Diffusion: The problem of tracking ideas

Jean Mark Gawron

1 Introduction

Let's start with the notion of a sign. A sign is something with a meaning and a form of expression, say, a word. Signs don't have to be linguistic; they may be images, pieces of music, bodily ornamentations, or articles of clothing. In these latter cases, it may be difficult to articulate what is "meant" by the sign, but let us take meaning in a very broad sense, to include especially things like signaling membership in a particular group.

Assuming the idea of a sign, I will draw a distinction between ideas and memes, although in the discussion below I will often have occasion to refer to shared properties. By an idea I mean something that can be meant by a sign, although, rightly or wrongly, I will speak of ideas as if they can be divorced from signs. By a meme I mean something that is very like a sign; that is, it has a form (a sequence of words, a particular image), and an associated content. For my purposes, the chief difference between a meme and a sign is that memes can be quite extended. A 300-word Facebook post can be a meme. That meme expresses one idea, although the idea may be complex. A very different Facebook post with very different words might be used to express the same idea. Thus, we would have two memes and one idea.

A key premise of social science is that both ideas and memes have beginnings. Until very recently, this has been an assumption made out out of necessity, but observational confirmation has been lacking. More importantly, observation of the process by which social innovations unfold, spread, and take hold has been lacking. For example, the theory that the spread of a new idea might be epidemic-like, beginning with a long period of relatively slow development among a small community of early adopters, or even stasis, followed by a sudden period of explosive growth, had relatively few case studies for support. Now, with the web and social media we have the opportunity for massively documented studies tracking every stage in the establishment of an innovation. Thus, Yee et al. (2007) and Friedman et al. (2007) study the emergence of conventions governing avatars in the online community Second Life. Kooti et al. (2012) meticulously documents the spread of the retweeting convention of prefixing "RT" to a retweet.

Because we have a complete record of all tweets since the dawn of Twitter, we know the identity of the first tweeter to use "RT" (@TDavid); we know the date (01/25/08); we know who the early adopters were and who their followers were. We know there was period of long slow growth; we also know there were a number of memes competing to signal the same idea throughout that period, and that three were dominant. This means we can study the structure of competing diffusion networks and determine how many "parents" (influencers) adopters had. This structure strongly suggests that early adopters had multiple exposures before becoming adopters. Thus the particular variety of transmission model that is relevant is one involving what is called *complex contagion*, which in turn may require bridges between network clusters that are "wide enough to transmit strong social reinforcement" (Centola and Macy 2007). Kooti et al. (2012) also shows that the constraints of the medium were relevant. The first appearance of *RT* was in a message right at the 140 character limit; and DaveT had until then been usinng the fully spelled out *Retweet* variant.

Studies like Kooti et al. (2012) are foundational works in a new science, with results of much relevance for old sciences like Sociology, Political Science, and Linguistics. It remains to be seen how much the example of a Twitter innnovation can tell us about us phenomena like new words, the spread of extremism, urban myths, or new product adoption; but case studies like this one provide opportunities to observe components of theories of all of these. Lab experiments have previously been used only to study aspects of the innovation process at the micro level (Wilkes-Gibbs and Clarke 1992, Selten and Warglien 2007) and mathematical models and computational simulations (Boyd and Richerson 1995, Walker and Wooldridge 1995) have been used for study at the macro level. What social networks and online communities final offer is a laboratory in which we can relate the macro level and micro level.

2 Diffusion

Diffusion Theory seeks to explain how ideas or memes spread. The premise is the assumption of the previous section. There is an innnovation, a moment of birth, and suddely the idea or meme has a set of early adopters. A central question of diffusion theory is: What happens baxt? What are the necessary and sufficient conditions of adoption. What brings one to the tipping point? Centola and Macy summarize a key result and a set of classical studies that support it:

> People may hear about a movement to "think globally, act locally," but it is when they see people they know getting involved

> that they become most susceptible to recruitment. Similarly, many people may hear about a new fashion, but it is not until they see their friends display it that they are persuaded to go along (Crane 1999). From hybrid corn (Ryan and Gross 1943) to medical innovations (Coleman et al. 1966), the pattern is well-documented. The decisive event is not hearing about an innovation, but observing enough people participating to be convinced that the innovation should be adopted (James 1990, Simmel and Wolff 1950, Rogers 1995).

Bakshy et al. (2012) cite a body of evidence showing that the result is robust in the online world, but they also articulate the key problem in interpreting it.

> One particularly salient characteristic of diffusion behavior is the correlation between the number of friends engaging in a behavior and the probability of adopting the behavior. This relationship has been observed in many online contexts, from the joining of Live- Journal groups ... to the bookmarking of photos ... and the adoption of user-created content ... However, as Anagnostopoulos et al. (2008) point out, individuals may be more likely to exhibit the same behavior as their friends because of homophily rather than as a result of peer influence. Statistical techniques such as permutation tests and matched sampling ... help control for confounds, but ultimately cannot resolve this fundamental problem ...
> – Bakshy et al. (2012)

Thus, key problems are inferring relationships that lead to the *decision to adopt*, and identifying key properties of networks that promote widespread adoption. Centola and Macy, for example, argue that "[network] structures that are highly efficient for the rapid dissemination of information are often not conducive to the diffusion of collective action based on the information."

At this point I would like to point out that much of the recent groundbreaking work on diffusion has either been on theoretical issues of network structure or on memes. This has been matter of necessity. The problem of working on the spread of ideas is that a single idea can be cloaked in many forms, and this raises numerous practical issues for a detailed diffusion study. Let me emphasize this point by returning to the simple example of *RT*, which is actually an exception. On careful inspection, it is not quite so simple. First of all, Kooti et al. (2012) do not simply study a meme;

this is because, as their title declares, they are studying the emergence of a *convention*, a standardized way of getting a particular thing done. To do this, they track a cluster of memes (which they call *variants*) *competing* to express the same idea. Thus, they collect data on *RT, R/T, HT, via,* ♲*, Retweet,* and *Retweeting*. This already requires a fair amount of careful "fieldwork" (eyeballing data), since you can't do a disjunctive Twitter search until you know all the disjuncts you're searching for. But during that fieldwork, the decision had to be made, "Does this item before me express the relevant idea? Is this a legitimate competitor for the relevant part of meaning space?" That was not a trivial decision.[1] Second, the cluster of memes they study do not in fact always express the same idea. As they point out, there are really two types of construction, one typically followed by an "@" and the source of the retweet, the other not followed by anything. Thus the function of the first is attribution, the function of the second to annnounce that something is being passed on (it is a copy). The ideas are related, but they are not the same idea. The RT that eventually emerged as an established convention can be used both ways.

The two issues that arise with the RT study are the two fundamental difficulties in tracking ideas. First, ideas come cloaked in many forms. If one cannot simply write a regular expression or a google query to capture the idea, what is the best way of coming up with an operational definition (a disjunction of variants)? So far, careful fieldwork seems to be the best answer. Second, the forms that express an idea may also serve to express related ideas. Sometimes these related ideas are of interest, sometimes not. And sometimes we have hybrids of problem one and problem two. That is, what appears to be a variant is actually expressing a "slightly" different idea, often intentionally so, as is most extremely the case in parodies, which by definition resemble what they are ridiculing.[2]

A very interesting set of transitional examples is given in Simmons et al. (2011), which discusses the phenomenon of *meme mutation*, the systematic alteration of meme forms, even in newswire contexts where direct quotation is asserted. What they show is quotation errors made early in the news cycle are often propagated faithfully through a number of newssources, demonstrating that few of them are actually returning to the source for their data.

[1] No doubt it helped that variants were often but not always followed by a *source*, a website or Twitter feed.

[2] See for example the *Dove Evolution* You-Tube video `http://www.youtube.com/watch?v=qhib8XiDc9Y&feature=related` and the parody *Slob Evolution* `http://www.youtube.com/watch?v=kpXInB9aV7A`.

This is a work whose intention is to study memes, not ideas, but which demonstrates quite clearly that the problem of variants inevitably arises in all diffusion studies.

All these problems become far worse when we turn to the study of ideas like climate change or Hurricane Katrina. Here the most promising line of research seems to be one that takes us one level of abstraction away from the words themselves. Topic Detection and Tracking (TDT) is an area of research that has focused on finding temporally and thematically similar events in large news corpora, as a tool for helping analysts cope with vast amounts of data. A number of researchers have used Latent Dirichlet Allocation (LDA) as a way of clustering similar documents (Blei et al. 2003, Heinrich 2005, Griffiths and Steyvers 2004, Pan and Mitra 2011), or alternatively, of grouping documents according to their topics. The idea is closely related to the idea of Latent Semantic Analysis (LSA) and probabilistic LSA (Deerwester et al. 1990, Hofmann 1999), but it has undergone a number of computational improvements along the way. In LDA, a topic is a probability distribution over words; words which have a high probability of occurring in the topic are words which often co-occurred in the training set. Thus, the abstraction is to move away from words to their co-occurrence patterns, and to treat as alike words whose co-occurrence statistics are similar. What I am proposing then, is that in order to track ideas, we have to treat ideas as the TDT community has treated topics, as probability distributions over words. More precisely for complex ideas like climate change, I propose we treat an idea as a set of "topics", each a probability distribution. The first thing to try is obviously LDA itself.

The paradigm suggested by the TDT work is to assemble a group of training documents of interest, expressing one point of view, or idea, and to choose a number of topics, say 7, and train an LDA model on that data. The number 7, then, captures the set of variants, of interest. To use the models, we train alternative LDA models on different data sets expressing different ideas. One model, for example, can be trained on pro climate change documents, another on anti climate change documents. Since a document in LDA is represented as a mixture of topics, we can apply both models to new test documents and compute which "idea" model best explains a new document.

3 Conclusion

The new field of Diffusion studies tracks the spread of ideas annd memes. In both cases a kind of slippage has to be dealt with. First, when studying memes, we can not always rely on a form to be used in the same way; ambiguity creeps into even the simplest sign systems, as the example of *RT* shows; RT is used for both attribution and notification that a quotation is being transmitted. Second, memes mutate, and such motations are worth inncorporating into our diffusion networks. In addition, the mutation processes may become part of what we are studying, since "fitter" variants may diffuse more efficiently.

With ideas the problems grow exponentially worse. The number of variations we have to deal with is large, and it becomes impossible to draw boundaries bectween closely related ideas. Nor is it always clear that we want to, if closely related ideas diffuse inn similar ways. I have thus argued that "soft" idea models are appropriate, and put forth some speculative thoughts about drafting LDA models to fill this role.

References

Anagnostopoulos, A., R. Kumar, and M. Mahdian. 2008. Influence and correlation in social networks. In *Proceedings of the 14th International Conference on Knowledge Discovery and Data Mining*, 7–15. New York, NY.

Bakshy, E., I. Rosenn, C. Marlow, and L. Adamic. 2012. The role of social networks in information diffusion. *Arxiv preprint arXiv:1201.4145*.

Blei, D.M., A.Y. Ng, and M.I. Jordan. 2003. Latent dirichlet allocation. *The Journal of Machine Learning Research* 3:993–1022.

Boyd, R., and P.J. Richerson. 1995. Why does culture increase human adaptability? *Ethology and Sociobiology* 16(2):125–143.

Centola, D., and M. Macy. 2007. Complex contagions and the weakness of long ties1. *American Journal of Sociology* 113(3):702–734.

Coleman, J.S., E. Katz, and H. Menzel. 1966. *Medical innovation: A diffusion study*. Indianapolis, IN: Bobbs-Merrill.

Crane, D. 1999. Diffusion models and fashion: a reassessment. *The Annals of the American Academy of Political and Social Science* 566(1):13–24.

Deerwester, S., S.T. Dumais, G.W. Furnas, T.K. Landauer, and R. Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American Society for Information Science* 41(6):391–407.

Friedman, D., A. Steed, and M. Slater. 2007. Spatial social behavior in second life. In *Proceedings of the 7th International Conference on Intelligent Virtual Agents (IVA 2007)*, 252–263.

Griffiths, T.L., and M. Steyvers. 2004. Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America* 101(Suppl 1):5228.

Heinrich, G. 2005. Parameter estimation for text analysis. *Web: http://www.arbylon.net/publications/text-est. pdf.*

Hofmann, T. 1999. Probabilistic latent semantic indexing. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, 50–57. ACM.

James, Coleman. 1990. *Foundations of social theory.* Cambridge, MA: Harvard University Press.

Kooti, F., H. Yang, M. Cha, K. Gummadi, and W.A. Mason. 2012. The emergence of conventions in online social networks. In *Proceedings of the 6th iternational AAAI conference on weblogs and social media*, 194–201. AAAI.

Pan, C.C., and P. Mitra. 2011. Event detection with spatial latent dirichlet allocation. In *Proceeding of the 11th annual international ACM/IEEE joint conference on Digital libraries*, 349–358. ACM.

Rogers, Everett M. 1995. *Diffusion of Innovations.* New York: Free Press.

Ryan, B., and N.C. Gross. 1943. The diffusion of hybrid seed corn in two iowa communities. *Rural sociology* 8(1):15–24.

Selten, R., and M. Warglien. 2007. The emergence of simple languages in an experimental coordination game. *Proceedings of the National Academy of Sciences* 73617366.

Simmel, Georg, and Kurt H. Wolff. 1950. *The Sociology of Georg Simmel.* Glencoe, IL: Free Press.

Simmons, M.P., L.A. Adamic, and E. Adar. 2011. Memes online: Extracted, subtracted, injected, and recollected. *Proceedings of International Conference on Weblogs ad Social Media ICWSM 2011.*

Walker, A., and M. Wooldridge. 1995. Understanding the emergence of conventions in multi-agent systems. In *Proceedings of the First International Conference on Multi-Agent Systems*, 384–389, San Francisco, CA. MIT Press.

Wilkes-Gibbs, D., and H. Clarke. 1992. Coordinating beliefs in conversation. *Journal of Memory and Language* 31(2):183–194.

Yee, N., J.N. Bailenson, M. Urbanek, F. Chang, and D. Merget. 2007. The unbearable likeness of being digital: The persistence of nonverbal social norms in online virtual environments. *CyberPsychology & Behavior* 10(1):115–121.