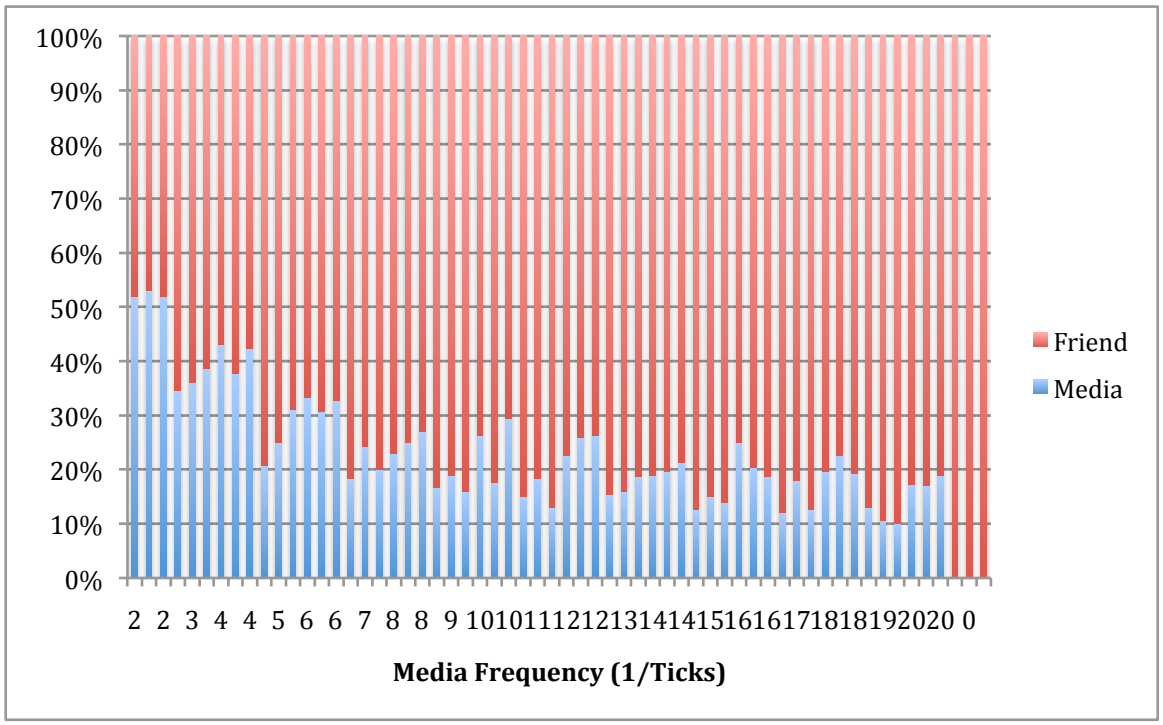

Figure 4: The effect of media frequency on the sources of meme influence on people. Red shows the percentage of people who accepted a meme from a friend, and blue shows the percentage that accepted it from the media. Note that the three rightmost data represent spreads with zero media influence.

Figure 4 shows the percentage of people that accept the meme from each source: a friend or the media. It is clear that a higher media frequency generally correlates with a higher percentage of people that accept the meme from the media instead of a friend. However, even with the media flashing at its higher frequency, the media never particularly overtakes the influence of friendships in terms of spreading the meme. At its most frequent occurrence, media

influences about half of the people in the network, and when it is decreased even a small step, friend influence completely wins over the majority of the people in the social network.

The Popularity of the Trend-Setter

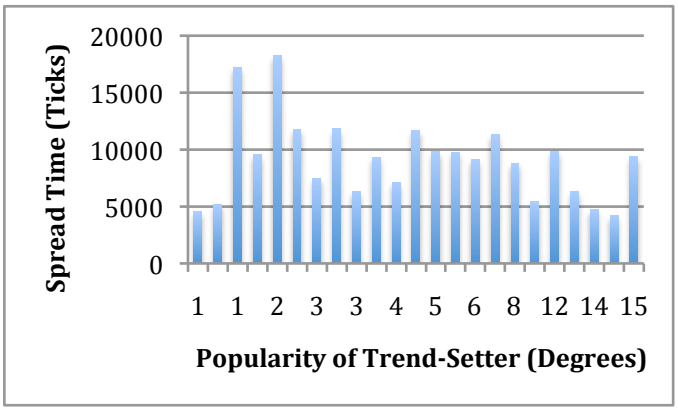

Figure 5: The effect of the trend-setter's popularity on the time it takes for a meme to spread across the network.

Figure 5 shows the results of an experiment in which all facets of the network were kept equal, except for the popularity of the trend-setter. These runs measured how many ticks it took for a meme to spread across an entire network of 300 people. The figure shows that the popularity of the trend-setter does not necessarily have a clear effect on the completion time of the spread. Some of the one-degree trend-setters had significantly shorter completion times than other one-degree trend-setters, as well as some of the higher-degree trend-setters.

It is most likely that other characteristics of the network are at play here. Since the experiment measured the completion time of the full spread of the meme, rather than measuring how many people were "trendy" after a given number of initial ticks, it was more of a measurement of the entire network rather than just the initial influence of the trend-setter. Due to the architecture of the preferential attachment network, there will likely be some roadblocks where a more central node will be highly resistant to a meme, taking several ticks to finally accept it and pass it on to more nodes. This sort of obstruction would cause the meme to take longer to spread to completion.

The Effects of the Interest-Categories of Trend-Followers

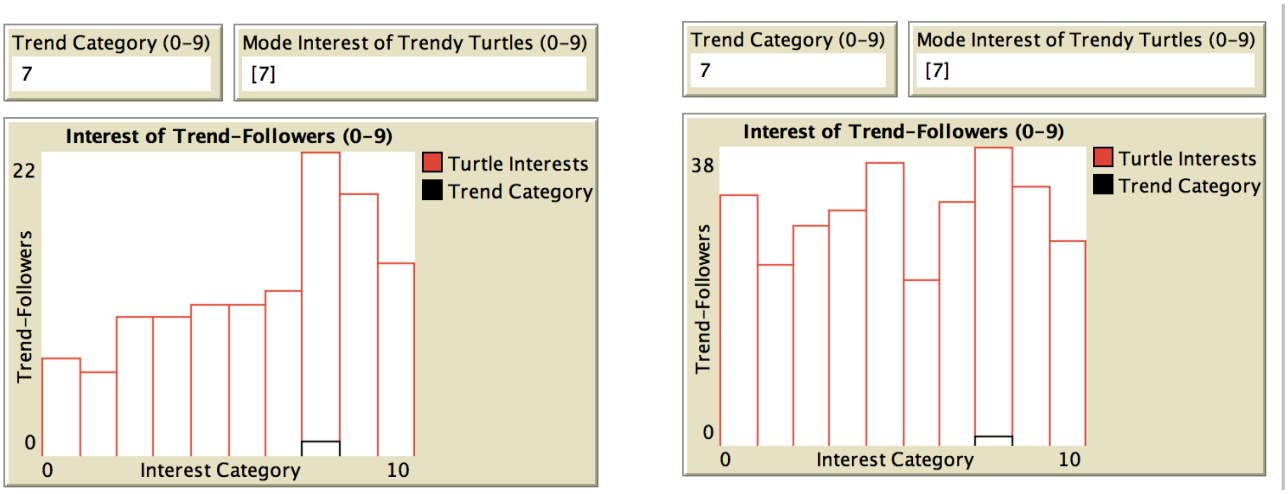

Figure 6: The interest-categories of trend-followers compared to the trend-category at 41.2% trendy (left) and 100% trendy (right).

Figure 6 shows the effect a person's interest-category has on their meme acceptance. The people with interest-category values that are closest to the trend-category value are most likely to accept the meme earlier on. So at the beginning of a spread, the distribution of the interest-categories of trendy people generally shows higher values closer to the trend-category value, and lower values farther away from the trend-category. While the interest-category is important in the beginning of the meme's spread, as the spread reaches completion, the distribution will deviate from the trend-category because all people in the network will eventually accept the meme.

Conclusion

This model has revealed some interesting behavior of the spread of a meme across a social network. First, it has demonstrated that a person's social interactions are more influential than the media in terms of their willingness to accept a meme. While the inclusion of the media generally speeds up the meme-spreading process by allowing the meme to reach more people at

once, it does not necessarily impact the an individual more than that individual's friends. Overall, an individual is more likely to accept a meme from a friend rather than the media.

Multiple experimental runs with this model show that the popularity of the trend-setter does not necessarily influence the time it takes for a meme to spread. While the popularity of the trend-setter might impact the initial spread speed of a meme, other network factors must be taken into account when analyzing the total spread time. Regardless of who starts the meme, the many properties of the network structure will all affect the ultimate number of ticks it takes for a meme to spread.

Finally, a person's interest-category does influence their decision to accept a meme, causing them to be more resistant to memes that are farther away from their interests. This is most obvious at the beginning of the spread when the distribution of interest-categories tend to favor the categories nearest to the trend-category. However, as a spread reaches completion, all people will eventually follow the meme, causing the overall distribution of interest categories to deviate from the trend-category. As a person hears a meme several times from their neighbors, they will eventually be persuaded into accepting it.

References

[1] Wilensky, U. (2005). NetLogo Preferential Attachment model. http://ccl.northwestern.edu/netlogo/models/PreferentialAttachment. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

[2] Wilensky, U. (1999). NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.