

The Emergence of Conventions in Online Social Networks

Farshad Kooti
MPI-SWS

Haeryun Yang
KAIST

Meeyoung Cha
KAIST

Krishna P. Gummadi
MPI-SWS

Winter A. Mason
Stevens Institute of Technology

Abstract

The way in which social conventions emerge in communities has been of interest to social scientists for decades. Here we report on the emergence of a particular social convention on Twitter—the way to indicate a tweet is being reposted and to attribute the content to its source. Initially, different variations were invented and spread through the Twitter network. The inventors and early adopters were well-connected, active, core members of the Twitter community. The diffusion networks of these conventions were dense and highly clustered, so no single user was critical to the adoption of the conventions. Despite being invented at different times and having different adoption rates, only two variations came to be widely adopted. In this paper we describe this process in detail, highlighting insights and raising questions about how social conventions emerge.

Introduction

To function properly, societies require massive amounts of coordination. This is especially true with complex and expansive modern industrial societies, but is also true even with much smaller groups of people. In some cases, these coordination problems are solved by institutions, for instance, the establishment of the International System of Units (the metric system) for units of measurements. Many times, however, there is no established institutional code or even an institution, in which case a *social convention* emerges. Social conventions guide us through complex interactions: eating dinner out (differing between restaurants and friends' homes), driving (on which side of the road to drive), altruistic gift exchange, etc.

Most definitions of “social convention” suggest that a behavior must occur with regularity and be common to at least some members of the society to be considered a social convention (Delgado 2002; Shoham and Tennenholtz 1997). Social conventions are distinct from social norms, as social conventions do not have a proscriptive component; that is, someone who does not follow a convention may be seen as eccentric or different, but not as a bad member of society. In fact, it is possible for conventions to evolve into

norms: once a social convention has been sufficiently established in a community, members of the community may begin to chastise those who do not utilize the convention. This rebuke serves multiple purposes; it reinforces the system that acts as a guide in social interactions (Feldman 1984; Opp 2001), enhances group identity and feelings of belongingness (Riley and Burke 1995), and facilitates group performance on tasks by establishing conventional roles (Levine and Moreland 1990).

Social conventions can be established in many different ways. We focus on the type that *evolves*, arising from interpersonal interaction and spreading from person to person within a community (Opp 2001). While there are many studies on the top-down process by which institutions come to create and enforce social norms, we are interested in the bottom-up process. How do conventions arise naturally? Who are the inventors and the early adopters of the convention? How do conventions spread from individual to individual? Do some conventions come to dominate over others?

Previously, researchers have taken one of three approaches to answer these questions. One approach is to focus on the micro-level processes with small-scale laboratory experiments (Wilkes-Gibbs and Clark 1992). Another approach has been to create mathematical models that embody theoretical principles and generate large-scale predictions that resemble real-world outcomes (Boyd and Richerson 1995). A third way is to actually observe the emergence of norms in naturally occurring groups. Since capturing the natural emergence of a social convention can be extremely difficult, this third approach has rarely been employed. Only recently has this become more feasible, as conventions have begun to be studied in online communities (Yee et al. 2007; Friedman, Steed, and Slater 2007).

In this research, we employ the third method with data gathered from a large online social network. We observe the emergence of a social convention and document the characteristics of the convention, the environment, and the dynamics of competing conventions that led to only two candidate conventions becoming widely popular. Specifically we focus on the online community in Twitter and the convention used to indicate that one is reposting a tweet to one's followers and attributing the content to the source—an act commonly known as “retweeting” (boyd, Golder, and Lotan 2010). Utilizing a near-complete collection of tweets, we

track the birth of the retweeting convention over the first 3.5 years of Twitter's existence. Several different *variations* of this convention emerged during the first few years of Twitter, some with great success and some with little impact. These variations arose organically and became widely adopted by many individuals and third-party applications over time, until (and even after) Twitter rolled out an official, built-in "retweet" button in November 2009.

There are several unique advantages that come from studying the retweeting convention. To begin with, before any particular form of the convention was established on Twitter, a user would most likely first encounter a particular variation through his or her contacts on Twitter. Consequently, another advantage of studying retweeting is that one's contacts on Twitter are explicit, so we can observe how a convention spreads from one user to another. A third advantage is that the convention is not usable outside of Twitter, so nearly all uses of the convention will have been on Twitter, limiting potential exposures outside of the environment. Because we have nearly all of the tweets, this means we also have nearly all of the uses of the convention. In other words, by studying this particular convention we can explore not just the micro-level processes of adoption or the macro-level outcomes, but actually explore how one led to the other.

Some of our key results are as follows:

1. Several retweeting conventions arose organically, because of the perceived need to forward other people's tweets efficiently in Twitter.
2. Early adopters of the retweeting conventions were more active and well-connected than the remaining adopters or typical users. They are genuinely influential or core users, who also adopted other new features of Twitter early.
3. A great majority of early and later adopters had a fellow Twitter friend who adopted the same convention prior to them. This demonstrates that conventions mostly spread through the internal social links in Twitter.
4. The conventions spread through a dense network, so there were no "bottleneck" users who were critical to the spread of the convention.
5. As the conventions continued to spread, only two of them (RT and via) spread past the boundary of the core users, a property that we have identified as a necessary characteristic feature of the conventions that come to dominate.

Related Work

The theory on the emergence of social conventions can be traced to the earliest studies of social influence. One of the first comprehensive reviews outlined a theory of how group norms develop and are maintained (Feldman 1984). This theory stated that norms can develop top-down, bottom-up, or through a critical event in the group's history, but the maintenance of a norm depends on whether it is instrumental; that is, if it helps the group continue to exist. A more recent review on norm formation has a slightly different view (Opp 2001). It says that the "instrumentality hypothesis" as suggested in (Bettenhausen and Murnighan 1985;

Levine and Moreland 1990) should be refined so that norms created top-down serve an explicit, group-maintaining function, while emergent norms typically are fulfilling some other internal group function—still instrumental, but not in the way initially suggested.

One approach to study the emergence of conventions is through small-scale experimental work. One of the first to explicitly investigate how people solve coordination problems focused on linguistic conventions for reference. In (Wilkes-Gibbs and Clark 1992), participants were asked to develop short-hand verbal code for referring to images that allowed them to communicate more efficiently and therefore complete tasks with the images more quickly. In a similar paradigm (Selten and Warglien 2007), participants attempted to coordinate by sending messages that "describe" a symbol or set of symbols shown independently to each participant. Participants paid a cost for each letter used in the messages, but were rewarded if the symbols matched. Conventions matching letters to symbols or symbol features quickly emerged, resembling grammars when the symbols were complex. In both of these studies, the need to coordinate efficiently led to linguistic or pseudo-linguistic conventions.

A more popular approach to study the emergence of conventions has been to create mathematical or computational models that embody a theory, to demonstrate how a specific mechanism could lead to the emergence of social conventions. For instance, in the model described in (Boyd and Richerson 1995), agents are exploring an environment with costs and rewards, and can learn from other agents through observation. In this model, conventions evolved because of imitation and proved to be beneficial to the group. A different agent-based model was proposed in (Walker and Wooldridge 1995) to delineate mechanisms that could lead to the emergence of social conventions. Taking a game-theoretic approach, the work in (Shoham and Tennenholtz 1997; Goyal and Janssen 1997) outline features that most efficiently lead to conventions in coordination games. Noting that conventions typically arise in the context of social networks, one model looked at coordination problems in social networks (Delgado 2002), and found that conventions emerge nearly as rapidly in scale-free networks as they do in fully-connected ones, and that social conventions arise more efficiently in complex networks than in regular graphs. Another recent simulation model found that having different roles in networks can also facilitate the emergence of social conventions (Savarimuthu, Cranefield, and Purvis 2007).

The empirical study of naturally occurring norms is uncommon, but was encouraged by the development of online communities. The existence of social conventions has been observed in collaborative virtual environments such as MUDs (Multi-user Dungeons) and MOOs (MUD, object-oriented), although the primary purpose of these studies was to demonstrate that the virtual social environments were valid for social science research, something that is largely accepted today, and did not focus on how the conventions came to be (Becker and Mark 1998; 1999). For the most part, it seemed the online conventions were imported from existing social conventions such as greetings and the so-

cially appropriate distance for interpersonal interactions. This latter convention, also known as proxemics, is carried into the more modern virtual environment, *Second Life*. Two sets of researchers (Friedman, Steed, and Slater 2007; Yee et al. 2007) observed or experimentally manipulated the virtual distance between avatars in *Second Life* and found people moved their avatar to maintain an appropriate social distance. The work in (Pankoke-Babatz and Jeffrey 2002) aggregated the documents that explicitly define norms for online behavior—what can be viewed as the outcome of the emergent process of norm formation.

Methodology

We used data gathered from Twitter as reported in previous work (Cha et al. 2010), which comprises the following three types of information: profiles of 52 million users, 1.9 billion directed follow links among these users, and 1.7 billion public tweets posted by the collected users. The oldest tweet in the dataset is from March 2006, when the Twitter service was publicly launched. The follow link information is based on a snapshot taken at the time of data collection in September 2009. However, the user and tweet information is near-complete because user IDs were sequentially queried from all possible ranges (0–80 million) at the time of data collection. Therefore, this dataset provides a unique opportunity to study the birth of new collective behaviors in Twitter.

We focus on the convention for indicating a message is reposted while attributing the source, commonly known as “retweeting”. The convention typically had a syntax of “marker @username repeated-text” (boyd, Golder, and Lotan 2010), so we searched for this syntax in the tweet dataset and identified all potential variations. Among them, we study the four most frequently used variations (“RT”, “via”, “Retweet”, and “Retweeting”) and three lesser used ones (“HT”, “R/T”, and the recycle icon, shown below).



These seven retweet variations each gained popularity and spread quickly. Some Twitter users only used one variation, while others tried out multiple variations. In this work, we define *adoption* as the use of a given marker (e.g., RT) at least once. The longitudinal dataset documents the times when individuals first adopt any particular variation and how the seven different variations compete for popularity over time. The competition is natural because there is almost always a need to establish regularized patterns of behavior in social interactions in order to set common ground and to decrease conflicts (Shoham and Tennenholtz 1997). In the following sections we examine the birth and the spreading process in detail.

Origin of Conventions

Every variation that arises from the interaction of individuals is, at some point, *invented*. It is possible for the same variation to be invented multiple times independently—this could be called “convergent evolution.” However, it is impossible to know for certain whether the subsequent inventions were

truly independent or if the information about their use was transmitted outside of Twitter. Therefore, in this section we focus on the very first invention of each of the seven variations for reposting on Twitter: how they were first used and what was the context for their invention.

The first variation ever used to indicate a tweet came from another user was *via*, followed by the original poster “@kosso”, as shown in Table 1. This variation is sensible, as it is immediately understandable to most English speakers. The very first use was in March of 2007, only 12 months from the launch of Twitter, and only 4 months from the first “@username” reference appeared in Twitter. This use, and the many subsequent uses of this and other variations, establishes that there is a need on Twitter to indicate a message is passed on from another source and to attribute the message to the source.

The second variation that we observed in the dataset was HT, which is a shortening of the words “heard through”. Here we note some features that are common in Twitter: the shortening of a word or phrase because of the constraint on the number of characters allowed in a single tweet, and the inclusion of a URL. Many times the act of reposting is explicitly for the purpose of sharing information that one believes will be of interest to one’s followers. An alternative interpretation of HT is “hat tip”, and it is possible that this variation, like the greetings in the community studied in (Becker and Mark 1998), was imported from a convention established on blogs, where it was considered polite to give credit to a source by “tipping one’s hat” to the source using the abbreviation HT. This interpretation suggests that the actual content of the tweet did not have to be identical, and instead HT is more about attributing content to its source.

In contrast, the first use of the Retweet variation does not attribute the source—it is merely alerting the user’s followers that the message was passed on. Thus, the initial usage was slightly different than the usage of *via* or HT. With those, it was not necessarily the case that the message was reproduced exactly, whereas with Retweet, there was no attribution to the source. Subsequently Retweet came to also be accompanied by the attribution to the source. Moreover, this is the first variation that was invented and was community-specific, as “Retweet” only makes sense if you understand what a “tweet” is. This variation laid the foundation for what became the most widely adopted variation, and is the term Twitter ultimately chose to describe the act of reposting a message.

The longer variant of Retweet, Retweeting, first appears in a discussion about how to appropriately repost a message *and* attribute the source. This tweet was from @twhirl, a popular application for managing Twitter accounts (Twitter client): “what would you like for including sender in retweet? ‘Retweet from @sender:’, ‘Retweeting @sender:’, ‘Retweet @sender:’, or any other idea?” This variation of Retweet not only marks the evolution of the usage of the term Retweet, it also indicates a transition in the community. At this point, there was an awareness that the need existed, and discussion began about what the correct variation should be.

The evolution of the Retweet variation continued less than

Variation	Username	Date	Text
via	@tagami	2007-03-16	@JasonCalacanis (via @kosso) - new Nokia N-Series phones will do Flash, Video and YouTube
HT	@TravisSeitler	2007-10-22	The Age Project: how old do I look? http://tweetl.com/21b (HT @technosailor)
Retweet	@kevinks	2007-11-01	Retweet: @AHealthyLaugh is in the Boston Globe today, for a Stand up show she's doing tonight. Add the funny lady on Tweeter!
Retweeting	@musicdt	2008-01-05	Retweeting @Bwana: Is anyone streaming live from CES? #ces
RT	@TDavid	2008-01-25	RT @BreakingNewsOn: "LV Fire Department: No major injuries and the fire on the Monte Carlo west wing contained east wing nearly contained."
R/T	@samflemming	2008-06-20	r/t: @danwei Live online chat with Chinese President Hu Jintao. http://tinyurl.com/5qqecp . He claims he uses net to know netizen concerns
recycle icon	@claynewton	2008-09-16	[recycle icon] @ev of @biz re: twitterkeys [star] http://twurl.nl/fc6trd

Table 1: *The very first tweet that used each variation, its date, and the tweet content*

a month after the post discussing the appropriate variation to use, when a user named @TDavid shortened the word Retweet to the simple RT, in the tweet shown in Table 1. The user @TDavid had been exposed to other variations, including via, Retweet, and Retweeting, and had used two of them (via and Retweet) before. This particular retweet has exactly 140 characters, the length limit set by Twitter. This strongly suggests that the invention of RT was a result of a previous variation being adapted to the constraints of the social environment (i.e., the 140 character limit).

In most of these cases, the creation of the variation seems to have arisen naturally. One variation, the recycle symbol, seems to have been created for the purpose of improving the existing variations, and was advocated for explicitly in the discussion of which convection should be used on Twitter. On Twitter, the discussion of which variation should be used is necessarily local; that is, one can only have that discussion on Twitter with one's followers. It is perhaps because the variation was advocated by one of the original founders of Twitter, and subsequently reposted by the other founder, that the recycle symbol was adopted at all. However, most likely because it was not easy to replicate, this was one of the least used variations we observed (see Table 7).

It is worth noting that the first user of Retweet, RT, R/T, and recycle icon had been exposed to the usage of other variations at the time of their first usage. And, as it is shown in Table 2, first users of Retweet, RT, and recycle icon used another variation before starting the new variation.

One might predict that the inventors of the variations would be more active and central members of the community, both because they would be the members who are most committed to establishing the group's identity (Hackman

Variation	Exposed to	Prev. used
via	—	—
HT	—	—
Retweet	via	via
Retweeting	—	—
RT	via, Retweet, Retweeting	via, Retweet
R/T	via, Retweet, Retweeting, RT	—
recycle icon	via, Retweet, Retweeting, RT, R/T, HT	via

Table 2: *The variations that the first adopter of the variation had been exposed to or used previously*

1976), and because the most active members would be most likely to encounter the coordination problem that leads to a variation. Indeed, the data support this hypothesis. The inventors of these variations posted more tweets on average, had higher in-degree and out-degree, and were more likely to have been exposed to another variation than the average contemporaneous user.

Early Adopters

Having looked at the inventors, we now focus on a larger set of *early adopters* of each variation. We investigate their characteristics and connectivity to understand the early stages of the emergence of variations. Unless specified otherwise, we define the first 1000 adopters of each variation as the early adopters. The results in this section are similar if restricted to the first 100 or 500 adopters.

Early adopters are core users

To understand early adopters better, we investigated their network characteristics and profile information. In general, early adopters are much more popular and active than the remaining adopters and typical non-adopter users. Figure 1 shows the in-degree (followers) distribution of early adopters, the remaining adopters, and all Twitter users. Early adopters have two orders of magnitude more followers than normal users and an order of magnitude more than the later adopters. Additionally, when we computed the PageRank values, most of the early adopters (80%) ranked in top 1% PageRank, suggesting that they are not only popular Twitter users but are also central topologically.

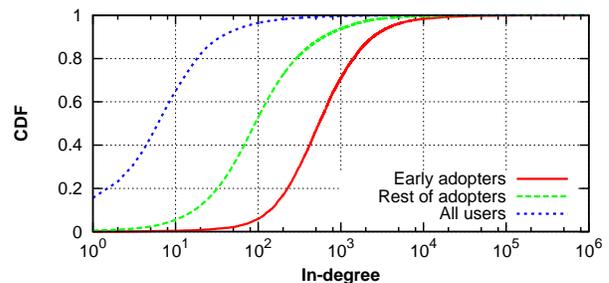


Figure 1: *In-degree of early adopters compared to rest of adopters of the retweeting convention and all Twitter users*

	Has Bio	Has URL	Profile Pic	Changed profile theme	Has Location	Has Lists
Early adopters	94%	85%	99%	91%	95%	57%
Random sample	25%	14%	50%	40%	36%	4%

Table 3: Characteristics from profile and activity of early adopters

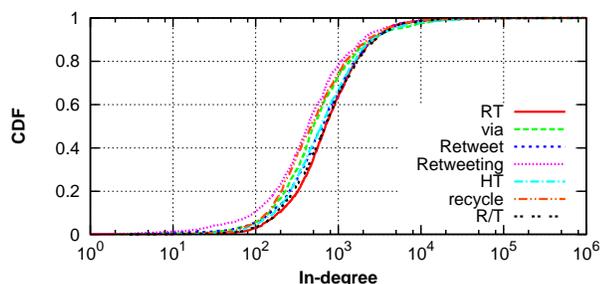


Figure 2: In-degree of early adopters across variations

In addition to the differences in the structural positions of early adopters, we also found differences in their profile information. We crawled detailed information on the profile pages of early adopters in January 2012. This profile information includes each user’s bio (a short description of a user posted by the user herself), listing information (a grouping mechanism that users can control for managing follow links), location, profile picture, and page theme.

We compared the profile information of early adopters with a random sample of 300,000 users, as shown in Table 3. The two groups differed in many ways. While 94% and 85% of early adopters provide bio and URLs (links) to their external web pages, only 25% and 14% of the random users did so. A similar trend was confirmed for profile picture, profile theme, location, and list information. These results suggest that the early adopters of the retweeting convention are active and innovative users, who explore more features provided by Twitter than the average user.

Early adopters also differed from a typical Twitter user in the content of their bio information. Users in the random sample, which represent the general Twitter population, describe themselves using words such as: *love*, *life*, *live*, and *music*. In contrast, early adopters introduce themselves with words such as *media*, *developer*, *geek*, *web*, and *entrepreneur*. This finding is in tune with the recent work that looked at the adoption pattern of the Twitter service itself, where cities with the most early adopters of Twitter tended to be those with young and tech-savvy populations (Toole, Cha, and Gonzalez 2012).

Although the early adopters are different from both the rest of adopters and from typical users, they are not particularly different from variation to variation. Figure 2 shows that early adopters have a similar number of followers independent of which variation they adopted. Although these users have similar characteristics, there is little overlap between early adopters of different variations; there are only 616 who adopted two or more, and 89 who adopted three or more variations.

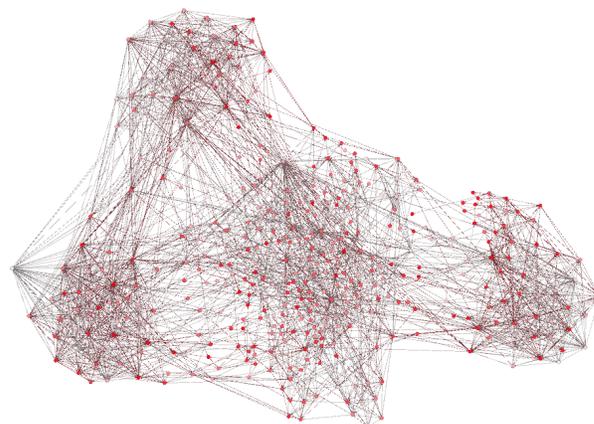


Figure 3: Diffusion network of first 500 adopters of RT

Diffusion network of early adopters

There are three possible scenarios in which a user adopts a variation: either the user was influenced by someone she was following through Twitter, the user was influenced by some external force, or the user independently invented the variation. We refer to adopters in the former case as *internal adopters*, and we refer to adopters in the latter two cases as *external adopters*.

Given the temporal ordering of users’ adoptions and the direction of following links between them, we can construct a network of early adopters representing how the variation spread in its early stages. In this *diffusion network* there is a link from user A to user B if and only if (1) user B follows user A and (2) user B adopts the variation after user A. The diffusion networks represent those instances where the adoption of the variation was most likely due to exposure through Twitter. Internal adopters naturally appear downstream in the diffusion network, whereas external adopters may appear at a root or as singletons.

Figure 3 shows an example diffusion network for the first 500 adopters of the RT variation. For simplicity, we only show the largest connected component and do not show sin-

Variation	first 100	first 500	first 1000
via	36%	64%	69.7%
HT	36%	63.4%	74%
Retweet	44%	66.4%	77.2%
Retweeting	35%	56.4%	66.8%
RT	65%	78.6%	86.1%
R/T	60%	78.4%	83.3%
recycle icon	73%	81.2%	84.8%

Table 4: Percentage of the early adopters who were internal

Variation	% in largest CC	Ave. # of links	Most critical	Ave. depth of internal adopters	Ave. clustering coefficient in LCC
via	83.2%	3.58	2%	1.48	0.233
HT	86.9%	4.22	1.8%	1.56	0.253
Retweet	79.2%	2.91	1.1%	1.44	0.241
Retweeting	91.7%	5.00	2.4%	1.46	0.225
RT	95.3%	6.37	0.5%	1.61	0.293
R/T	92.9%	5.00	1.6%	1.53	0.319
recycle icon	91.5%	4.07	4.9%	1.65	0.320

Table 5: Properties of diffusion network of early adopters

gletons or smaller components of the network. Except for 53 users, 47 of whom were singletons, the 447 remaining users formed a single large connected component, indicating that RT diffused in the network of early adopters. While we do not explore community structure in this paper, the figure indicates the existence of four communities of early adopters. It is likely that as users connecting multiple communities adopt a variation, it spread from one community to another, thereby reaching new audience.

Table 4 shows the percentage of internal adopters for the first 100, 500, and 1000 adopters of each variation. The percentage of internal adopters is low for early adopters, and as variations got more popular, more users adopted the variation after being exposed to it through Twitter. The percentage of internally influenced adopters always increases as the set of adopters gets larger, and thus the earliest variations had a much lower percentage of internal adopters. One possible explanation for this finding is that the initial variations were more natural and therefore more likely to be independently invented. Alternatively, it could be that there was much more off-site discussion in the early days of Twitter. Unfortunately, we do not have a way to disambiguate these hypotheses.

Next we take a closer look at the structure of the diffusion network amongst early adopters. Table 5 shows different structural properties of the network of each variation. We observe that most of the nodes are part of a single large connected component and the average clustering coefficient of each network is high, which means that the variations spread through multiple paths between well-connected users. Moreover, the average number of links per node in the largest connected component (LCC) is quite high. This suggests that the diffusion networks were dense and most adopters had multiple parents, as illustrated in Figure 4.

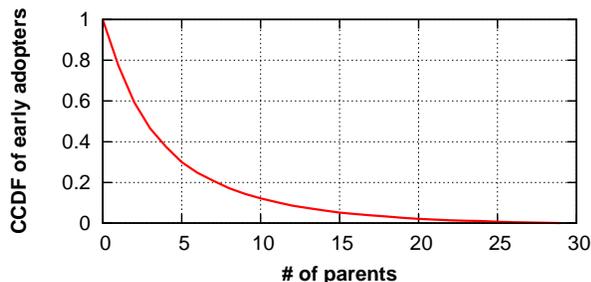


Figure 4: CCDF of number of parents of early adopters

The fact that adopters needed multiple exposures before they adopt a variation suggests that diffusion of retweeting variations is a “complex contagion” (Centola and Macy 2007).

For each internal adopter there can be one or more chains of influence, each of which leads back to one or more external adopters. We consider the length of the shortest chain to be the depth of a node. The average depth of all internal adopters in a diffusion network indicates how many adopters a variation typically went through. Table 6 shows the distribution of depths of adopters, and Table 5 shows the average depths of internal adopters for each variation. The average depth for all variations is less than two, corroborating the result that early adopters have a dense diffusion network and are well-connected to each other. This high density is also evident in Figure 3, which shows the diffusion network of early adopters of RT.

Finally, we define the *criticality* of a node as the percentage of nodes that adopted a variation after they were exposed to the variation exclusively from that node or its descendants in the diffusion graph. In other words, if we remove that specific node from the network, how many other nodes would not have been exposed to the variation? Table 5 shows the node with the highest criticality for each variation. The criticality values are in general very low; this could be explained by the fact that most of the nodes have multiple parents, so very few nodes were only exposed to the variation because of one user.

In summary, we investigated early adopters in this section. We showed that early adopters are highly active and popular core users. They are tightly connected to each other and most of them have been influenced by multiple other adopters, and as a result there are no critical early adopters that any variation relied on to become popular.

Spread of Conventions

The longitudinal tweet dataset allows us to track how each one of the variations gained popularity from their first use to

Variation	0	1	2	3	4
via	303	401	262	31	3
HT	260	383	308	39	10
Retweet	228	461	288	21	2
Retweeting	332	397	236	32	3
RT	139	394	416	46	5
R/T	167	452	328	46	7
recycle icon	152	385	376	84	3

Table 6: Number of early adopters with specific depth

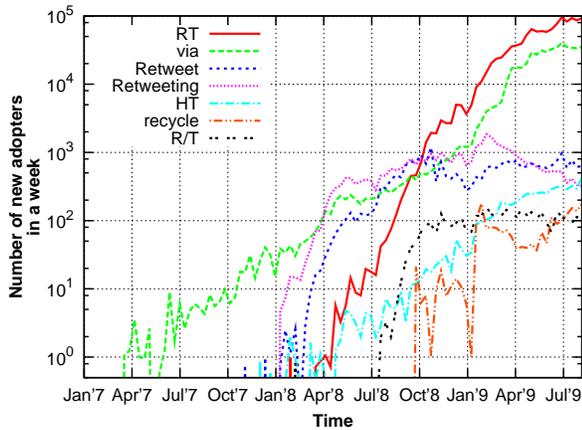


Figure 5: *New adopters of variations over time*

the last day. Figure 5 shows the time series of week-to-week user gains over the first 3.5 years of Twitter’s existence. It is clear that the different variations experienced very different patterns of growth. By the end of mid-2009, only two variations, RT and via, had achieved widespread usage. The recycle icon, HT, and R/T continued to add new users, but their popularity nearly stabilized. Retweet and Retweeting began losing popularity, as the rate of new adopters declined, potentially because of their long length, which is costly given the 140 character limit.

Table 7 shows the final number of adopters and the number of times each variation was used in our dataset. RT and via reached a total of 1.8 million and 750,000 adopters respectively, while the others only reached on the order of thousands or tens of thousands of adopters. In total, over 2 million users adopted at least one of these variations in Twitter and an impressive 59 million or an estimated 3.5% of all tweets (59 million out of 1.7 billion tweets) contain a retweeting variation.

Interestingly, the final reach of the variations does not seem to be strongly related to either the amount of time the variation had to grow or the rate at which it grew. As can be seen in Figure 5, via started the earliest and had slow growth (relative to the other variations), but ended with the second-highest number of adopters; Retweet and Retweeting grew as fast or faster than RT and started earlier, but never approached the reach of RT. It is an interesting question whether there are features of the variations, the inventors, or the early dynamics (or some combination thereof)

Variation	# of adopters	# of retweets
RT	1,836,852	53,221,529
via	751,547	5,367,304
Retweeting	50,400	296,608
Retweet	36,601	110,616
HT	8,346	22,657
R/T	5,300	28,658
recycle icon	3,305	18,255
Total	2,059,350	59,065,627

Table 7: *Number of adopters and retweets per variation*

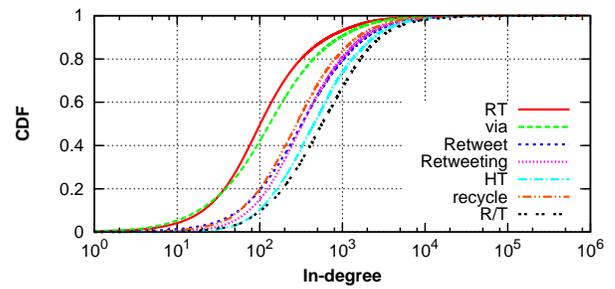


Figure 6: *In-degree of all adopters*

that can be used to predict which variations would come to dominate.

One notable feature is the distribution of in-degree of all adopters, as the fairly large gap between popular variations and the others suggests it is related to their popularity. Figure 6 shows the distribution of the in-degree of all eventual adopters of each variation. We see that RT and via, the top two variations, were adopted by less active and less popular users than the other variations.

This finding may suggest that when a variation becomes extremely popular, its adopters start to reach beyond the circle of highly-connected core users and reach more peripheral users—a finding that is in tune with the famous “diffusion of innovations” theory (Rogers 1983). This pattern of diffusion could similarly be explained by classical “two-step flow” theory, which has been shown to hold in Twitter (Wu et al. 2011). In contrast, adopters of the less popular variations may never break out of the core group of Twitter users who pay attention to the new trends and technologies of the service. Whether this shift is a leading or lagging variable to the increase in popularity is a question for future work.

As Twitter grew and the variations spread, the probability that a user would be exposed to a variation before adopting it increased as well, continuing the trend in Table 4. All seven variations had a high percentage of internally influenced users (81.97%–97.93%). This could be partially explained by the fact that the act of reposting is increasing itself, and therefore an increase in the use of any variation and likelihood that someone would be exposed to it. It also implies that as time went on, more and more users were learning about the variations through the Twitter network rather than independently inventing variations or learning about them from outside channels.

Conclusion and Future Work

In this work, we described the context in which a social convention emerged, capturing the natural evolution of the convention at a level of detail and scale essentially impossible until recently. Specifically, we observed the very first acts of reposting a message and attributing it to its source on the microblogging site Twitter. At first, variations of the convention were borrowed from natural language (“via”) or other online communities (“HT”, from blogs). Eventually, more community-specific variations (Retweet and Retweeting) were invented, followed shortly by discussion of which

variations were best. However, because of an environmental constraint—the 140 character limit on tweets—a more concise community-specific variation emerged (RT). Interestingly, it was this variant that came to be the most popular, despite its late introduction and despite subsequent attempts to explicitly introduce “better” variations like the recycle symbol.

The inventors of these variations were not the typical user. They posted more tweets, had higher network degree, and were more likely to describe themselves with words like “geek” and “founder”; in other words, they were the core members of the Twitter community. The early adopters were also more active and innovative. The variations spread through densely connected networks, bouncing from person to person in a way that meant most adopters of the variation were fewer than two hops from someone who had never been exposed to the variation on Twitter when they first used it. This could be a general finding, that social conventions are more likely to arise in the active and densely connected parts of a community.

As Twitter grew, the variations spread at very different rates and with very different outcomes. Two of the variations that eventually became extremely popular were used more than a million times. Their adopters extended beyond the cluster of highly-connected core users into the periphery of the Twitter network. The remaining variations were adopted by orders of magnitude fewer users. It is not clear why the variations that were widely adopted were so successful relative to the other variations, and this is a crucial question for future research.

This work was mostly descriptive, and there are many important predictive questions that remain unanswered. For instance, given characteristics of the variations, the users, and their network, can one predict which variation a particular user will choose to adopt? Do different variations have large-scale collisions as they diffuse through the network? Do variations become associated with subgroups, leading to schisms in groups, or can they coexist? By answering these questions, we would have a much better understanding of how social conventions emerge in human society.

Acknowledgements Meeyoung Cha was supported by Basic Science Research Program through the National Research Foundation, Korea (2011-0030936).

References

- Becker, B., and Mark, G. 1998. Social conventions in collaborative virtual environments. In *Proceedings of Collaborative Virtual Environments*, 17–19.
- Becker, B., and Mark, G. 1999. Constructing social systems through computer-mediated communication. *Virtual reality*.
- Bettenhausen, K., and Murnighan, J. K. 1985. The Emergence of Norms in Competitive Decision-Making Groups. *Administrative Science Quarterly* 30:350–372.
- Boyd, R., and Richerson, P. J. 1995. Why does culture increase human adaptability? *Ethology and Sociobiology* 16:125–143.
- boyd, d.; Golder, S.; and Lotan, G. 2010. Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter. In *Hawaii International Conference on System Sciences*.
- Centola, D., and Macy, M. 2007. Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology* 113(3):702–734.
- Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, K. 2010. Measuring User Influence in Twitter: The Million Follower Fallacy. In *International AAAI Conference on Weblogs and Social Media*.
- Delgado, J. 2002. Emergence of social conventions in complex networks. *Artificial Intelligence* 141(1-2):171–185.
- Feldman, D. C. 1984. The development and enforcement of group norms. *The Academy of Management Review* 9(1):47–53.
- Friedman, D.; Steed, A.; and Slater, M. 2007. Spatial social behavior in second life. *Intelligent virtual agents*.
- Goyal, S., and Janssen, M. C. W. 1997. Non-Exclusive Conventions and Social Coordination* 1. *Journal of Economic Theory* 77(1):34–57.
- Hackman, J. R. 1976. *Group Influences on Individuals in Organizations*, volume 3. Consulting Psychologists Press, Inc.
- Levine, J., and Moreland, R. L. 1990. Progress in Small Group Research. *Annual Reviews in Psychology* 41:585–634.
- Opp, K.-D. 2001. How do norms emerge? An outline of a theory. *Mind & Society* 3(2):101–128.
- Pankoke-Babatz, U., and Jeffrey, P. 2002. Documented Norms and Conventions on the Internet. In Luczak, H., ed., *International Journal of Human-Computer Interaction*. Lawrence Erlbaum Association. 219–235.
- Riley, A., and Burke, P. J. 1995. Identities and Self-Verification in the Small Group. *Social Psychology Quarterly* 58(2):61–73.
- Rogers, E. M. 1983. *Diffusion of Innovations*. Free Press.
- Savarimuthu, B.; Cranefield, S.; and Purvis, M. 2007. Role Model Based Mechanism for Norm Emergence in Artificial Agent Societies. In *AAMAS '07 Workshop on Coordination, Organization, Institutions and Norms in Agent Systems*, 1–12.
- Selten, R., and Warglien, M. 2007. The emergence of simple languages in an experimental coordination game. In *Proceedings of the National Academy of Sciences*, 7361–7366.
- Shoham, Y., and Tennenholtz, M. 1997. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence* 94(1-2):139–166.
- Toole, J. L.; Cha, M.; and Gonzalez, M. C. 2012. Modeling the adoption of innovations in the presence of geographic and media influences. In *Public Library of Science One*.
- Walker, A., and Wooldridge, M. 1995. Understanding the emergence of conventions in multi-agent systems. *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)* 384–389.
- Wilkes-Gibbs, D., and Clark, H. 1992. Coordinating beliefs in conversation. *Journal of Memory and Language* 31(2):183–194.
- Wu, S.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Who says what to whom on twitter. In *Proceedings of the 20th international conference on World wide web, WWW '11*, 705–714. New York, NY, USA: ACM.
- Yee, N.; Bailenson, J. N.; Urbanek, M.; Chang, F.; and Merget, D. 2007. The Unbearable Likeness of Being Digital: The Persistence of Nonverbal Social Norms in Online Virtual Environments. *CyberPsychology & Behavior* 10(1):115–121.