

Wayfinding in Social Networks

David Liben-Nowell

Abstract With the recent explosion of popularity of commercial social-networking sites like Facebook and MySpace, the size of social networks that can be studied scientifically has passed from the scale traditionally studied by sociologists and anthropologists to the scale of networks more typically studied by computer scientists. In this chapter, I will highlight a recent line of computational research into the modeling and analysis of the *small-world phenomenon*—the observation that typical pairs of people in a social network are connected by very short chains of intermediate friends—and the ability of members of a large social network to collectively find efficient routes to reach individuals in the network. I will survey several recent mathematical models of social networks that account for these phenomena, with an emphasis both on provable properties of these social-network models and on the empirical validation of the models against real large-scale social-network data.

1 Introduction

An intrepid graduate student at a university in Cambridge, MA leaves his department on a quest for cannoli in the North End, a largely Italian neighborhood of Boston. This mission requires him to solve a particular navigation problem in the city: from his background knowledge (his own “mental map” of Boston, including what he knows about the proximity of particular landmarks to his destination pastry shop, how he conceives of the margins of the North End district, where he thinks his office is) and what he gathers at his initial location (street signs, perhaps the well-trodden routes scuffed into the grass outside his office that lead to the nearest subway station, perhaps the smell of pesto), he must begin to construct a path towards his destination. *Wayfinding*, a word coined by the urban planner Kevin Lynch in the

David Liben-Nowell

Department of Computer Science; Carleton College; Northfield, MN 55057, USA

e-mail: dlibenno@carleton.edu

1960's [33], refers to the processes by which a person situates him or herself in an urban environment and navigates from point to point in that setting. (Incidentally, wayfinding is particularly difficult in Boston, with its often irregular angles and patterns: one may become confused about directions—the North End is almost due east of the West End—or encounter impermeable boundaries like rivers or highways even while heading in the right cardinal direction.)

This chapter will be concerned with wayfinding in networks—specifically in *social networks*, structures formed by the set of social relationships that connect a set of people. The issues of navigability of a social network are analogous to the issues of navigability of a city: can a source person s reliably identify a step along an efficient path “towards” a destination person t from her background knowledge (a “mental map” of the social network) plus whatever information she gathers from her immediate proximity? How does one find routes in a social network, and how easy is it to compute those routes? To borrow another term from Kevin Lynch, is the “legibility” of a real social network more like Manhattan—a clear grid, easy navigation—or like Boston?

In later sections, I will describe a more formal version of these notions and questions, and describe some formal models that can help to explain some of these real-world phenomena. (The interested reader may also find the recent surveys by Kleinberg [25] and Fraigniaud [17] valuable.) I will begin with a grid-based model of social networks and social-network routing due to Jon Kleinberg [23, 26], under which Kleinberg has fully characterized the parameter values for which the resulting social network supports the construction of short paths through the network in a decentralized, distributed fashion. Geography is perhaps the most natural context in which to model social networks via a regular grid, so I will then turn to empirical observations of geographic patterns of friendship in real large-scale online social networks. I will then describe some modifications to Kleinberg's model suggested by observed characteristics of these real-world networks, most notably the widely varying population density across geographic locations. Finally, I will turn to other models of social networks and social-network routing, considering both network models based on notions of similarity that are poorly modeled by a grid (e.g., occupation) and network models based simultaneously on multiple notions of person-to-person similarity.

2 The Small-World Phenomenon

Although social networks have been implicit in the interactions of humans for millennia, and social interactions among humans have been studied by social scientists for centuries, the academic study of social networks *qua* networks is more recent. Some of the early foundational contributions date from the beginning of the twentieth century, including the “web of group affiliations” of Georg Simmel [45], the “sociograms” of Jacob Moreno [40], and the “topological psychology” of Kurt Lewin [31]. In the 1950's, Cartwright and Zander [9] and Harary and Norman [20]

described an explicitly graph-theoretic framework for social networks: nodes represent individuals and edges represent relationships between pairs of individuals. We will use this graph-theoretic language throughout the chapter.

The line of social-network research that is the focus of this chapter can be traced back to an innovative experiment conceived and performed by the social psychologist Stanley Milgram in the 1960's [38]. Milgram chose 100 “starter” individuals in Omaha, NE, and sent each one of them a letter. The accompanying instructions said that the letter holder s should choose one of her friends to whom to forward the letter, with the eventual goal of reaching a target person t , a stockbroker living near Boston. (For Milgram’s purposes, a “friend” of s was anyone with whom s was on a mutual first-name basis.) Each subsequent recipient of the letter would receive the same instructions, and, presumably, the letter would successively home in on t with each step. What Milgram found was that, of the chains that reached the stockbroker, on average they took about six hops to arrive. That observation was the origin of the phrase “six degrees of separation.” Of course, what this careful phrasing glosses over is the fraction of chains—about 80%—that failed to reach the target; Judith Kleinfeld has raised an interesting and compelling set of critiques about the often overbroad conclusions drawn from these limited data [27]. Still, Milgram’s small-world results have been replicated in a variety of settings—including a recent large-scale email-based study by Dodds, Muhamad, and Watts [13]—and Milgram’s general conclusions are not in dispute.

The *small-world problem*—why is it that a random resident of Omaha should be only a few hops removed from a stockbroker living in a suburb of Boston?—was traditionally answered by citing, explicitly or implicitly, the voluminous body of mathematical literature that shows that random graphs have small diameter. But that explanation suffers in two important ways: first, social networks are poorly modeled by random graphs; and, second, in the language of Kevin Lynch, Milgram’s observation is much more about wayfinding than about diameter. The first objection, that social networks do not look very much like random graphs, was articulated by Duncan Watts and Steve Strogatz [49]. Watts and Strogatz quantified this objection in terms of the high *clustering* in social networks: many pairs of people who share a common friend are also friends themselves. The second objection, that the Milgram experiment says something about the efficiency of wayfinding in a social network and not just something about the network’s diameter, was raised by Jon Kleinberg [23, 26]. Kleinberg observed that Milgram’s result is better understood not just as an observation about the existence of short paths from source to target, but rather as an observation about distributed algorithms: somehow people in the Milgram experiment have collectively managed to *construct* short paths from source to target.

How might people be able to accomplish this task so efficiently? In a social network—or indeed in any network—finding short paths to a target t typically hinges on making some form of measurable progress towards t . Ideally, this measure of progress in a social network would simply be graph distance: at every step, the path would move to a node with smaller graph distance to t —i.e., along a shortest path to t . But the highly decentralized nature of a social network means that only a handful of nodes (the target t himself, the neighbors of t , and perhaps a few neigh-

bors of the neighbors of t) genuinely know their graph distance to t . Instead one must use some guide other than graph distance to home in on t .

The key idea in routing in this context—frequently cited by the participants in real small-world experiments as their routing strategy [13, 21]—is to use similarity of characteristics (geographic location, hobbies, occupation, age, etc.) as a measure of progress, a proxy for the ideal but unattainable graph-distance measure of proximity. The success of this routing strategy hinges on the sociological observation of the crucial tendency towards *homophily* in human relationships: the friends of a typical person x tend to be similar to x . This similarity tends to occur with respect to race, occupation, socioeconomics, and geography, among other dimensions. See the survey of McPherson, Smith-Lovin, and Cook [36] for an excellent review of the literature on homophily. Homophily makes characteristic-based routing reasonable, and in fact it also gives one explanation for the high clustering of real social networks: if x 's friends tend to be similar to x , then they also tend to be (somewhat less) similar to each other, and therefore they also tend to know each other directly with a (somewhat) higher probability than a random pair of people.

Homophily suggests a natural *greedy algorithm* for routing in social networks. If a person s is trying to construct a path to a target t , then s should look at all of her friends $\Gamma(s)$ and, of them, select the friend in $\Gamma(s)$ who is “most like” the target t . This notion is straightforward when it comes to geography: the source s knows both where her friends live and where t lives, and thus s can just compute the geographic distance between each $u \in \Gamma(s)$ and t , choosing the u minimizing that quantity. Routing greedily with respect to occupation is somewhat murkier, though one can imagine s choosing u based on distance within an implicit hierarchy of occupations in s 's head. (Milgram's stockbroker presumably falls into something like the service industry \rightarrow financial services \rightarrow investment \rightarrow stocks.) Indeed, the greedy algorithm is well founded as long as an individual has sufficient knowledge of underlying person-to-person similarities to compare the distances between each of her friends and the target.

3 Kleinberg's Small-World Model: The Navigable Grid

Although homophily is a key motivation for greedy routing, homophily alone does not suffice to ensure that the greedy algorithm will find short paths through a social network. As a concrete example, suppose that every sociologist studying social networks knows every other such sociologist and nobody else, and every computer scientist studying social networks knows every other such computer scientist and nobody else. This network has an extremely high degree of homophily. But the network is not even connected, let alone navigable by the greedy algorithm. For the greedy algorithm to succeed, the probability of friendship between people u and v should somehow vary more smoothly as the similarity of u and v decreases. Intuitively, there is a tension between having “well-scattered” friends to reach faraway targets and having “well-localized” friends to home in on nearby targets. Without

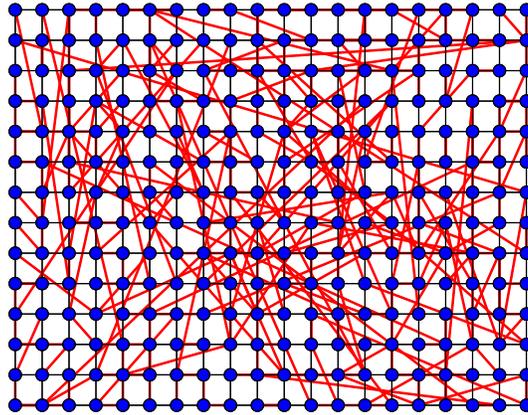


Fig. 1 Kleinberg’s small-world model [23, 26]. A population of n people is arranged on a k -dimensional grid, and each person u is connected to her immediate neighbors in each direction. Each person u is also connected to a long-range friend v , chosen with probability $\propto d(u, v)^{-\alpha}$, where $d(\cdot, \cdot)$ denotes Manhattan distance and $\alpha \geq 0$ is a parameter to the model. The example two-dimensional network here was generated with $\alpha = 2$.

the former, a large number of steps will be required to span the large gap from a source s to an especially dissimilar target t ; without the latter, similarity will be only vaguely related to graph-distance proximity, and thus the greedy algorithm will be a poor approximation to a globally aware shortest-path algorithm.

A rigorous form of this observation was made by Jon Kleinberg [23, 26], through formal analysis of this tradeoff in an elegant model of social networks. Here is Kleinberg’s model, in its simplest form. (See Section 6 for generalizations.) Consider an n -person population, and arrange these people as the points in a regular k -dimensional grid. Each person u in the network is connected to $2k$ “local neighbors,” the people who live one grid point above and below u in each of the k cardinal directions. (People on the edges of the grid will have fewer local neighbors, or we can treat the grid as a torus without substantively affecting the results.) Each person u will also be endowed with a “long-range link” to one other person v in the network. That person v will be chosen probabilistically, where $\Pr[u \rightarrow v] \propto d(u, v)^{-\alpha}$, where $d(\cdot, \cdot)$ denotes Manhattan distance in the grid, and $\alpha \geq 0$ is a parameter to the model. (Changing the model to endow each person with any constant number of long-range links does not qualitatively change the results.) See Figure 1 for an example network, with $k = \alpha = 2$. Notice that the parameter α operationalizes the tradeoff between highly localized friends and highly scattered friends: setting $\alpha = 0$ yields links to a v chosen uniformly from the network, while letting $\alpha \rightarrow \infty$ yields only links from person u to person v if $d(u, v) = 1$.

A *local-information algorithm* is one that computes a path to a target without global knowledge of the graph. When a person u chooses a next step v in the path to the target t , the person u has knowledge of the structure of the grid, including the grid locations of u herself, u ’s local neighbors, u ’s long-range contact, and the

target t . However, the remaining structure of the graph—that is, the long-range links for nodes other than u —are not available to u when she is making her routing choice. (The results are not affected by expanding the knowledge of each node u to include the list of all people previously on the path from the original source s to u , or even the list of long-range links for each of those people.)

Kleinberg was able to give a complete characterization of the navigability of these networks by local-information algorithms:

Theorem 1 (Kleinberg [23, 26]). *Consider an n -person network with people arranged in a k -dimensional grid, where each person has $2k$ local neighbors and one long-range link chosen with parameter $\alpha \geq 0$, so that $\Pr[u \rightarrow v] \propto d(u, v)^{-\alpha}$. For an arbitrary source person s and an arbitrary target person t :*

- *If $\alpha \neq k$, then there exists some constant $\varepsilon > 0$, where ε depends on α and k but is independent of n , such that the expected length of the path from s to t found by any local-information algorithm is $\Omega(n^\varepsilon)$.*
- *If $\alpha = k$, then the greedy algorithm—i.e., the algorithm that chooses the next step in the path as the contact closest to the target t under Manhattan distance in the grid—finds a path from s to t of expected length $O(\log^2 n)$.*

The proof that greedy routing finds a path of length $O(\log^2 n)$ when $\alpha = k$ proceeds by showing that the probability of halving the distance to the target at any step of the path is $\Omega(1/\log n)$. Thus, in expectation, the distance to the target is halved every $O(\log n)$ steps. The path reaches the target after the distance is halved $\log n$ times, and therefore $O(\log^2 n)$ total steps suffice to reach the target in expectation.

For our purposes, we will broadly treat paths of length polynomial in the logarithm of the population size as “short,” and paths of length polynomial in the population size as “long.” (We will use standard terminology in referring to these “short” paths as having *polylogarithmic* length—that is, length $O(\log^c n)$ for some constant exponent c , in a population of size n .) There has been significant work devoted to tightening the analysis of greedy routing in Kleinberg’s networks—for example, [7, 35]—but for now we will focus on the existence of algorithms that find paths of polylogarithmic length, without too much concern about the precise exponent of the polynomial. A network in which a local-information algorithm can find a path of polylogarithmic length is called *navigable*. Theorem 1, then, can be rephrased as follows: a k -dimensional grid-based social network with parameter α is navigable if and only if $k = \alpha$.

Note that this definition of navigability, and hence Kleinberg’s result, describes routing performance asymptotically in the population size n . Real networks, of course, are finite. Aaron Clauset and Christopher Moore [11] have shown via simulation that in finite networks, greedy routing performs well even under what Theorem 1 identifies as “non-navigable” values of α . Following [11], define α_{opt} as the value of α that produces the network under which greedy routing achieves the shortest path lengths. Clauset and Moore’s simulations show that α_{opt} is somewhat less than k in large but finite networks; furthermore, although α_{opt} approaches k as the population grows large, this convergence is relatively slow.

4 Geography and Small Worlds

Now that we have formal descriptions of mathematical models of social networks, we can turn to evaluating the correspondence between these models' predictions and real social-network data. To begin, consider physical-world geographic proximity as the underlying measure of similarity between people. Geographic distance is a natural metric on which to focus both because of its simplicity and because of its observed importance in empirical studies: participants in real-world Milgram-style experiments of social-network routing frequently cite geographic proximity as the reason that they chose a particular friend as the next step in a chain [13, 21]. (Typically, people who were early in these chains cited geography as a principal reason for their choice, and those who appeared later in the chains more frequently reported that their choice was guided by similarity of occupation.)

Although the description of Kleinberg's model in Section 3 is couched in abstract terms, there is a very natural geographic interpretation to the model's underlying grid. Geographic distance on the surface of the earth is well modeled by a 2-dimensional grid under Manhattan distance, where we imagine grid points as the intersections of evenly spaced lines of longitude and latitude. The grid is of course a simplification of real proximity in the real world, but it is plausible as a first approximation. Thus we have a mathematical model of social-network routing based on geographic proximity, and real-world evidence that people who (at least partially) successfully route through social networks do so (at least partially) based on geographic proximity. We are now in a position to test the mathematical model against real social networks.

Much of the empirical work described here will be based on data from the LiveJournal blogging community, found online at livejournal.com. LiveJournal is an appealing domain for study in part because it contains reasonably rich data about its users, even if one ignores the detailed accounts of users' personal lives frequently found in their blog posts. LiveJournal users create profiles that include demographic information such as birthday, hometown, and a list of interests/hobbies. Each user's profile also includes an explicit list of other LiveJournal users whom that user considers to be a friend.

The analysis that follows—performed in joint work with Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins [32]—was based on a crawl of LiveJournal performed in February 2004, comprising about 1.3 million user profiles. (As of this writing, there are slightly more than 17 million LiveJournal accounts.) Of those 1.3 million users, approximately 500,000 users declared a hometown that we were able to locate in a database of longitudes and latitudes in the continental United States. These data yield a large-scale social network with geographic locations, using the explicitly listed friendships to define connections among this set of 500,000 people. Figure 2 contains a visual representation of this network.

Using this 500,000-person network, we can simulate the Milgram experiment, using the purely geographic greedy routing algorithm. There is some subtlety in setting up this simulation—for example, what happens if the simulated chain reaches a person u who has no friends closer to the target than u herself?—and, because

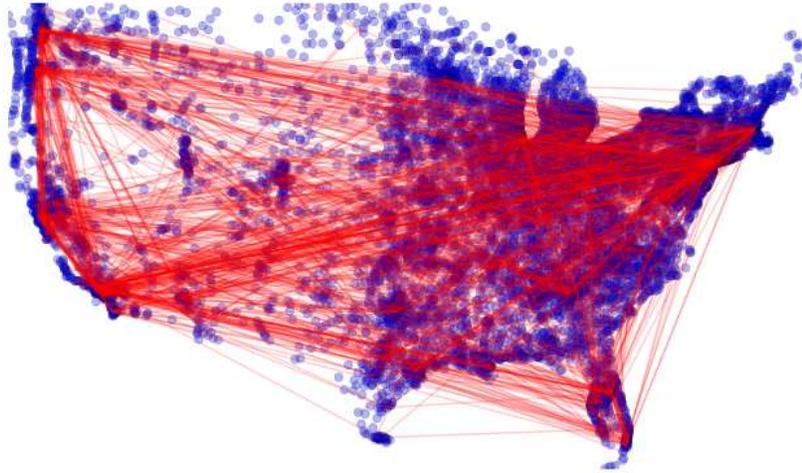


Fig. 2 The LiveJournal social network [32]. A dot is shown for each geographic location that was declared as the hometown of at least one of the $\approx 500,000$ LiveJournal users whom we were able to locate at a longitude and latitude in the continental United States. A random 0.1% of the friendships in the network are overlaid on these locations.

the resolution of geographic locations is limited to the level of towns and cities, we try only to reach the city of the target t rather than t herself. We found that, subject to these caveats, the geographic greedy algorithm was able to find short paths connecting many pairs of people in the network. (See [32] for more detail.)

With the above observations (people are arranged on a 2-dimensional geographic grid; greedy routing based on geography finds short paths through the network) and Theorem 1, we set out—in retrospect, deeply naïvely—to verify that the probability of friendship between people u and v grows asymptotically as $d(u, v)^{-2}$ in the LiveJournal network. In other words, in the language of Kleinberg’s theorem, we wanted to confirm that $\alpha = 2$. The results are shown in Figure 3, which displays the probability $P(d)$ of friendship between two people who live a given distance d apart—i.e., the fraction of pairs separated by distance d who declare a friendship in LiveJournal.

One immediate observation from the plot in Figure 3 is that the probability $P(d)$ of friendship between two people in LiveJournal separated by distance d really does decrease smoothly and markedly as d increases. This relationship already reveals a nonobvious fact about LiveJournal; there was no particular reason to think that geographic proximity would necessarily play an important role in friendships in a completely virtual community like this one. Section 7 includes some discussion of a few possible reasons why geography remains so crucial in this virtual setting, but for now it is worth noting that the “virtualization” of real-world friendships (that is, the process of creating digital records of existing physical-world friendships) seems to explain only some of the role of geography. For example, it seems hard for

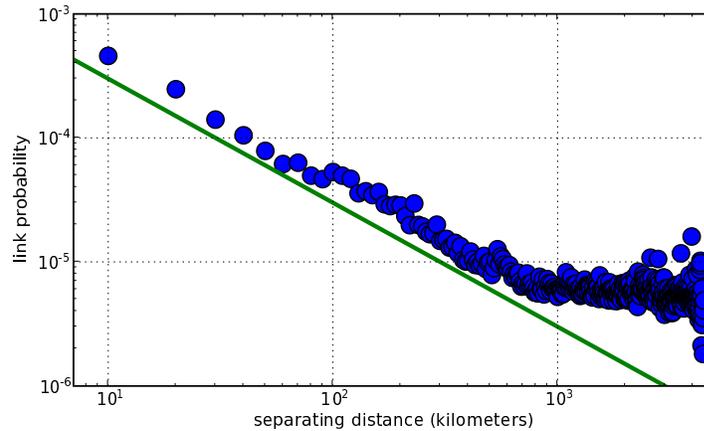


Fig. 3 The probability $P(d)$ of a friendship between two people in LiveJournal as a function of the geographic distance d between their declared hometowns [32]. Distances are rounded into 10-kilometer buckets. The solid line corresponds to $P(d) \propto 1/d$. Note that Theorem 1 requires $P(d) \propto 1/d^2$ for a network of people arranged in a regular 2-dimensional grid to be navigable.

this process to fully account for the marked difference in link probability between people separated by 300 versus 500 kilometers, a range at which regular physical-world interactions seem unlikely.

A second striking observation from the plot in Figure 3 is that the probability $P(d)$ of friendship between people separated by distance d is very poorly modeled by $P(d) \propto 1/d^2$, the relationship required by Theorem 1. This probability is better modeled as $P(d) \propto 1/d$, and in fact is even better modeled as $P(d) = \varepsilon + \Theta(1/d)$, for a constant $\varepsilon \approx 5.0 \times 10^{-6}$. Apropos the discussion in the previous paragraph, this additive constant makes some sense: the probability that people u and v are friends can be thought of as the sum of two probabilities, one that increases with their geographic proximity, and one that is independent of their geographic locations. But, regardless of the presence or absence of the additive ε , the plot in Figure 3 does not match—or even come close to matching—the navigable exponent required by Kleinberg’s theorem.

Similar results have also been observed in another social-networking context. In a study of the email-based links among about 450 members of Hewlett-Packard Research Labs [1], Lada Adamic and Eytan Adar found that the link probability $P(d)$ between two HP Labs researchers was also closely matched by $P(d) \propto 1/d$, where d measured the Manhattan distance between the cubicle locations of the employees. In this setting, too, geographic greedy routing found short paths to most targets—though not as short as those found by routing greedily according to proximity in the organizational hierarchy of the corporation (see Section 7)—again yielding a greedily navigable network that does not match Theorem 1.

5 Variable Population Density and Rank-Based Friendship

The observations from the previous section lead to a seeming puzzle: a navigable two-dimensional grid, which must have link probabilities decaying as $1/d^2$ to be navigable according to Theorem 1, has link probabilities decaying as $1/d$. But another look at Figure 2 reveals an explanation—and reveals the naïvete of looking for $P(d) \propto 1/d^2$ in the LiveJournal network. Although a 2-dimensional grid is a reasonable model of geographic location, a *uniformly distributed* population on a 2-dimensional grid is a very poor model of the geographic distribution of the LiveJournal population. Population density varies widely across the United States—from over 10,000 people/km² in parts of Manhattan to approximately 1 person/km² in places like Lake of the Woods County, in the far northern reaches of Minnesota. Two Manhattanites who live 500 meters apart have probably never even met; two Lake of the Woods residents who live 500 meters apart are probably next-door neighbors, and thus they are almost certain to know each other. This wide spectrum suggests that distance cannot be the whole story in any reasonable geographic model of social networks: although $\Pr[u \rightarrow v]$ should be a decreasing function of the geographic distance between u and v , intuitively the rate of decrease in that probability should reflect something about the population in the vicinity of these people.

One way to account for variable population density is *rank-based friendship* [6, 28, 32], which models social networks as follows. The grid-based model described here is the simplest version of rank-based friendship; as with Kleinberg’s distance-based model, generalizations that do not rely on the grid have been studied. (See Section 6.) We continue to measure person-to-person distances using Manhattan distance in a k -dimensional grid, but we will now allow an arbitrary positive number of people to live at each grid point. Each person still has $2k$ local neighbors, one in each of the two directions in each of the k dimensions, and one long-range link, chosen as follows. Define the *rank* of a person v with respect to u as the number of people who live at least as close to u as v does, breaking ties in some consistent way. (In other words, person u sorts the population in descending order of proximity to u ; the rank of v is her index in this sorted list.) Now each person u chooses her long-range link according to rank, so that $\Pr[u \rightarrow v]$ is inversely proportional to the rank of v with respect to u . See Figure 4 for an example rank-based network.

Rank-based friendship generalizes the navigable $\alpha = k$ setting in the distance-based Theorem 1: in a k -dimensional grid with constant population at each point, the rank of v with respect to u is $\Theta(d(u, v)^k)$. But even under non-uniform population densities, social networks generated according to rank-based friendship are navigable by the greedy algorithm:

Theorem 2 (Liben-Nowell, Novak, Kumar, Raghavan, Tomkins [28, 32]). *Consider an n -person network where people are arranged in a k -dimensional grid so that at least one person lives at every grid point x . Suppose each person has $2k$ local neighbors and one long-range link chosen via rank-based friendship. Fix any source person s and choose a target person t uniformly at random from the popula-*

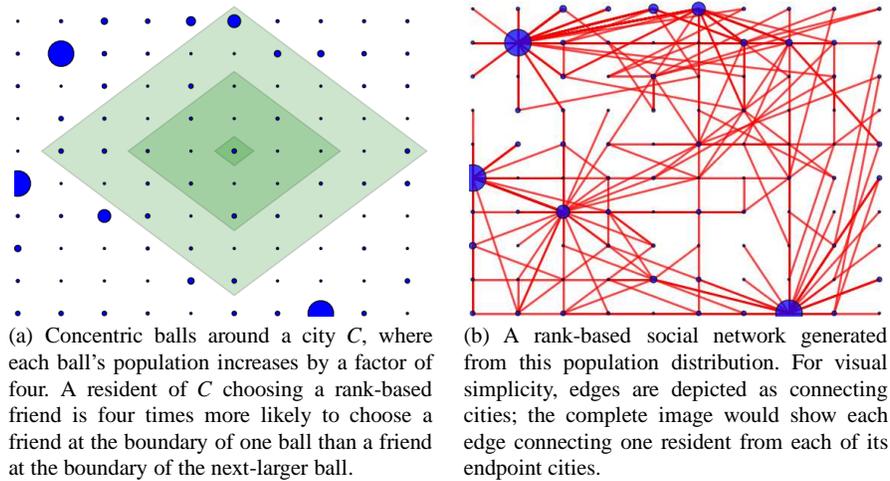


Fig. 4 Two images of a sample rank-based social network with variable population density. Each blue circle represents a city with a population whose size is proportional to the circle's radius. Distances between cities, and hence between people, are computed using Manhattan distance. A rank-based friendship for each person u is formed probabilistically, where $\Pr[u \rightarrow v]$ is inversely proportional to the number of people who live closer to u than v is, breaking ties consistently. The local neighbors—for each person u , one friend in the neighboring city in each cardinal direction—are not shown.

tion. Then under greedy routing the expected length of the path from s to the point x_t in which t lives is $O(\log^3 n)$.

A few notes about this theorem are in order. First, relative to Theorem 1, rank-based friendship has lost a logarithmic factor in the length of the path found by greedy routing. Recently, in joint work with David Barbella, George Kachergis, Anna Sallstrom, and Ben Sowell, we were able to show that a “cautious” variant on greedy routing finds a path of expected length $O(\log^2 n)$ in rank-based networks [6], but the analogous tightening for greedy routing itself remains open.

Second, Theorem 2 makes a claim about the expected length of the path found by the greedy algorithm for a randomly chosen target t , where the expectation is taken over both the random construction of the network *and* the random choice of the target. In contrast, Theorem 1 makes a claim about the expected length of the path found by the greedy algorithm for *any* target, where the expectation is taken only over the random construction of the network. Intuitively, some targets in a rank-based network may be very difficult to reach: if a person t lives in a region of the network that has a comparatively very sparse population, then there will be very few long-range links to people near t . Thus making progress towards an isolated target may be very difficult. However, the difficulty of reaching an isolated target like t is offset by the low probability of choosing such a target; almost by definition, there

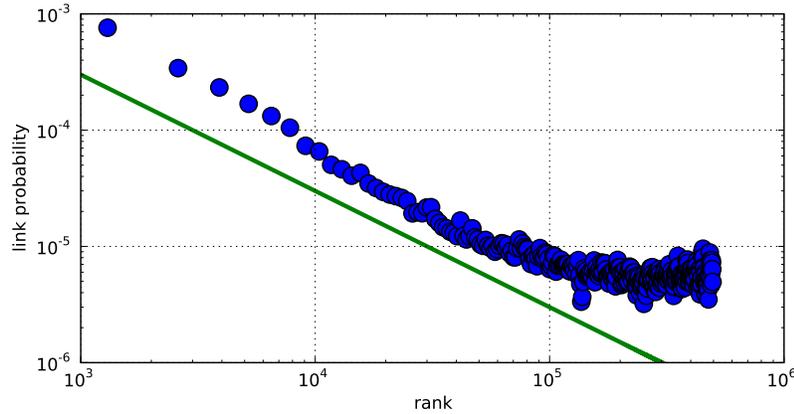


Fig. 5 The probability $P(r)$ of a friendship between two people u and v in LiveJournal as a function of the rank of v with respect to u [32]. Ranks are rounded into buckets of size 1300, which is the LiveJournal population of the city for a randomly chosen person in the network, and thus 1300 is in a sense the “rank resolution” of the dataset. (The unaveraged data are noisier, but follow the same trend.) The solid line corresponds to $P(r) \propto 1/r$. Note that Theorem 2 requires $P(r) \propto 1/r$ for a rank-based network to be navigable.

cannot be very many people who live in regions of the network that have unusually low density. The proof of Theorem 2 formalizes this intuition [28, 32].

This technical difference in the statements of Theorems 1 and 2 in fact echoes points raised by Judith Kleinfeld in her critique of the overly expansive interpretation of Milgram’s experimental results [27]. Milgram’s stockbroker was a socially prominent target, and other Milgram-style studies performed with less prominent targets—the wife of a Harvard Divinity School student, in one study performed by Milgram himself—yielded results much less suggestive of a small world.

It is also worth noting that, although the “isolated target” intuition suggests why existing proof techniques are unlikely to yield a “for all targets” version of Theorem 2, there are no known population distributions in which greedy routing fails to find a short path to any particular target in a rank-based network. It is an interesting open question to resolve whether there are population distributions and source–target pairs for which greedy routing fails to find a path of short expected length in rank-based networks (where, as in Theorem 1, the expectation is taken only over the construction of the network).

These two ways in which Theorem 2 is weaker than Theorem 1 are of course counterbalanced by the fact that Theorem 2 can handle varying population densities. But the real possible benefit is the potential for a better fit with real data. Figure 5 is the rank analogue of Figure 3: for any rank r , the fraction of LiveJournal users who link to their r th-most geographically proximate person is displayed. (Some averaging has been done in Figure 5: because a random person in the LiveJournal network

lives in a city with about 1300 residents, the data do not permit us to adequately distinguish among ranks that differ by less than this number.)

As it was with distance, the link probability $P(r)$ between two people is a smoothly decreasing function of the rank r of one with respect to the other. And just as before, link probability levels off to about $\varepsilon = 5.0 \times 10^{-6}$ as the rank gets large, so $P(r)$ is well modeled by $P(r) = \Theta(1/r) + \varepsilon$. But unlike the distance-based model of Figure 3 and Theorem 1, the fit between Figure 5 and Theorem 2 is notable: people in the LiveJournal network really have formed links with a geographic distribution that is a remarkably close match to rank-based friendship.

6 Going off the Grid

Until now, our discussion has concentrated on models of proximity that are based on Manhattan distance in an underlying grid. We have argued that these grid-based models are reasonable for geographic proximity. Even in the geographic context, though, they are imperfect: the 2-dimensional grid fails to account for real-world geographic features like the third dimension of a high-rise apartment complex or the imperfect mapping between geographic distance and transit-time distance between two points. But in a real Milgram-style routing experiment, there are numerous other measures of proximity that one might use as a guide in selecting the next step towards a target: occupation, age, hobbies, and alma mater, for example. The grid is a very poor model for almost all of these notions of distance. In this section, we will consider models of social networks that better match these non-geographic notions of similarity. Our discussion will include both non-grid-based models of social networks and ways to combine multiple notions of proximity into a single routing strategy.

Non-Grid-Based Measures of Similarity

Excluding geographic proximity to the target, similarity of occupation is the most-cited reason for the routing choices made by participants in Milgram-style routing experiments [13, 21]. Consider, then, modeling person-to-person proximity according to occupation. A hierarchical notion of similarity is natural in this context: imagine a tree T whose leaves correspond to particular occupations (“cannoli chef” or “urban planner,” perhaps), where each person u “lives” at the leaf ℓ_u that represents her occupation. The occupational proximity of u and v is given by the height of the least common ancestor (LCA) of ℓ_u and ℓ_v in T , which we will denote by $\text{lca}(u, v)$.

Hobbies can be modeled in a similar hierarchical fashion, though modeling hobby-based proximity is more complicated: a typical person has many hobbies but only one occupation. Measuring the similarity of two alma maters is more complicated still. There are many ways to measure the similarity of two schools, paralleling many ways to measure the similarity of two people: geography, “type of school” like

liberal arts college versus research university, athletic conference, strength of computer science department, etc. But even with these complications, similarity of any of occupation, hobbies, or alma mater is more naturally modeled with a hierarchy than with a grid.

Navigability in social networks derived from a hierarchical metric has been explored through analysis, through simulation, and through empirical study of real-world interactions. Kleinberg has shown a similar result to Theorem 1 for the tree-based setting, characterizing navigable networks in terms of a single parameter that controls how rapidly the link probability between people drops off with their distance [24]. As in the grid, Kleinberg’s theorem identifies an optimal middle ground in the tradeoff between having overly parochial and overly scattered connections: if T is a regular b -ary tree and $\Pr[u \rightarrow v] \propto b^{-\beta \cdot \text{lca}(u,v)}$, then the network is navigable if and only if $\beta = 1$. Watts, Dodds, and Newman [48] have explored a similar hierarchical setting, finding the ranges of parameters that were navigable in simulations. (Their focus was largely on the combination of multiple hierarchical measures of proximity, an issue to which we will turn shortly.) Routing in the hierarchical context has also been studied empirically by Adamic and Adar, who considered the role of proximity in the organizational structure of Hewlett–Packard Labs in social links among HP Labs employees [1]. (Because a company’s organizational structure forms a tree where people more senior in the organization are mapped to internal nodes instead of to leaves, Adamic and Adar consider a minor variation on LCA to measure person-to-person proximity.) Adamic and Adar found that, as with geography, there is a strong trace of organizational proximity in observed connections, and that, again as with geography, greedy routing towards a target based on organizational proximity was generally effective. (See Section 7 for some discussion.)

The question of navigability of a social network derived from an underlying measures of distance has also been explored beyond the contexts of the grid and the tree. Many papers have considered routing in networks in which person-to-person distances are measured by shortest-path distances in an underlying graph that has some special combinatorial structure. These papers then typically state bounds on navigability that are based on certain structural parameters of the underlying graph; examples include networks that have low treewidth [16], bounded growth rate [14, 15, 42], or low doubling dimension [19, 47]. The results on rank-based friendship, including generalizations and improvements on Theorem 2, have also been extended to the setting of low doubling dimension [6, 28]. However, a complete understanding of the generality of these navigability results in terms of properties of the underlying metric remains open.

Another way to model person-to-person proximity—and also to model variation in population density, in a different way from rank-based friendship—is the very general *group-structure* model of Kleinberg [24]. Each person in an n -person population is a member of various groups (perhaps defined by a shared physical neighborhood, an employer, a hobby), and $\Pr[u \rightarrow v]$ is a decreasing function of the size of the smallest group containing both u and v . Kleinberg proved that the resulting network is navigable if $\Pr[u \rightarrow v]$ is inversely proportional to the size of the smallest group including both u and v , subject to two conditions on the groups.

Informally, these conditions are the following. First, every group g must be “covered” by relatively large subgroups (so that once a path reaches g it can narrow in on a smaller group containing any particular target t). Second, groups must satisfy a sort of “bounded growth” condition (so that a person u has only a limited number of people who are in a group of a particular size with u , and thus u has a reasonable probability of “escaping” from small groups to reach a faraway target t).

Simultaneously Using Many Different Notions of Similarity

One of the major advantages of the group-structure model is that it allows us to model proximity between two people based on many “dimensions” of possible similarity, simply by defining some groups in terms of each of these multiple dimensions. Combining knowledge of various measures of the proximity—age, geography, and occupation, say—of one’s friends to the target is natural, and, indeed, something that real-world participants in small-world studies do [13, 21, 38]. Identifying plausible models for social networks and good routing algorithms to find short paths in these networks when there are many relevant notions of similarity remains an interesting and fertile area for research.

We have already implicitly considered one straightforward way of incorporating additional dimensions of similarity by modeling proximity in a k -dimensional grid for $k > 2$. (Even $k = 2$ uses two types of similarity—for geography, longitude and latitude—and computes person-to-person similarity by the combination of the two.) Because the grid-based model uses Manhattan distance, here the various dimensions of proximity are combined simply by summing their measured distances. Martel and Nguyen [35, 41], Fraigniaud, Gavoille, and Paul [18], and Barrière et al. [7] have performed further work in analyzing the grid-based setting for general k . These authors have shown that if people are given a small amount of additional information about the long-range links of their friends, then k -dimensional grids result in shorter paths as k increases. (From Theorem 1, we need to have $\Pr[u \rightarrow v] \propto d(u, v)^{-k}$ to achieve polylogarithmic path lengths; these results establish improvements in the polylogarithmic function as k increases, for a slightly souped-up version of local-information routing.)

Watts, Dodds, and Newman have explored a model of person-to-person similarity based on multiple hierarchies [48]. They consider a collection of k different hierarchies, where each is a regular b -ary tree in which the leaves correspond to small groups of people. The similarity of people u and v is given by their most similar similarity: that is, if $\text{lca}_i(u, v)$ denotes the height of the LCA of u and v in the i th hierarchy, then we model $d(u, v) := \min_i \text{lca}_i(u, v)$. People are mapped to each hierarchy so that a person’s position in each hierarchy is determined independently of positions in other hierarchies. Watts, Dodds, and Newman show via simulation that using $k > 1$ hierarchies yields better performance for greedy routing than using just one. In particular, using $k \in \{2, 3\}$ hierarchies gave the best performance. These experiments show that, for these values of k , the resulting network appears to be searchable for a broader range of parameters for the function giving friend-

ship probability as a function of distance. As with Theorem 1, there is provably a single exponent $\beta = 1$ under which greedy routing produces polylogarithmic paths when there is one hierarchy; for two or three hierarchies, these simulations showed a wider range of values of β that yield navigable networks.

The results in the Watts–Dodds–Newman setting are based on simulations, and giving a fully rigorous theoretical analysis of routing in this context remains an interesting open challenge. So too do a variety of generalizations of that setting: dependent hierarchies, or a combination of grid-based and hierarchy-based measures of proximity, or the incorporation of variable population density into the multiple-hierarchy setting. Broader modeling questions remain open, too. One can conceive of subtler ways of combining multiple dimensions of similarity than just the sum or the minimum that seem more realistic. For example, it seems that making significant progress towards a target in one dimension of similarity at the expense of large decreases in similarity in several other dimensions is a routing mistake, even if it reduces the minimum distance to the target over all the dimensions. Realistically modeling these multidimensional scenarios is an interesting open direction.

From Birds of a Feather to Social Butterflies (of a Feather)

The generalizations that we have discussed so far are still based on greedy routing under broader and more realistic notions of proximity, but one can also consider enriching the routing algorithm itself. For example, algorithms that endow individuals with additional “semi-local” information about the network, such as awareness of one’s friends’ friends, have also been studied (e.g., [18, 30, 34, 35, 47]). But there is another natural and simple consideration in Milgram-style routing that we have not mentioned thus far: some people have more friends than others. This is a significant omission of the models that we have discussed; people in these models have a constant or nearly constant number of friends. In contrast, degrees in real social networks are well modeled by a *power-law distribution*, in which the proportion of the population with f friends is approximately $1/f^\gamma$, where γ is a constant around 2.1 to 2.4 in real networks (see, e.g., [5, 8, 12, 29, 39]). In the routing context, a popular person can present a significant advantage in finding a shorter path to the target. A person with more friends has a higher probability of knowing someone who is significantly closer to any target—in virtue of having drawn more samples from the friendship distribution—and thus a more popular person will likely be able to find a shorter path to a given target.

Strategies that choose high-degree people in routing have been studied in a number of contexts, and, largely through simulation, these strategies have been shown to perform reasonably well [1–3, 22, 46]. Of these, perhaps the most promising algorithm for homophilous power-law networks is the *expected-value navigation (EVN)* algorithm of Şimşek and Jensen [46], which explicitly combines popularity and proximity in choosing the next step in a chain. Under EVN, the current node u chooses as the next node in the path its neighbor v whose probability of a direct link to the target is maximized. The node u computes this probability using the knowl-

edge of v 's proximity to t as well as v 's outdegree δ_v . (An underlying model like the grid, for example, describes the probability p_v that a particular one of v 's friendships will connect v to t ; one can then compute the probability $1 - (1 - p_v)^{\delta_v}$ that one of the δ_v friendships of v will connect v to t . EVN chooses the friend maximizing this probability as the next step in the chain.) Although Şimşek and Jensen give empirical evidence for EVN's success, no theoretical analysis has been performed. Analyzing this algorithm—or other similar algorithms that incorporate knowledge of node degree in addition to target proximity—in a rigorous setting is an important and open problem. Although a precise rigorous account of EVN has not yet been given, it is clear that EVN captures something crucial about real routing: the optimal routing strategy is some combination of getting close to a target in terms of similarity (the people who are more likely to know others most like the target) and of getting to popular intermediate people who have a large social circle (the people who are more likely to know many others in general). The interplay between popularity and proximity—and incorporating richer notions of proximity into that understanding—is a rich area for further research.

7 Discussion

It is clear that the wayfinding problem for real people in real social networks is only approximated by the models of social networks and of social-network routing discussed in this chapter. In many ways, real wayfinding is easier than it is in these models: we know which of our friends lived in Japan for a year, or tend to be politically conservative, or have a knack for knowing people in many walks of life, and we also have some intuitive sense of how to weight these considerations in navigating the network towards a particular target person. But real wayfinding is harder for real people in many ways, too: for example, even seemingly simple geography-based routing is, at best, a challenge for the third of college-age Americans who were unable to locate Louisiana on a map of the United States, even after the extensive press coverage of Hurricane Katrina [43].

The models of similarity and network knowledge that we have considered here are simplistic, and studying more realistic models—models with richer notions of proximity, or models of the errors or inconsistencies in individuals' mental maps of these notions of proximity, for example—is very interesting. But there is, of course, a danger of trying to model “too well”: the most useful models do not reproduce all of the fine-grained details of a real-world phenomenon, but rather shed light on that phenomenon through some simple and plausible explanation of its origin.

With this perspective in mind, I will highlight just one question here: why and how do social networks become navigable? A number of models of the evolution of social networks through the “rewiring” of long-range friendships in a grid-like setting have been defined and analyzed [10, 11, 44]; these authors have shown that navigability emerges in the network when this rewiring is done appropriately. We have seen here that rank-based friendship is another way to explain the navigability of so-

cial networks, and we have seen that friendships in LiveJournal, viewed geographically, are well approximated by rank-based friendship. One piece is missing from the rank-based explanation, though: why is it that rank-based friendship should hold in a real social network, even approximately? Figure 3 shows that geography plays a remarkably large role in friendships even in LiveJournal’s purely virtual community; friendship probability drops off smoothly and significantly as geographic proximity decreases. Furthermore Figure 5 shows that rank-based friendship is a remarkably accurate model of friendship in this network. But are there natural processes that can account for this behavior? Why should geographic proximity in the flesh-and-blood world resonate so much in the virtual world of LiveJournal? And why should this particular rank-based pattern hold?

One explanation for the important role of geography in LiveJournal is that a significant number of LiveJournal friendships are online manifestations of existing physical-world friendships, which crucially rely on geographic proximity for their formation. This “virtualization” is undoubtedly an important process by which friendships appear in a virtual community like LiveJournal, and it certainly explains some of geography’s key role. But accounting for the continued slow decay in link probability as geographic separation increases from a few hundred kilometers to a thousand kilometers, beyond the range of most spontaneous physical-world interactions, seems to require some additional explanation. Here is one speculative possibility: many interests held by LiveJournal users have natural “geographic centers”—for example, the city where a professional sports team plays, or the town where a band was formed, or the region where a particular cuisine is popular. Shared interests form the basis for many friendships. The geographic factor in LiveJournal could perhaps be explained by showing that the “mass” of u and v ’s shared interests (appropriately defined) decays smoothly as the geographic distance between u and v increases. Recent work of Backstrom et al. [4] gives some very intriguing evidence related to this idea. These authors have shown results on the geographic distribution of web users who issue various search queries. They characterize both the geographic “centers” of particular search queries and the “spread” of those queries, in terms of how quickly searchers’ interest in that query drops off with the geographic distance from the query’s center. Developing a comprehensive model of friendship formation on the basis of this underlying geographic nature of interests is a very interesting direction for future work.

To close, I will mention one interesting perspective on the question of an underlying mechanism by which rank-based friendship might arise in LiveJournal. This perspective comes from two other studies of node linking behavior as a function of node-to-node similarity, in two quite different contexts. Figure 6(b) shows the results of the study by Adamic and Adar [1] of the linking probability between HP Labs employees as a function of the distance between them in the corporate hierarchy. Their measure of similarity is a variant of LCA, modified to allow the calculation of distances to an internal node representing a manager in the corporate hierarchy. LCA distance is in a sense implicitly a logarithmic measure: for example, in a uniformly distributed population in the hierarchy, the number of people at distance d grows exponentially with d . Thus this semilog plot of link probabilities is on the

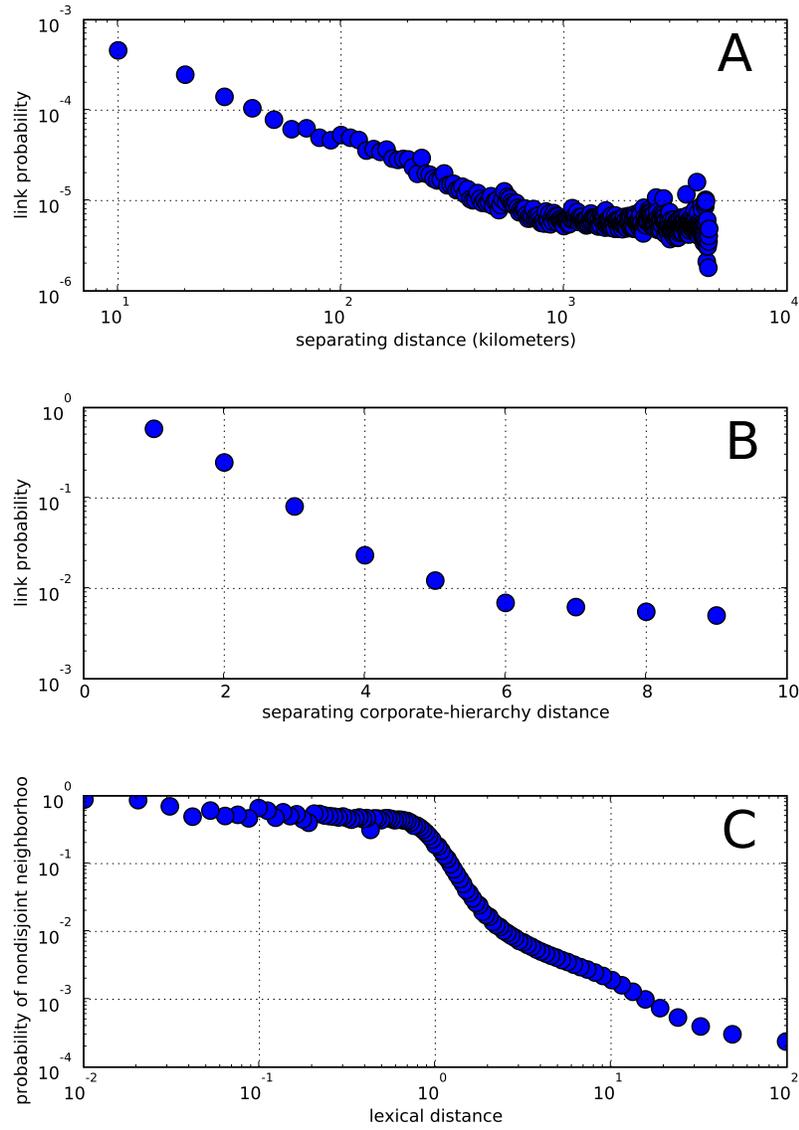


Fig. 6 Three plots of distance versus linking probability: (a) the role of geographic distance between LiveJournal users [32], a reproduction of Figure 3; (b) the role of corporate-hierarchy distance between HP Labs employees, from a study by Lada Adamic and Eytan Adar [1]; and (c) the role of lexical distance between pages on the web, from a study by Filippo Menczer [37].

same scale as the other log–log plots in Figure 6. Figure 6(c) shows the analogous plot from a study by Filippo Menczer [37] on the linking behavior between pages on the web. Here the similarity between two web pages is computed based on the lexical distance of the pages’ content. Because the raw link probabilities are so small, here the plot shows the probability that neighborhoods of two pages have nonempty overlap, where a page p ’s neighborhood consists of the page p itself, the pages to which p has a hyperlink, and pages that have a hyperlink to p .

Intriguingly, the LiveJournal linkage pattern, reproduced as Figure 6(a), and the HP Labs plot in Figure 6(b) show approximately the same characteristic shape in their logarithmic plots: a linear decay in link probability for comparatively similar people, leveling off to an approximately constant link probability for comparatively distant pairs. Figure 6(c) shows the opposite pattern: the probability of connection between two comparatively similar web pages is roughly constant, and then begins to decay linearly (in the log–log plot) once the pages’ similarity drops beyond a certain level. Figures 6(a) and 6(b) both plot link probability between people in a social network against their (geographic or corporate) distance; Figure 6(c) plots link probability for web pages. Understanding why linking patterns in social networks look different from the web—and, more generally, making sense of what might be generating these distributions—remains a fascinating open question.

Acknowledgements Thanks to Lada Adamic and Filippo Menczer for helpful discussions and for providing the data used to generate Figures 6(b) and 6(c). I would also like to thank the anonymous referees for their very helpful comments. This work was supported in part by NSF grant CCF-0728779 and by grants from Carleton College.

References

1. Lada A. Adamic and Eytan Adar. How to search a social network. *Social Networks*, 27(3):187–203, July 2005.
2. Lada A. Adamic, Rajan M. Lukose, and Bernardo A. Huberman. Local search in unstructured networks. In *Handbook of Graphs and Networks*. Wiley-VCH, 2002.
3. Lada A. Adamic, Rajan M. Lukose, Amit R. Puniyani, and Bernardo A. Huberman. Search in power-law networks. *Physical Review E*, 64(046135), 2001.
4. Lars Backstrom, Jon Kleinberg, Ravi Kumar, and Jasmine Novak. Spatial variation in search engine queries. In *Proceedings of the 17th International World Wide Web Conference (WWW’08)*, pages 357–366, April 2008.
5. Albert-László Barabási and Eric Bonabeau. Scale-free networks. *Scientific American*, 288:50–59, May 2003.
6. David Barbella, George Kachergis, David Liben-Nowell, Anna Sallstrom, and Ben Sowell. Depth of field and cautious-greedy routing in social networks. In *Proceedings of the 18th International Symposium on Algorithms and Computation (ISAAC’07)*, pages 574–586, December 2007.
7. Lali Barrière, Pierre Fraigniaud, Evangelos Kranakis, and Danny Krizanc. Efficient routing in networks with long range contacts. In *Proceedings of the 15th International Symposium on Distributed Computing (DISC’01)*, pages 270–284, October 2001.

8. Béla Bollobás, Oliver Riordan, Joel Spencer, and Gábor Tusnády. The degree sequence of a scale-free random graph process. *Random Structures and Algorithms*, 18(3):279–290, May 2001.
9. Dorwin Cartwright and Alvin Zander. *Group Dynamics: Research and Theory*. Row, Peterson, 1953.
10. Augustin Chaintreau, Pierre Fraigniaud, and Emmanuelle Lebhar. Networks become navigable as nodes move and forget. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP'08)*, pages 133–144, July 2008.
11. Aaron Clauset and Christopher Moore. How do networks become navigable? Manuscript, 2003. Available as `cond-mat/0309415`.
12. Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. Manuscript, 2007. Available as `arXiv:0706.1062`.
13. Peter Sheridan Dodds, Roby Muhamad, and Duncan J. Watts. An experimental study of search in global social networks. *Science*, 301:827–829, 8 August 2003.
14. Philippe Duchon, Nicolas Hanusse, Emmanuelle Lebhar, and Nicolas Schabanel. Could any graph be turned into a small world? *Theoretical Computer Science*, 355(1):96–103, 2006.
15. Philippe Duchon, Nicolas Hanusse, Emmanuelle Lebhar, and Nicolas Schabanel. Towards small world emergence. In *Proceedings of the 18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'06)*, pages 225–232, August 2006.
16. Pierre Fraigniaud. Greedy routing in tree-decomposed graphs. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA'05)*, pages 791–802, October 2005.
17. Pierre Fraigniaud. Small worlds as navigable augmented networks: Model, analysis, and validation. In *Proceedings of the 15th Annual European Symposium on Algorithms (ESA'07)*, pages 2–11, October 2007.
18. Pierre Fraigniaud, Cyril Gavoille, and Christophe Paul. Eclecticism shrinks even small worlds. In *Proceedings of the 23rd Symposium on Principles of Distributed Computing (PODC'04)*, pages 169–178, July 2004.
19. Pierre Fraigniaud, Emmanuelle Lebhar, and Zvi Lotker. A doubling dimension threshold $\theta(\log \log n)$ for augmented graph navigability. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA'06)*, pages 376–386, September 2006.
20. Frank Harary and Robert Z. Norman. *Graph Theory as a Mathematical Model in Social Science*. University of Michigan, 1953.
21. P. Killworth and H. Bernard. Reverse small world experiment. *Social Networks*, 1:159–192, 1978.
22. B. J. Kim, C. N. Yoon, S. K. Han, and H. Jeong. Path finding strategies in scale-free networks. *Physical Review E*, 65(027103), 2002.
23. Jon Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the 32nd Annual Symposium on the Theory of Computation (STOC'00)*, pages 163–170, May 2000.
24. Jon Kleinberg. Small-world phenomena and the dynamics of information. In *Advances in Neural Information Processing Systems (NIPS'01)*, pages 431–438, December 2001.
25. Jon Kleinberg. Complex networks and decentralized search algorithms. In *International Congress of Mathematicians (ICM'06)*, August 2006.
26. Jon M. Kleinberg. Navigation in a small world. *Nature*, 406:845, 24 August 2000.
27. Judith Kleinfeld. Could it be a big world after all? The “six degrees of separation” myth. *Society*, 39(61), April 2002.
28. Ravi Kumar, David Liben-Nowell, and Andrew Tomkins. Navigating low-dimensional and hierarchical population networks. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA'06)*, pages 480–491, September 2006.
29. Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, D. Sivakumar, Andrew Tomkins, and Eli Upfal. Stochastic models for the web graph. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS'00)*, pages 57–65, November 2000.
30. Emmanuelle Lebhar and Nicolas Schabanel. Close to optimal decentralized routing in long-range contact networks. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming (ICALP'04)*, pages 894–905, July 2004.

31. Kurt Lewin. *Principles of Topological Psychology*. McGraw Hill, 1936.
32. David Liben-Nowell, Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Geographic routing in social networks. *Proceedings of the National Academy of Sciences*, 102(33):11623–11628, August 2005.
33. Kevin Lynch. *The Image of the City*. MIT Press, 1960.
34. Gurmeet Singh Manku, Moni Naor, and Udi Wieder. Know thy neighbor’s neighbor: the power of lookahead in randomized P2P networks. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC’04)*, pages 54–63, June 2004.
35. Chip Martel and Van Nguyen. Analyzing Kleinberg’s (and other) small-world models. In *Proceedings of the 23rd Symposium on Principles of Distributed Computing (PODC’04)*, pages 179–188, July 2004.
36. Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, August 2001.
37. Filippo Menczer. Growing and navigating the small world web by local content. *Proceedings of the National Academy of Sciences*, 99(22):14014–14019, October 2002.
38. Stanley Milgram. The small world problem. *Psychology Today*, 1:61–67, May 1967.
39. Michael Mitzenmacher. A brief history of lognormal and power law distributions. *Internet Mathematics*, 1(2):226–251, 2004.
40. Jacob L. Moreno. *Who Shall Survive? Foundations of Sociometry, Group Psychotherapy and Sociodrama*. Nervous and Mental Disease Publishing Company, 1934.
41. Van Nguyen and Chip Martel. Analyzing and characterizing small-world graphs. In *Proceedings of the 16th ACM–SIAM Symposium on Discrete Algorithms (SODA’05)*, pages 311–320, January 2005.
42. Van Nguyen and Chip Martel. Augmented graph models for small-world analysis with geographical factors. In *Proceedings of the 4th Workshop on Analytic Algorithms and Combinatorics (ANALCO’08)*, January 2008.
43. Roper Public Affairs and National Geographic Society. 2006 geographic literacy study, May 2006. <http://www.nationalgeographic.com/roper2006>.
44. Oskar Sandberg and Ian Clarke. The evolution of navigable small-world networks. Manuscript, 2006. Available as `cs/0607025`.
45. Georg Simmel. *Conflict And The Web Of Group Affiliations*. Free Press, 1908. Translated by Kurt H. Wolff and Reinhard Bendix (1955).
46. Özgür Şimşek and David Jensen. Decentralized search in networks using homophily and degree disparity. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI’05)*, pages 304–310, August 2005.
47. Aleksandrs Slivkins. Distance estimation and object location via rings of neighbors. In *Proceedings of the 24th Symposium on Principles of Distributed Computing (PODC’05)*, pages 41–50, July 2005.
48. Duncan J. Watts, Peter Sheridan Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296:1302–1305, 17 May 2002.
49. Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.